Measuring inconsistency in probabilistic knowledge bases

Glauber De Bona

Thesis submitted
to the
Institute of Mathematics and Statistics
of the
University of São Paulo
for
obtention of the title
of
Doctor of Sciences

Program: Computer Science

Supervisor: Prof. Dr. Marcelo Finger

During the development of this work, the author was financially supported by CAPES

São Paulo, March 2016

Measuring inconsistency in probabilistic knowledge bases

This version of the thesis contains the corrections and changes suggested by the Committee Members during the work's original version defence, which occurred on 22/1/2016. A copy of the original version is available at the Insitute of Mathematics and Statistics of the University of São Paulo.

Committee Members:

- Prof. Dr. Marcelo Finger (supervisor) - IME-USP
- Prof. Dr. Renata Wasserman - IME-USP
- Prof. Dr. Fábio Gagliardi Cozman - POLI-USP
- Prof. Dr. Walter Alexandre Carnielli - Unicamp
- Prof. Dr. João Marcos de Almeida - UFRN

Acknowledgements

In first place, I would like to thank my wife, Gleyce. Her admirable affection and kindness look after me almost all the time, and her complicity keeps me passionate about life; and, as Pascal said, nothing great is done without passion.

If I am here, it is because a couple had overflowing love to bring me to existence and tireless dedication to raise me. Effort which was always focused on education, either at home, or in schools that did not bound my intellectual development. It is ultimately impossible to repay this debt of an earned life, thus I can only offer my eternal gratitude to my parents, Neivor and Neiva.

I thank my academic father, Professor Marcelo Finger, who has supervised me since my master's degree pursue, in an extremely healthy relationship. Marcelo has always been willing to help me in anything I needed, has trusted in my work and pushed me forward. His experience and intuition have complemented me, in a partnership with plenty of results, conjectures and refutations. Besides that, he believed in my productivity in spite of my particular characteristics, giving me freedom to work from home and to live months at the beach. Whether the formal supervision terminates here, that our contribution last for decades!

Other people have technically collaborated with my work, not only with ideas and suggestions, but also with co-authorships on results and publications. Especially, I would like to thank Professor Fábio Gagliardi Cozman, who has always contributed in my oral examinations, and with whom I have had the pleasure of working together more than once. I also thank Professor Márcio Moretto Ribeiro, who, besides suggesting ideas, has been helping in their development, together with Yuri David Santos and Professor Renata Wasserman. Renata also was in my final *viva voice* examination, with Professor João Marcos de Almeida and Professor Walter Alexandre Carnielli, to whom I'm also thankful. I would like to thank Eduardo Menezes de Morais for using techniques here elaborated to solve real-world problems. I thank Julia Staffel for the help in applying to Formal Epistemology some ideas investigated in this study. I thank yet Carlos Stein Naves de Brito for the mathematical methods suggested.

Many are the people with whom I share part of my life, and I am thankful to all the friends that turn existence into a lighter burden. Namely, I would like to thank Abelardo, Alexandredalus (or Zaratesta) and Tutu, my cousin Preto and my sister, Aracelli. I would also like to thank my spouse's family, Roberta, Renan and Ylana, who has embraced me like its part.

Not rarely, theists start or end this section with acknowledgements to a supposed god. As a Bayesian agent, I have amassed evidences that neither confirm nor disconfirm the hypothesis of such entity's existence, and my atheism (in spite of my first name) is determined by a probability that is inversely proportional to the complexity of such hypothesis — or by Ockham's razor. Thus, I do not thank god, but I do thank those whose texts opened my eyes to the critical, rational thought. In awaking me to philosophy, great authors made me see the world in a different way, when I was

reborn to search in the academic career a path towards the truth. Hence, I would probably have never obtained a doctor's degree in this life if Socrates and Plato had not taken me out the naive reality with their cave myth, Descartes had not rescued me from the absolute skepticism through his *cogito* argument, Kant had not told the cognoscible phenomenon from the thing-in-itself for me, Schopenhauer had not shown me that there is will beyond representation and Nietzche had not warned me that god is dead. Thanks for saving my life!

Resumo

DE BONA, G. **Medindo inconsistência em bases de conhecimento probabilístico**. 2016. 152 f. Tese (Doutorado) - Instituto de Matemática e Estatística, Universidade de São Paulo, São Paulo, 2016.

Em termos de raciocínio probabilístico clássico, para se realizar inferências de uma base de conhecimento, normalmente é necessário garantir a consistência de tal base. Quando nos deparamos com um conjunto de probabilidades que são inconsistentes entre si, interessa-nos saber onde está a inconsistência, quão grave esta é, e como corrigi-la. Medidas de inconsistência têm sido recentemente propostas como uma ferramenta para endereçar essas questões na comunidade de Inteligência Artificial. Este trabalho investiga o problema da medição de inconsistência em bases de conhecimento probabilístico.

Postulados básicos de racionalidade têm guiado a formulação de medidas de inconsistência na lógica clássica proposicional. No caso probabilístico, o carácter quantitativo da probabilidade levou a uma propriedade desejável adicional: medidas de inconsistência devem ser contínuas. Para atender a essa exigência, a inconsistência em bases de conhecimento probabilístico tem sido medida através da minimização de distâncias. Nesta tese, demonstramos que o postulado da continuidade é incompatível com propriedades desejáveis herdadas da lógica clássica. Como algumas dessas propriedades são baseadas em conjuntos inconsistentes minimais, nós procuramos por maneiras mais adequadas de localizar a inconsistência em lógica probabilística, analisando os processos de consolidação subjacentes. A teoria AGM de revisão de crenças é estendida para englobar a consolidação pelo ajuste de probabilidades. As novas formas de caracterizar a inconsistência que propomos são empregadas para enfraquecer alguns postulados, restaurando a compatibilidade de todo o conjunto de propriedades desejáveis.

Investigações em estatística Bayesiana e em epistemologia formal têm se interessado pela medição do grau de incoerência de um agente. Nesses campos, probabilidades são geralmente interpretadas como graus de crença de um agente, determinando seu comportamento em apostas. Agentes incoerentes possuem graus de crença inconsistentes, que o expõem a transações de apostas desvantajosas — conhecidas como *Dutch books*. Estatísticos e filósofos sugerem medir a incoerência de um agente através do prejuízo garantido a qual ele está vulnerável. Nós provamos que estas medidas de incoerência via *Dutch books* são equivalentes a medidas de inconsistência via minimização de distâncias da comunidade de IA.

Palavras-chave: raciocínio probabilístico, lógica probabilística, medidas de inconsistência.

Abstract

DE BONA, G. **Measuring inconsistency in probabilistic knowledge bases**. 2016. 152 f. Tese (Doutorado) - Instituto de Matemática e Estatística, Universidade de São Paulo, São Paulo, 2016.

In terms of standard probabilistic reasoning, in order to perform inference from a knowledge base, it is normally necessary to guarantee the consistency of such base. When we come across an inconsistent set of probabilistic assessments, it interests us to know where the inconsistency is, how severe it is, and how to correct it. Inconsistency measures have recently been put forward as a tool to address these issues in the Artificial Intelligence community. This work investigates the problem of measuring inconsistency in probabilistic knowledge bases.

Basic rationality postulates have driven the formulation of inconsistency measures within classical propositional logic. In the probabilistic case, the quantitative character of probabilities yielded an extra desirable property: that inconsistency measures should be continuous. To attend this requirement, inconsistency in probabilistic knowledge bases have been measured via distance minimisation. In this thesis, we prove that the continuity postulate is incompatible with basic desirable properties inherited from classical logic. Since minimal inconsistent sets are the basis for some desiderata, we look for more suitable ways of localising the inconsistency in probabilistic logic, while we analyse the underlying consolidation processes. The AGM theory of belief revision is extended to encompass consolidation via probabilities adjustment. The new forms of characterising the inconsistency we propose are employed to weaken some postulates, restoring the compatibility of the whole set of desirable properties.

Investigations in Bayesian statistics and formal epistemology have been interested in measuring an agent's degree of incoherence. In these fields, probabilities are usually construed as an agent's degrees of belief, determining her gambling behaviour. Incoherent agents hold inconsistent degrees of beliefs, which expose them to disadvantageous bet transactions — also known as Dutch books. Statisticians and philosophers suggest measuring an agent's incoherence through the guaranteed loss she is vulnerable to. We prove that these incoherence measures via Dutch book are equivalent to inconsistency measures via distance minimisation from the AI community.
Keywords: probabilistic reasoning, probabilistic logic, inconsistency measures.

Contents

List of Abbreviations **ix**

List of Symbols **xi**

List of Tables **xiii**

1 Introduction **1**
1.1 Objectives . . . 4
1.2 Contributions . . . 4
1.3 Bird's-Eye View of the Thesis . . . 5
1.4 List of Publications . . . 6

2 Preliminaries **9**
2.1 Classical Propositional Logic . . . 9
2.2 Propositional Probabilistic Logic . . . 10
2.2.1 Syntax and Semantics . . . 11
2.2.2 Deciding the Consistency of Probabilistic Bases . . . 12
2.2.3 Alternative Semantics for Probabilistic Consistency . . . 15

3 Inconsistency Measures in Probabilistic Logic and Rationality Postulates **17**
3.1 General Inconsistency Measures and Postulates . . . 17
3.2 Inconsistency Measures Tailored to the Probabilistic Logic . . . 21
3.2.1 Measuring Inconsistency as Distance between Probability Bounds . . . 21
3.2.2 Measuring Inconsistency as the Distance to being Satisfied . . . 25
3.3 Related Work . . . 28
3.4 The Incompatibility of Postulates . . . 30

4 Consolidating Probabilistic Knowledge Bases **33**
4.1 Consolidation Operators and Strategies . . . 34
4.2 Removing Formulas . . . 35
4.3 Changing Probability Bounds . . . 38
4.3.1 Minimising Distances between Probabilities . . . 39
4.3.2 Minimising a Discrepancy to a Probabilistic Interpretation . . . 41
4.4 Adding Violations . . . 44
4.5 Related Work . . . 47
4.6 Back to the Postulates Incompatibility . . . 50

5 A Brief Review on Belief Revision: Quantitative Consolidation within the AGM Paradigm **51**
5.1 Consolidation and Contraction in AGM Theory . . . 52
5.2 Changing Probability Bounds versus Discarding Formulas . . . 55
5.3 Liftable Contractions . . . 58
5.4 Probabilistic Contraction . . . 62
5.5 AGM Paradigm and ε-Consolidations . . . 65

6 Localising Inconsistency and Fixing the Postulates **67**
6.1 Minimal Inconsistent Sets and Free Conditionals . . . 68
6.2 Refining the Free Conditional Concept . . . 70
6.3 Refining the Primitive Conflict Concept . . . 72
6.4 Primitive Conflicts and ε-Consolidations . . . 74
6.5 Reconciling the Postulates . . . 76

7 Inconsistency Measures via Dutch Books **79**
7.1 Introducing Dutch Books . . . 79
7.2 Minimal Violation Measures as Guaranteed Sure Losses . . . 83
7.3 Other Feasible Principled Measures . . . 85
7.4 Measuring Inconsistency with Confidence Factors . . . 87

8 Conclusion **95**
8.1 Summary . . . 95
8.2 Concluding Remarks . . . 96
8.3 Future Work . . . 97

A Proofs of Technical Results **99**
A.1 Technical Results from Chapter 3 . . . 99
A.2 Technical Results from Chapter 4 . . . 103
A.3 Technical Results from Chapter 5 . . . 107
A.4 Technical Results from Chapter 6 . . . 119
A.5 Technical Results from Chapter 7 . . . 126

Bibliography **131**

List of Abbreviations

AGM	Alchourrón, Gärdenfors e Makinson
CONDSAT	PSAT with conditional probabilities.
GPSAT	Generalised PSAT.
iff	"if, and only if".
IC	Inescapable Conflict
MAXSAT	Maximum Satisfiability Problem.
MIS	Minimal Inconsistent (Sub)set.
NP	A Computational Complexity Class, see for instance (Garey and Johnson, 1979).
OPSAT	PSAT optimisation version.
P	A Computational Complexity Class, see for instance (Garey and Johnson, 1979).
PSAT	Probabilistic Satisfiability Problem.
PSAT-IP	PSAT with imprecise probabilities.
SAT	Satisfiability problem for Classical Propositional Logic.

List of Symbols

$\wedge$	Classical logical conjunction.
$\vee$	Classical logical disjunction.
$\neg$	Classical logical negation.
$\rightarrow$	Classical logical implication.
$\leftrightarrow$	Classical logical bi-implication.
$\models$	Classical logical entailment.
$\top$	Tautology.
$\perp$	Contradiction.
$\mathcal{L}_{X_n}$	Propositional logic over X_n.
$\mathcal{L}^P_{X_n}$	Probabilistic language over $\mathcal{L}_{X_n}$.
$\bar{\mathcal{L}}^P_{X_n}$	The set $\{\alpha \vert \neg\alpha \in \mathcal{L}^P_{X_n}\}$.
Π_n	The set of all probability masses $\pi : W_n \rightarrow [0, 1]$.
$\mathbb{K}$	The set of all probabilistic knowledge bases.
$\mathbb{K}_c$	The set of all canonical probabilistic knowledge bases.
$\Lambda_\Gamma(.)$	Function that changes the probability lower bounds of Γ.
$\Lambda^P_\Gamma(.)$	Like $\Lambda_\Gamma(.)$, but resulting in precise probabilities.
$f \circ g$	f composed with g.
$\Vert.\Vert_p$	The p-norm of a vector.
$\mathbb{N}_{>0}$	Natural numbers set excluding 0.
$\overline{\mathbb{N}}_{>0}$	$\mathbb{N}_{>0} \cup \{\infty\}$.
$-$	A contraction operation.
!	A consolidation operation.
$\Gamma \perp \alpha$	Collection of maximal subsets of Γ not implying α.
$\Gamma \perp\!\!\!\perp \alpha$	Collection of minimal subsets of Γ implying α.

List of Tables

7.1 Prizes earned by Alice (in $) for each bet and each possible world. 81
7.2 Prizes earned by Alice (in $) for each bet and each possible world. 81

Chapter 1

Introduction

"when you can measure what you are speaking about, you know something about it; but when you cannot [...] your knowledge is of a meagre and unsatisfactory kind;"

— Lord Kelvin (Thomson, 1891)

Representing real-world knowledge and performing inference usually demand formalisms that can cope with uncertainty. Probabilistic logics combine the deductive power of logical systems with the well-founded Theory of Probability to attend to this need. Boole (1854) had already assigned probabilities to logical formulas, analytically deriving the probability bounds a set of premises entail for a given sentence. de Finetti (1930, 1974) employed logic to develop his theory of subjective probability, which yielded a probabilistic logic based on the concept of coherence (Coletti and Scozzafava, 2002). Hailperin (1965, 1986) revisited the work of Boole, using a linear algebra formulation to perform probabilistic reasoning. Nilsson (1986) presented the topic to the Artificial Intelligence (AI) community, introducing the problems of probabilistic satisfiability (PSAT) and entailment. Kavvadias and Papadimitriou (1990) and Hansen and Jaumard (2000) made important contributions to these problems on the computational aspects, providing practicable algorithms. Halpern (2003) and Ognjanovic, Raškovic, and Markovic (2009) formalised different probabilistic logics, providing them with proof systems and investigating their completeness. Kern-Isberner and Lukasiewicz (2004) and Haenni, Romeijn, Wheeler, and Williamson (2011) among others, studied entailment in probabilistic logic based on different semantics, for instance employing the maximum entropy principle.

All approaches mentioned above rely on the consistency of the set of premises in order to perform probabilistic inference, which is a standard practice throughout logical reasoning. Nonetheless, many are the possible sources of inconsistency in a probabilistic knowledge base: it may contain statistical data from different samples, it could have been formed by the opinion of different experts, or even a single expert could lack the resources to check his own consistency while building the base, etc. To restore consistency in such cases, the inconsistency may be analysed, which calls for a way to measure it. This work mainly investigates measures of inconsistency for knowledge bases over probabilistic logic, whose consistency checking (PSAT) was the focus of the author's master's dissertation.

The problem of measuring inconsistency in knowledge bases[1] over logical languages has increasingly received attention during the last years. As a relatively recent research branch, most proposals

[1] Classically, any inconsistent (logically closed) theory contains all formulas in the language, so that inconsistency measuring focus on (finite) bases.

that address the question of how to gauge the inconsistency severity in knowledge bases have focused on classical propositional logic. Knight (2002) put forward a method to measure inconsistency by attaching probabilities to formulas. Hunter and Konieczny (2006) combined measures based on how many formulas are required to produce a contradiction with measures based on the proportion of the language affected by it. A different approach by Hunter and Konieczny (2008) is entirely grounded in minimal inconsistent sets and distributes the inconsistency "culpability" among the formulas. Grant and Hunter (2013) quantified inconsistency by employing distances between models (valuations, or truth assignments). While developing these measures, Hunter and Konieczny (2006, 2008, 2010) have laid out a set of basic properties an inconsistency measure for bases should hold, which has become standard rationality postulates.

In AI, one of the main uses of measuring inconsistency in a knowledge base is to guide the consolidation of inconsistent pieces of information. For instance, within propositional logic, Grant and Hunter (2011) showed how inconsistency measures can be used to direct the stepwise resolution of conflicts via the weakening or the discarding of formulas.

In probabilistic bases, inconsistencies are more common, especially when knowledge is gathered from different sources. To fix these probabilistic knowledge bases, one can, for instance, delete pieces of information, or change the probabilities' numeric values (or bounds). In this case, an inconsistency measure helps one to detect whether or not a change approximates consistency. In other areas, inconsistency measures for probabilistic logic have found applications in merging conflicting opinions, leading to an increased predictive power (Karvetski *et al.*, 2013; Wang *et al.*, 2011), and in quantifying the incoherence of procedures from classical statistical hypothesis testing (Schervish *et al.*, 2002a).

Example 1.0.1. Consider we are devising an expert system to assist medical diagnosis. Suppose a group of experts on a disease D is required to quantify the relationship between D and its symptoms. Suppose three conditional probabilities (among others) are presented:

- the probability of a patient exhibiting symptom S_1 given he/she has disease D is 50%;
- the probability of a patient exhibiting symptom S_2 given he/she exhibits symptom S_1 and has disease D is 80%;
- the probability of a patient exhibiting symptom S_2 given he/she has disease D is 30%.

A knowledge engineer, while checking those facts, finds that they are inconsistent: according to the first two items, the probability of symptom S_2, given disease D, should be at least $50\% \times 80\% = 40\%$, instead of 30%. He does not even know where each probability came from, but plans to adjust them, since consistency is a requirement. How should he proceed? Which probabilities is the degree of inconsistency most sensitive to? Once chosen which number to change, should it be raised or lowered in order to approximate consistency? These are the kind of questions an inconsistency measure can help to answer. □

The problem of measuring inconsistency in probabilistic knowledge bases has more recently been tackled by Thimm (2013), Muiño (2011) and Potyka (2014). All three authors developed measures based on distance minimisation, tailored to the probabilistic case. Muiño (2011) was driven by the real knowledge base CADIAG-2 (Adlassnig and Kolarz, 1986), showing its infinitesimal inconsistency degree. Thimm (2013) adapted Hunter and Konieczny's desirable properties for inconsistency

measuring, searching for measures that satisfy a set of postulates. Potyka (2014) focused on computational aspects, looking for efficiently computable measures.

An inconsistency measure, for a general logic, is a function taking knowledge bases to non-negative numbers, which must obey some postulates. The first one, introduced by Hunter and Konieczny (2006) for classical logic, is *(Consistency)*, which claims that an inconsistency measurement is zero if, and only if, the corresponding base is consistent. Another desirable property suggested by Hunter and Konieczny (2006) is *(Independence)*, stating that the withdrawal of a free formula of the base — *i.e.*, a formula that does not belong to any minimal inconsistent set — should not change the inconsistency measurement. Thimm (2013) brought these postulates, among others, to the probabilistic context, adding *(Continuity)* to the list, which intuitively says that small changes in probabilities lead to small changes at the value of the inconsistency measure. We prove that (Consistency), (Independence) and (Continuity) cannot hold together, and some of these postulates must be abandoned or exchanged for jointly satisfiable ones. Since (Independence) is based on minimal inconsistent sets, a problem related to the postulates reconciliation is how to characterise "atomic" inconsistencies in probabilistic logic.

Example 1.0.2. Consider again a situation similar to that from Example 1.0.1, in which a medical expert system is being designed. Suppose a group of experts on a disease D is required to quantify the relationship between D and its symptoms. Suppose three conditional probabilities (among others) are presented:

- the probability of a patient exhibiting both symptom S_1 and symptom S_2 given he/she has disease D is at least 60%;
- the probability of a patient exhibiting symptom S_1 but not symptom S_2 given he/she exhibits has disease D is at least 50%;
- the probability of a patient exhibiting symptom S_1 given he/she has disease D is at most 80%.

Once more, a knowledge engineer finds that these pieces of information are inconsistent: according to the first two items, the probability of symptom S_1, given disease D, should be at least 110%. Instead of adjusting these probabilities by himself, this knowledge engineer decides that the experts who elicit them should do the job. It happens that each of these probabilistic assessments has come from a different expert. The engineer then intends to schedule a meeting among the physicians responsible for the inconsistency in order for them to reassess their assignments in a consistent way. These experts are very busy, and their time, expensive, thus the knowledge engineer wants to invite for the meeting only the very physicians whose probabilistic assessments are collaborating, or causing, the inconsistency of the whole base. Who should the engineer call? Or, which pieces of information can be "blamed" for causing the inconsistency? It seems clear that the doctors responsible for the two first statements should be invited, but what about the third? These questions can be answered by *localising* the inconsistency in the base. □

While computer scientists are investigating the problem of measuring inconsistency in probabilistic knowledge bases, statisticians and philosophers have been interested in evaluating the incoherence of formal agents that assign probability to events or propositions. In Bayesian epistemology, probabilities are usually construed as an agent's degrees of belief, which can be operationally

defined as the relative prices she would be willing to pay for gambles. Under these assumptions, Schervish, Seidenfeld, and Kadane (2002b) proposed to measure the incoherence of an agent by quantifying the maximum sure loss she would be exposed to via a disadvantageous set of bets — a *Dutch book*. Staffel (2015) analysed these measures of incoherence based on Dutch books, formulating a rationality principle they should obey and arguing in favour of a specific measure. To the best of our knowledge, these proposals for measuring incoherence of Bayesian agents have been ignored within the AI community, even though they correspond to measures on probabilistic knowledge bases.

1.1 Objectives

The first main objective of this thesis is to reconcile the rationality postulates for measuring the inconsistency of probabilistic knowledge bases. To achieve this, the characterisation of problematic sets of formulas in a probabilistic knowledge base — those causing the inconsistency — must be analysed, since these sets are the basis of some incompatible desirable properties. That is, before fixing the postulates, we have to understand how to localise inconsistency in probabilistic logic.

As a second major aim, this work intends to link different approaches from different communities to very similar problems: on the one hand, distance-based methods to measure the inconsistency of probabilistic knowledge bases in Artificial Intelligence; on the other hand, sure losses via Dutch books to quantify the incoherence of agents in Bayesian statistics and formal epistemology.

1.2 Contributions

Identifying and fixing the incompatibility of the rationality postulates for inconsistency measures in probabilistic logic is the first major contribution of this work. We prove that the desirable properties are not jointly satisfiable and provide suitable weakenings for two of them. A derived achievement of this thesis is the proposal of alternative forms of localising the inconsistency in probabilistic bases, which are employed to reconcile the postulates. In this development, we point out the close link between the tasks of localising and repairing inconsistency.

These new characterisations of conflicting formulas resulted from the analysis of probabilistic consolidation methods. In classical logic, consolidation is grounded in the standard AGM paradigm of belief revision (Alchourrón *et al.*, 1985), but its qualitative character fails to capture consolidation via probabilities adjustment. An additional contribution of the present work lies in extending the AGM paradigm to allow for probabilistic consolidation via adjusting probability bounds. As a consequence, the scope of probabilistic contraction is defined, together with a method for performing it.

In the third contribution of this thesis, we prove a formal equivalence between inconsistency measures via distance minimisation and incoherence measure via Dutch books. This provides the former with a meaningful interpretation, through the betting behaviour induced by the probabilities. Furthermore, our results open the path for the reciprocal interchange of ideas and techniques between the corresponding communities, as we show they are tackling the same formal problem.

1.3 Bird's-Eye View of the Thesis

After introducing probabilistic knowledge bases and fixing notation in Chapter 2, this work develops its three main contributions in different chapters. In the following, we present the organisation of the thesis, together with the contextualisation of its contributions within the existing literature.

Inconsistency measures for probabilistic knowledge bases are presented in Chapter 3. We start with the adaptation of inconsistency measures initially proposed for classical logic and their rationality postulates, as (Consistency) and (Independence). Another desirable property to be proved problematic is (MIS-separability), which deals with decomposability — through Minimal Inconsistent Sets — and implies (Independence). We proceed by investigating measures tailored to the probabilistic case, adding the postulate of (Continuity), proposed by Thimm (2013). We review the measures proposed by Thimm (2013) (and Muiño (2011)), Potyka (2014) and Capotorti, Regoli, and Vattari (2010), although the latter adopt a different semantics, while we adapt them to the more general setting of imprecise probabilities. We point out that two measures from Potyka (2014) are practicable, being computable through linear programs. The final section brings part of this work's first contribution: the presented postulates are shown to be incompatible. Afterwards, we suggest that (Independence) might be the requirement to be weakened, together with the stronger property of (MIS-separability), and the following chapter searches for their reasons.

The postulates of (Independence) and (MIS-Separability) depends on the concepts of free formula and minimal inconsistent set, respectively. In Chapter 4, we investigate methods to restore consistency (to *consolidate*) in order to show that these concepts are linked to classical consolidation and contraction — *i.e.*, discarding formulas to reach consistency. However, there are other methods to consolidate a probabilistic base, and the rest of the chapter is devoted to them. For instance, a natural consolidation procedure for probabilistic bases is changing the probabilities, or probability bounds, instead of ruling them out. In particular, we analyse the consolidation procedures that somehow reflect the continuous inconsistency measures tailored to probabilistic bases shown in Chapter 3. To investigate if free formulas and minimal inconsistent sets are also linked to these intrinsically probabilistic consolidation procedures is the task of the Chapter 6.

A detour is taken in Chapter 5, discussing consolidation via probabilities changing and the AGM theory — named after Alchourrón, Gärdenfors, and Makinson (1985). A consolidation in such paradigm, as a contraction by the contradiction, should be a subset of the knowledge (or belief) base, thus precluding adjustment in the probability bounds. Hansson (1999) defines pseudo-contractions as containing formulas implied by the original base, allowing for decreasing lower (or increasing upper) probability bounds. Nevertheless, under the standard semantics, the whole logical language is entailed by an inconsistent knowledge base, and any consistent base could be its pseudo-contraction by the contradiction. A possible direction to circumvent this issue was recently suggested by Santos, Ribeiro, and Wassermann (2015), by employing weaker, subclassical consequence operations to build the pseudo-contractions. We build on their ideas to introduce *liftable contractions*, in order that the consolidation methods seen in Chapter 4 may be encompassed in the AGM paradigm — which is our second contribution. Probabilistic contraction is investigated and we show that it is not well-behaved, unless the formula being contracted is in a specific language.

Chapter 6 begins by showing how free formulas and minimal inconsistent sets do no fit the consolidation procedures related to continuous measures. Transposing to these consolidation meth-

ods the link found between free formulas and classical consolidation, we introduce *innocuous* and *ε-innocuous* conditionals. In a similar way, we define *inescapable conflict* and *ε-inescapable conflict* to characterise primitive conflicts in probabilistic knowledge bases. Grounded in these new concepts, we propose weakened versions of (Independence) and (MIS-Separability). At the end of the Chapter, the first contribution of this thesis is completed when the new postulates are shown to be compatible with (Consistency) and (Continuity).

In Chapter 7, we investigate the relation between the inconsistency measures seen and incoherence measures based on Dutch books. We review two measures from Schervish *et al.* (2002b) that quantify the incoherence of an agent using the betting concept of Dutch book, under the operational interpretation that the agent's degrees of belief (probabilities) determine her gambling behaviour. The third main contribution of this thesis lies in showing the connection between inconsistency measures for probabilistic knowledge bases and incoherence measures for agents from formal epistemology. We prove that the two measures adapted from (Potyka, 2014) that are computable through linear programs are equivalent to two measures from Schervish *et al.* — that is, these two measures are rather efficiently computable and have a meaningful interpretation. In a second moment, other measures from the work of Schervish *et al.* are shown to satisfy the postulates and shown to be computationally feasible. Finally, we propose a general framework to measure inconsistency with confidence factors, encompassing measures from both communities that satisfy the reconciled postulates.

1.4 List of Publications

Part of this thesis' results, as well as related research done by the author during his doctoral studies, has appeared in workshops, conferences and journals:

De Bona and Finger(2015) G. De Bona and M. Finger. Measuring inconsistency in probabilistic logic: rationality postulates and Dutch book interpretation. *Artificial Intelligence*, 227:140–164.

de Morais *et al.***(2015)** E.M. de Morais, G. De Bona and M. Finger. Non-markovian logic-probabilistic modeling and inference. In *IJCAI-2015 Workshop on Weighted Logics for Artificial Intelligence (WL4AI-2015)*, pages 48–55.

De Bona *et al.***(2015)** G. De Bona, F.G. Cozman and M. Finger. Generalized probabilistic satisfiability through integer programming. *Journal of the Brazilian Computer Society*, 21(1): 1–14.

Finger and De Bona(2015) M. Finger and G. De Bona. Probabilistic satisfiability: algorithms with the presence and absence of a phase transition. *Annals of Mathematics and Artificial Intelligence*, pages 1–39. ISSN 1012-2443.

De Bona *et al.***(2014)** G. De Bona, F.G. Cozman and M. Finger. Towards classifying propositional probabilistic logics. *Journal of Applied Logic*, 12(3):349–368.

De Bona *et al.***(2013)** G. De Bona, F.G. Cozman and M. Finger. Generalized probabilistic satisfiability. In *Intelligent Systems (BRACIS), 2013 Brazilian Conference on*, pages 182–188. IEEE.

Finger and De Bona(2011) M. Finger and G. De Bona. Probabilistic satisfiability: logic-based algorithms and phase transition. In *Proceedings of IJCAI'11*.

Chapter 2

Preliminaries

Probabilistic knowledge bases are formed by attaching conditional probabilities — or, more generally, probabilistic constraints — to formulas from a given object language. We focus in assigning probability bounds to formulas from classical propositional logic. In Section 2.1, we briefly review such logic before inserting probabilities in Section 2.2. After introducing the probabilistic logic we are interested in, we discuss its inherent problem of consistency checking.

2.1 Classical Propositional Logic

The propositional language we consider is formed by a finite set of *atomic propositions* (or *atoms* or *logical variables*) $X_n = \{x_1, x_2, \ldots, x_n\}$ combined by logical connectives, possibly with punctuation elements (parentheses). A *well-formed formula* (or, simply, a *formula*) of the propositional language over a fixed X_n can be defined recursively from atomic propositions and connectives:

- If $\varphi \in X_n$ is a formula;
- If φ is a formula, then $\neg\varphi$ is a formula;
- If φ and ψ are formulas, then $(\varphi \vee \psi)$, $(\varphi \wedge \psi)$ and $(\varphi \rightarrow \psi)$ are formulas.

Given a set X_n of atomic propositions, the *propositional language* over it, $\mathcal{L}_{X_n}$, is the smallest set containing all the corresponding well-formed formulas.

To simplify notation, parentheses in formulas may be omitted, since there is a precedence order among the connectives: $\neg$ (*negation*), $\wedge$ (*conjunction*), $\vee$ (*disjunction*) and $\rightarrow$ (*implication*). We will also use the connective $\leftrightarrow$ (*bi-implication*) as an abbreviation, with $\varphi \leftrightarrow \psi$ denoting $(\varphi \rightarrow \psi) \wedge (\psi \rightarrow \varphi)$. The symbol $\top$ denotes the formula $(x_1 \vee \neg x_1)$, and its negation is represented by $\perp$.

Each atomic proposition can be assigned a truth value, either $TRUE$ or $FALSE$, which we denote by 1 and 0, respectively. A *valuation* (*truth assignment*) is a function $v : X_n \rightarrow \{0, 1\}$ that takes atomic propositions and returns truth values.

The domain of valuations can be extended to the whole propositional language $\mathcal{L}_{X_n}$ according to the semantic definition of each connective within a given logic. Let φ and ψ be formulas of the propositional language; within classical propositional logic, we have:

- $v(\varphi \wedge \psi) = 1$ iff[1] $v(\varphi) = 1$ and $v(\psi) = 1$;

[1] Throughout the text, "iff" abbreviates "if, and only if".

- $v(\varphi \vee \psi) = 1$ iff $v(\varphi) = 1$ or $v(\psi) = 1$;
- $v(\neg\varphi) = 1$ iff $v(\varphi) = 0$;
- $v(\varphi \rightarrow \psi) = 1$ iff $v(\varphi) = 0$ or $v(\psi) = 1$.

We say a valuation v *satisfies* a formula $\varphi \in \mathcal{L}_{X_n}$ iff $v(\varphi) = 1$. A valuation satisfies a set of formulas if it satisfies all formulas within the set. When a formula $\varphi \in \mathcal{L}_{X_n}$ (or a set $B \subseteq \mathcal{L}_{X_n}$) is such that there is valuation v satisfying it, it is said to be *satisfiable*, or *consistent*; otherwise, φ (or the set B) is *unsatisfiable*, or *inconsistent*. *Boolean Satisfiability* is a decision problem concerned with whether or not a finite set of formulas $B = \{\varphi_i \in \mathcal{L}_{X_n} | 1 \leq i \leq m\}$ is satisfiable; Boolean satisfiability is an NP-complete problem (Cook, 1971).

The set of atomic propositions X_n defines a set of 2^n *possible worlds*, which are formulas in $\mathcal{L}_{X_n}$ in the form $\ell_1 \wedge \ell_2 \wedge \ldots, \ell_n$, where, for each $1 \leq i \leq n$, either $\ell_i = x_i$ or $\ell_i = \neg x_i$. For instance, for $X_2 = \{x_1, x_2\}$, there are 4 possible worlds: $x_1 \wedge x_2$, $x_1 \wedge \neg x_2$, $\neg x_1 \wedge x_2$ and $\neg x_1 \wedge \neg x_2$. We denote by $W_n = \{w_1, w_2, \ldots, w_{2^n}\}$ the set of possible worlds corresponding to $X_n = \{x_1, x_2, \ldots, x_n\}$.

As usual, an *entailment* relation $\models \subseteq \mathcal{L}_{X_n} \times \mathcal{L}_{X_n}$ is defined as the set of all pairs of formulas (φ, ψ) such that, for any valuation v, if $v(\varphi) = 1$, then $v(\psi) = 1$. We write $\varphi \models \psi$, which reads "φ entails ψ", to denote $(\varphi, \psi) \in \models$. We say φ and ψ are *equivalent* if $\varphi \models \psi$ and $\psi \models \varphi$. Any formula equivalent to $\top$ is called a *tautology*, and a formula is a *contradiction* when it is equivalent to $\bot$. We define Classical Propositional Logic as the pair $(\mathcal{L}_{X_n}, \models)$.

While discussing belief revision in Chapter 5, we use consequence operations instead of entailment relations, in order to be aligned with the corresponding literature. In particular, we define the classical consequence operation $Cn_{Cl} : 2^{\mathcal{L}_{X_n}} \rightarrow 2^{\mathcal{L}_{X_n}}$ such that $\varphi \in Cn_{Cl}(B)$ iff every valuation that satisfies B also satisfies φ.

Each valuation $v : \mathcal{L}_{X_n} \rightarrow \{0, 1\}$ corresponds to a particular possible world $w \in W_n$ in the following way: $v(x_i) = 1$ iff $w \models x_i$ for all $1 \leq i \leq n$. In this sense, one can alternatively define a formula $\varphi \in \mathcal{L}_{X_n}$ as satisfiable when there is a possible world $w \in W_n$ such that $w \models \varphi$.

2.2 Propositional Probabilistic Logic

Classical logic is well-known for its limitation in coping with uncertainty. In the real world, natural agents (human beings) often lack sufficient confidence to hold a full belief in a given proposition, like "it will rain tomorrow in São Paulo". It has been argued — in many different ways — that degrees of belief *should* be probabilities, in the sense that they must respect the probability axioms (de Finetti, 1974; Howson and Urbach, 1989; Joyce, 1998). Consequently, probabilistic logic has been a standard approach to modelling the rationality of natural agents in formal epistemology, a field whose aim "is to harness the power of formal methods to bring rigor and clarity to philosophical analysis" (Wheeler, 2012). In AI, practical applications of automated reasoning usually require a framework capable of representing and handling the uncertain information that is natural to some domains. Probability theory is a well-established mathematical tool to manipulate uncertainty quantitatively. Hence, combining probability and logic has been applied to the development of expert systems in Artificial Intelligence (Adlassnig and Kolarz, 1986; Rödder and Meyer, 1996).

Departing from classical propositional logic, there are several ways to extend its syntax and semantics with probabilities. One may want to express a nested probability, like "the probability of

the probability of φ being q is q'", or to encode disjunctions like "the probability of φ is either q or q'". Additionally, instead of talking about precise probabilities, it is possible to represent probability intervals, or even relate probability values as "the probability of φ is at least the probability of ψ squared"; for a survey on these possibilities, see (De Bona *et al.*, 2014). As there is a trade-off between expressivity power and computational complexity, most logical systems are relatively simple, not allowing probability nesting nor disjunction of probabilities and assigning only precise probabilities or probability bounds. Furthermore, these choices have practical, semantic reasons, for the meaning of higher-order probabilities is at least debatable (Uchii, 1973) and probabilities, when not precise, are almost always assessed via closed intervals.

Due to applicability, the commonest approaches embed conditional probabilities, which can encode probabilistic rules, or a form of "causation to some degree". For instance, consider propositions like "Disease D causes symptom S in q% of patients", which can naturally be captured by conditional probabilities. Expert systems usually encode knowledge about a given domain into conditional probabilities (Adlassnig and Kolarz, 1986; Rödder and Meyer, 1996), which are also present in the frameworks of Bayesian (Pearl, 2009) and credal networks (Cozman, 2000). Therefore, the present study focuses on sets of imprecise conditional probabilities assigned to formulas from classical propositional logic, which we now formally define.

2.2.1 Syntax and Semantics

Fixed a propositional language $\mathcal{L}_{X_n}$, a *probabilistic conditional* (or simply a *conditional*) is a formula of the form $P(\varphi|\psi) \geq q$, with the intended meaning "the probability that φ is true given that ψ is true is at least q", where $\varphi, \psi \in \mathcal{L}_{X_n}$ are propositional formulas and q is a real number in $[0, 1]$. If ψ is a tautology, a conditional like $P(\varphi|\psi) \geq q$ is called an *unconditional probabilistic assessment*, usually denoted by $P(\varphi) \geq q$. We denote by $\mathcal{L}^P_{X_n}$ the set of all conditionals over $\mathcal{L}_{X_n}$. To save notation, $P(\varphi|\psi) \leq q$ abbreviates $P(\neg\varphi|\psi) \geq 1 - q$ and $P(\varphi|\psi) = q$, a *precise probability assessment*, denotes the elements of the pair $\{P(\varphi|\psi) \leq q, P(\varphi|\psi) \geq q\}$; that is, $P(\varphi|\psi) = q \in \Gamma$ denotes $P(\varphi|\psi) \leq q, P(\varphi|\psi) \geq q \in \Gamma$, for any $\Gamma \subset 2^{\mathcal{L}}_{X_n}$. Furthermore, given $\alpha = P(\varphi|\psi) \geq q$ and $\beta = P(\varphi|\psi) \geq q'$, we use $\alpha \bowtie \beta$, with $\bowtie \in \{\leq, \geq, <, >\}$, to denote that $q \bowtie q'$.

A *probabilistic interpretation* $\pi : W_n \to [0, 1]$, with $\sum_j \pi(w_j) = 1$, is a probability mass over the set of possible worlds W_n, which induces a probability measure $P_\pi : \mathcal{L}_{X_n} \to [0, 1]$ by means of $P_\pi(\varphi) = \sum\{\pi(w_j)|w_j \in W_n, w_j \models \varphi\}$. A conditional $P(\varphi|\psi) \geq q$ is *satisfied* by π iff $P_\pi(\varphi \wedge \psi) \geq qP_\pi(\psi)$. Note that, when $P_\pi(\psi) > 0$, a probabilistic conditional $P(\varphi|\psi) \geq q$ is constraining the conditional probability of φ given ψ; but any π with $P_\pi(\psi) = 0$ trivially satisfies that conditional (this semantics is adopted by Halpern (1990), Frisch and Haddawy (1994) and Lukasiewicz (1999), for instance). A *probabilistic knowledge base* (or simply a base) is a set of probabilistic conditionals $\Gamma \subseteq \mathcal{L}^P_{X_n}$. When a base Γ can be written as a set of precise (or unconditional) probability assignments, we say it is *precise* (*unconditional*). We denote by $\mathbb{K}$ the set of all probabilistic knowledge bases. A knowledge base Γ is *consistent* (or *satisfiable*) if there is a probability interpretation satisfying all conditionals $P(\varphi|\psi) \geq q \in \Gamma$. For a probabilistic base Γ, we denote by $Cn_{Pr}(\Gamma)$ the set of all conditionals $\alpha \in \mathcal{L}^P_{X_n}$ such that, if π satisfies Γ, then π satisfies α — this is known as the *standard semantics* (Haenni *et al.*, 2011). That is, Cn_{Pr} is the consequence operation corresponding to the standard semantics of probabilistic logic.

An alternative semantics for probabilistic entailment, which has been adopted for instance by

Kern-Isberner and Lukasiewicz (2004), defines that a base Γ entails a conditional α if the probabilistic interpretation π that maximises entropy while satisfying Γ also satisfies α. This entropy maximisation is endorsed by the Objective Bayesianism Epistemology, for instance, by means of the Equivocation norm (Williamson, 2010). As the inconsistency of a probabilistic base is defined in both the standard and the maximum entropy semantics as the absence of a probabilistic interpretation satisfying it, the problems of localising, measuring and repairing the inconsistency are the same in both approaches.

A finite base $\Gamma \in \mathbb{K}$ is said to be *canonical* if it is such that, if $P(\varphi|\psi) \geq q, P(\varphi|\psi) \geq q' \in \Gamma$, then $q = q'$. That is, for each pair φ, ψ, only one probability lower bound can be assigned to $P(\varphi|\psi)$ in a canonical probabilistic knowledge base.[2] We denote by $\mathbb{K}_c$ the set of all canonical probabilistic knowledge bases. Any base $\Gamma \in \mathbb{K}$ is equivalent to a canonical base — one can simply pick the highest probability lower bound of each set of conditionals $P(\varphi|\psi) \geq q$ on the same pair φ, ψ; if the resulting base is not finite, one can additionally repeat this procedure modulo equivalence on $\mathcal{L}_{X_n}$.[3]

2.2.2 Deciding the Consistency of Probabilistic Bases

Probabilistic Satisfiability (PSAT)

To present the problem of consistency checking in the probabilistic logic just introduced, we begin with a simpler fragment of the language. The *Probabilistic Satisfiability* problem (*PSAT*) lies in verifying whether or not an unconditional base that can be written as $\{P(\varphi_i) = q_i | 1 \leq i \leq m\}$, where $\varphi_i \in \mathcal{L}_{X_n}$ and $q_1, \ldots, q_m \in [0, 1]$ for all $1 \leq i \leq m$, is consistent. That is, PSAT is concerned with the consistency of finite sets of unconditional, precise probability assignments. Each base in this format is called a *PSAT instance*. When the base has the form $\{P(\varphi_i) \geq q_i | 1 \leq i \leq m\}$, with unconditional, but imprecise assignments, the problem of checking consistency is called *Probabilistic Satisfiability with Imprecise Probabilities* (*PSAT-IP*); and the base is called a *PSAT-IP instance*. PSAT is a particular case of PSAT-IP, but the latter can be polynomially reduced to the former, at least when probability bounds are rational, by employing techniques as the atomic normal form, from (Finger and De Bona, 2015; Finger and De Bona, 2011).

Example 2.2.1. A physician investigates the relation among 3 genes and the occurrence of a disease D. His working hypothesis is that, for the disease to develop in a given patient, at least 2 of the 3 genes must be present. Experiments show that the presence of each gene was detected in 60% of the patients who developed disease D. The physician wants to know if the experimental data is consistent with his hypothesis.

We can associate atomic propositions x_1, x_2 and x_3 to the occurrence of each gene in a given patient of disease D and represent the physician's hypotheses as $P((x_1 \vee x_2) \wedge (x_1 \vee x_3) \wedge (x_2 \vee x_3)) = 1$. Assigning probability 1 to this formula corresponds to expecting that every patient of disease D hold at least 2 of the 3 genes. Associating the frequency of the presence of each gene to the probability of each being present in a given patient, the experimental data can be encoded into three probability

[2] Note that this requirement is not too restrictive. Since nothing was said about logically equivalent propositions, a knowledge base may contain different probability lower bounds assigned to φ and $\varphi \wedge \varphi$, for instance.

[3] In the literature, instead of using canonical bases, probabilistic knowledge bases have been defined as ordered sets, allowing for repeated formulas; but this loosens the link to the belief revision approach — explored in Chapter 5 — which usually operates on non-ordered sets.

assessments: $P(x_1) = 0.6$, $P(x_2) = 0.6$ and $P(x_3) = 0.6$. We have constructed the following PSAT instance:

$$\begin{aligned} \{\, P((x_1 \vee x_2) \wedge (x_1 \vee x_3) \wedge (x_2 \vee x_3)) &= 1, \\ P(x_1) &= 0.6, \\ P(x_2) &= 0.6, \\ P(x_3) &= 0.6 \,\} \end{aligned}$$

Each patient of disease D may or may not have each gene, which allows us to classify them in 8 different groups, corresponding to the possible worlds over $X_3 = \{x_1, x_2, x_3\}$. If this PSAT instance is satisfiable, there is a possible distribution of the patients who developed disease D in these 8 groups where each patient has 2 from the 3 genes, while each gene occurs in exactly 60% of the patients — and the physician's hypothesis is consistent with the experimental data. Otherwise, the absence of a solution to this PSAT instance is a certificate of the inconsistency between these experimental data and the physician's hypothesis. □

A PSAT instance can be mathematically expressed as an exponentially large linear programming problem, as Nilsson (1986) showed. Let $\Gamma = \{P(\varphi_i) = q_i | 1 \leq i \leq m\}$ be a PSAT instance and let $I_{w_j} : \mathcal{L}_{X_n} \to \{0, 1\}$ be the indicator function of the set $\{\varphi \in \mathcal{L}_{X_n} | w_j \models \varphi\}$, such that $I_{w_j}(\varphi) = 1$ if $w_j \models \varphi$, otherwise $I_{w_j}(\varphi) = 0$, for all $w_j \in W_n$. We define the matrix $A_{m \times 2^n} = [a_{ij}]$, such that $a_{ij} = I_{w_j}(\varphi_i)$, and the vector $q_{m \times 1} = [q_i]$. The base Γ is satisfiable iff there is a vector $\pi_{2^n \times 1} = [\pi_i]$ satisfying the following restrictions:

$$A\pi = q \tag{2.1}$$

$$\pi \geq 0 \tag{2.2}$$

$$\sum \pi = 1 \tag{2.3}$$

Constraints (2.2) and (2.3) force π to be a probability mass over the set W_n of possible worlds, which are represented by the columns in A. Restriction (2.3) may be ommited if a whole row of $1's$ is inserted into A, in such a way that $a_{m+1,j} = 1, 1 \leq j \leq 2^n$, and an element $q_{m+1} = 1$ is inserted into q.

When unconditional probabilities are assigned via bounds ($P(\varphi) \geq q$), the linear program above can be slightly modified to encode the satisfiability of a PSAT-IP instance. One only needs to substitute $A\pi \geq q$ for $A\pi = q$ as Restriction (2.1), and the resulting linear program has a solution iff the PSAT-IP instance is satisfiable.

The next example presents a satisfiable PSAT instance and its formulation via linear programming.

Example 2.2.2. A citizen from "Traffic City", who works from Monday to Friday, complains that, in average, only once a week he does not face traffic jams along the path to his work — and he wonders if the rain could be the cause. Meanwhile, a local newspaper points out that, in the city, it does not rain heavily in 90% of the days and only in 10% of the weekdays one comes across heavy traffic and heavy rain. We want to know whether the citizen's complaint is consistent with the newspaper's data.

To formalize this problem as a PSAT instance, we use two atomic propositions: for a given weekday in "Traffic City", $r = x_1$ is true iff it rains heavily, and $t = x_2$ is true iff there is heavy traffic. We can say that he does not face heavy traffic in 20% of the weekdays, or, interpreting probabilities as frequencies, $P(\neg t) = 20\%$. According to the newspaper, associating probabilities to frequencies again, and assuming that the probability of raining in a giving day is independent from whether or not it is a weekday, we have $P(\neg r) = 90\%$ and $P(r \wedge t) = 10\%$. The citizen's complaint is consistent with the newspaper's data iff the following PSAT instance is satisfiable:

$$\Gamma = \{P(\neg t) = 20\%, P(\neg r) = 90\%, P(r \wedge t) = 10\%\}\,.$$

To construct the corresponding linear program, we define a matrix $A_{3\times 2^2} = [a_{ij}]$ such that $a_{ij} = I_{w_j}(\varphi_i)$. With 2 atomic propositions, there are 4 columns in A, corresponding to the possible worlds: $w_1 = \neg r \wedge \neg t$, $w_2 = r \wedge \neg t$, $w_3 = r \wedge t$, $w_4 = \neg r \wedge t$. The 3 probability assessments yield 3 rows in A:

φ_i	w_1	w_2	w_3	w_4
$\neg t$	1	1	0	0
$\neg r$	1	0	0	1
$r \wedge t$	0	0	1	0

The question now is whether there is a vector π that satisfies Restrictions (2.1)-(2.3), given the matrix A above and the vector $q = [0.2\,0.9\,0.1]^T$. Inserting a row of 1's at the bottom of A and adding a 1 to q to represent Restriction (2.3), we show a vector $\pi \geq 0$ such that $A\pi = q$:

$$\overset{A}{\begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 1 & 1 & 1 & 1 \end{bmatrix}} \cdot \overset{\pi}{\begin{bmatrix} 0.2 \\ 0 \\ 0.1 \\ 0.7 \end{bmatrix}} = \overset{q}{\begin{bmatrix} 0.2 \\ 0.9 \\ 0.1 \\ 1 \end{bmatrix}}$$

We can finally conclude that the PSAT Γ is satisfiable, attesting the consistency between the citizen's complaint and the newspaper's data. This probability mass π also presents a possible scenario that gives a clue about the the relation between rain and traffic: although there is heavy traffic on every weekday during which it heavily rains ($P_\pi(r \wedge \neg t) = 0$), on most weekdays there is heavy traffic even without heavy rain ($P_\pi(\neg r \wedge t) = 70\%$). □

By Carathéodory's Lemma, if the linear program (2.1)-(2.3) has a solution π, while encoding a PSAT or PSAT-IP instance, there is some solution $\pi^\star$ with at most $m + 1$ positive elements (Prasolov and Tikhomirov, 2001). As Georgakopoulos *et al.* (1988) noted, when probabilities are restricted to rational numbers, this fact implies that PSAT is in NP. One can take a vector $\pi^\star_{m+1,1} \geq 0$, discarding $2^n - m + 1$ null elements of π, and a matrix $A_{m+1,m+1}$, selecting the corresponding columns from A and embedding the Restriction (2.3), as an NP-certificate, verifying the relations (2.1)-(2.3) in polynomial time. Furthermore, the problem of deciding the Boolean satisfiability of a set $\{\varphi_1, \ldots, \varphi_m\} \subseteq \mathcal{L}_{X_n}$ can be reduced to a PSAT instance in polynomial time on m: just take $P(\varphi_i) = 1$, for $1 \leq i \leq m$. It follows that PSAT with rational probabilities is NP-hard and, consequently, NP-complete (Georgakopoulos *et al.*, 1988); and so is PSAT-IP with rational probability bounds.

From the practical point of view, as all constraints in (2.1)-(2.3) are linear, this system can be solved by linear programming techniques as Simplex. Despite the exponential number of columns, column generation methods can be used to handle them implicitly (Jaumard *et al.*, 1991; Kavvadias and Papadimitr , 1990), keeping the computation efficient enough to solve large knowledge bases — thousands of probabilities in (Finger *et al.*, 2013; Hansen and Perron, 2008). An open source implementation of a PSAT solver through column generation from Finger and De Bona (2011) is available at http://psat.sourceforge.net.

PSAT-IP with Conditional Probabilities

To decide the consistency of arbitrary finite probabilistic bases in $\mathbb{K}$, we can extend the linear programming approach to conditional probabilities. Jaumard *et al.* (1991) showed how the approach of Nilsson (1986) can be adapted to decide consistency in this general case, which was called CONDSAT. We adapt their approach, considering only conditionals in the form $P(\varphi|\psi) \geq q \in \mathcal{L}^P_{X_n}$, which is the single primitive probability assignment in our language.

Consider a finite probabilistic knowledge base $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq m\}$. Under the semantics adopted, each assessment $P(\varphi_i|\psi_i) \geq q_i$ is satisfied by a probability mass π iff $P_\pi(\varphi_i \wedge \psi_i) - q_i P_\pi(\psi_i) \geq 0$. The knowledge base is consistent iff these m restrictions can be jointly satisfied by a probability measure P_π induced by a probability mass π. Consider the $(m \times 2^n)$-matrix $A = [a_{ij}]$, with $a_{ij} = I_{w_j}(\varphi_i \wedge \psi_i) - q_i I_{w_j}(\psi_i)$, in which $I_{w_j} : \mathcal{L}_{X_n} \to \{0, 1\}$ is the indicator function of the set $\{\varphi \in \mathcal{L}_{X_n} | w_j \models \varphi\}$ — I_{w_j} is the valuation relative to the possible world w_j. The knowledge base Γ is satisfiable iff there is a $(2^n \times 1)$-vector π satisfying the system:

$$A\pi \geq 0 \tag{2.4}$$

$$\pi \geq 0 \,. \tag{2.5}$$

$$\sum \pi = 1 \tag{2.6}$$

Again, Constraints (2.5) and (2.6) force π to be a probability mass over the set of possible worlds $\{w_1, w_2, \ldots, w_{2^n}\}$, and Restriction (2.6) can be embedded into A as an extra row. As all restrictions are kept linear, Simplex algorithm with column generation methods can also be applied in this case.

Given this formulation of CONDSAT as a linear program, it is also an NP-complete problem when probability bounds are rational numbers (Georgakopoulos *et al.*, 1988), for the same reasons as PSAT. Due to this fact, when we refer to PSAT computational aspects, we usually mean those of the most general CONDSAT.

2.2.3 Alternative Semantics for Probabilistic Consistency

Within the same language $\mathcal{L}^P_{X_n}$, there are at least two different possible consistency definitions for a probabilistic base, corresponding to alternative semantics. Note that the system formed by Restrictions (2.4)-(2.6) does not require $P_\pi(\psi_i) > 0$ for all $P(\varphi_i|\psi_i) \geq q_i \in \Gamma$. In other words, any π such that $P_\pi(\psi_i) = 0$ trivially satisfies the conditional $P(\varphi_i|\psi_i) \geq q_i \in \Gamma$ — which is in accordance with the semantics we adopt. In standard probability theory, $P(\varphi|\psi)$ is defined, as $P(\varphi \wedge \psi)/P(\psi)$, only when $P(\psi) > 0$. So it is reasonable to require that conditioning formulas have positive probability. This approach is followed, for instance, by Muiño (2011). To employ this, one could add

to the system (2.4)-(2.6) an extra restriction $P_\pi(\psi_i) > 0$ for each $P(\varphi_i|\psi_i) \geq q_i \in \Gamma$. This would lead to a problem that does not lie in the standard format of linear programming, due to the occurrence of strict inequalities — for linear programming, see for instance (Papadimitriou and Steiglitz, 1998). In some cases, such an inequality can be replaced by $P_\pi(\psi_i) \geq \varepsilon$, for a suitably chosen $\varepsilon > 0$, without changing the satisfiability (De Bona *et al.*, 2014). Nevertheless, in the general case it is not possible to approximate strict inequalities by non-strict ones without losing some precision.

Another approach to conditional probability arises when $P(\varphi|\psi)$ is well-defined even with $P(\psi) = 0$, by requiring consistency of probability assignments conditional to the same ψ. This is possible when Kolmogorov's approach to conditional probability, through the ratio definition, gives place to the subjective probability theory proposed by de Finetti (1949, 1974). In such theory, a conditional probability measure is a function P defined over pairs $\varphi|\psi$, with $\varphi, \psi \in \mathcal{L}_{X_n}$, such that ψ is satisfiable, respecting the following axioms:

1. $P(\varphi|\psi) \geq 0$ for any $\varphi, \psi \in \mathcal{L}_{X_n}$ such that ψ is satisfiable;
2. $P(\varphi|\psi) = 1$ for any pair $\varphi, \psi \in \mathcal{L}_{X_n}$ of equivalent formulas such that ψ is satisfiable;
3. $P(\varphi \vee \theta|\psi) = P(\varphi|\psi) + P(\theta|\psi)$ for all $\varphi, \theta, \psi \in \mathcal{L}_{X_n}$ such that ψ is satisfiable and $\varphi \wedge \theta \wedge \psi$ is unsatisfiable.
4. $P(\varphi \wedge \theta|\psi) = P(\varphi|\theta \wedge \psi)P(\theta|\psi)$ for all $\varphi, \theta, \psi \in \mathcal{L}_{X_n}$ such that $\theta \wedge \psi$ is satisfiable.

In this way, the conditional probability is primitive instead of the unconditional one, which may be understood as $P(\varphi|\top)$. The consistency of a set of probabilistic conditionals according to the axioms above is called *coherence*, and problem of coherence checking is equivalent to PSAT in the unconditional case.

Coletti (2002) investigated the coherence checking problem showing it can be reduced to the compatibility of a sequence of linear systems. This approach has an independent interest and is not in the focus of the present work; for studies on coherence checking, see e.g. (Capotorti and Vantaggi , 1998) and (Coletti and Scozzafava, 1996, 2002).

Chapter 3

Inconsistency Measures in Probabilistic Logic and Rationality Postulates

Approaches to measuring inconsistency in probabilistic knowledge bases have been put forward by Muiño (2011), Thimm (2013) and Potyka (2014), with different semantics for the conditionals. We follow the one adopted by Thimm and Potyka, in which a conditional is also satisfied by any measure assigning null probability to the conditioning formula. Following Thimm (2013), in Section 3.1 we extend to the probabilistic case the inconsistency measures and rationality postulates Hunter and Konieczny (2006) proposed in classical logic. Imposing a postulate of continuity, Thimm (and, independently, Muiño) suggested a family of inconsistency measures for probabilistic bases based on distance minimisation, which are presented in Section 3.2.1. Concerned with computational complexity, Potyka (2014) put forward feasible measures we will review in Section 3.2.2. Related work within the coherence setting and under a different semantics, from Capotorti, Regoli, and Vattari (2010), is reviewed and adapted in Section 3.3. Finally, in Section 3.4, we argue against the joint satisfiability of three basic rationality postulates for measuring inconsistency in probabilistic logic, revisiting a result firstly published in (De Bona and Finger, 2015).

3.1 General Inconsistency Measures and Postulates

An *inconsistency measure* is a function $\mathcal{I} : \mathbb{K}_c \to [0, \infty)$, which takes canonical probabilistic knowledge bases and returns non-negative real numbers. We focus on canonical bases, for they are finite, as any practical application requires, without losing expressivity power. The reason why negative measurements are avoided becomes clear when the first desired property for inconsistency measures is presented (Hunter and Konieczny, 2006):

Postulate 3.1.1 (Consistency). *For any $\Gamma \in \mathbb{K}_c$, $\mathcal{I}(\Gamma) = 0$ iff Γ is consistent.*

That is, we want a measure whose minimum value is reserved to consistent bases. An inconsistency measure that only differentiates inconsistent from consistent bases, satisfying the postulate of (Consistency), is the *drastic measure* $\mathcal{I}_d : \mathbb{K}_c \to [0, \infty)$ (Hunter and Konieczny, 2006):

$$\mathcal{I}_d(\Gamma) = \begin{cases} 1, & \text{if } \Gamma \text{ is inconsistent,} \\ 0, & \text{otherwise.} \end{cases} \tag{3.1}$$

Besides (Consistency), the drastic inconsistency measure satisfies another basic postulate, which states that adding information to a base cannot make it less inconsistent (Hunter and Konieczny, 2006):

Postulate 3.1.2 (**Monotonicity**). *For any* $\Gamma \cup \{\alpha\} \in \mathbb{K}_c$, $\mathcal{I}(\Gamma \cup \{\alpha\}) \geq \mathcal{I}(\Gamma)$.

The main feature of the drastic measure is its computational cost, which is equivalent to decide the satisfiability of the base — an NP-complete problem when probabilities are rational. Of course, NP-completeness is at the border of computational feasibility and does not imply tractability, but to satisfy (Consistency), any inconsistency measure must at least check the base consistency.

Although the drastic measure has some positive attributes, it is not of great usefulness, since it is not capable of discriminating inconsistent bases among themselves. For instance, one expects that the base $\Delta = \{P(x_1) \geq 0.5, P(\neg x_1) \geq 0.6\}$ be less inconsistent that the base $\Gamma = \Delta \cup \{P(\bot) \geq 0.1\}$, since the latter is formed by adding more contradictory information ($P(\bot) \geq 0.1$) to the former, which was already unsatisfiable. However, the drastic measure cannot tell such difference: $\mathcal{I}_d(\Delta) = \mathcal{I}_d(\Gamma) = 1$. Based on this, a first idea to refine such inconsistency measure is to look at the number of "atomic" inconsistencies, or primitive conflicts, in the base. The intuition says that a base's degree of inconsistency should increase according to the quantity of problematic sets of formulas that cause the inconsistency. Classically, minimal inconsistent sets are the characterisation of such problematic sets:

Definition 3.1.3. A probabilistic knowledge base $\Gamma \in \mathbb{K}$ is a *minimal inconsistent set (MIS)* if Γ is inconsistent and every set $\Gamma' \subsetneq \Gamma$ is consistent.

When a minimal inconsistent set Δ is a subset of a base $\Gamma \in \mathbb{K}$, we say Δ is a *minimal inconsistent subset* of Γ — a *MIS of* Γ. Let $\mathrm{MIS} : \mathbb{K} \to 2^{\mathbb{K}}$ be a function such that, for any $\Gamma \in \mathbb{K}$, $\mathrm{MIS}(\Gamma)$ is the set of all MISes of Γ.

Now, the measure $\mathcal{I}_{MIS} : \mathbb{K}_c \to [0, \infty)$ is defined simply by counting the number of minimal inconsistent sets in a base (Hunter and Konieczny, 2008; Thimm, 2013):

$$\mathcal{I}_{MIS}(\Gamma) = |\,\mathrm{MIS}(\Gamma)|\,. \tag{3.2}$$

As $\Delta = \{P(x_1) \geq 0.5, P(\neg x_1) \geq 0.6\}$ has one MIS while $\Gamma = \Delta \cup \{P(\bot) \geq 0.1\}$ has two, for $\{P(\bot) \geq 0.1\}$ is a MIS[1], it follows that $\mathcal{I}_{MIS}(\Delta) = 1 < 2 = \mathcal{I}_{MIS}(\Gamma)$.

Despite the fact that $\mathcal{I}_{MIS}$ considers the quantity of minimal inconsistent sets, it does not assess their severity. A way to accomplish that is grounded in the idea that the larger a MIS is, the less inconsistent it is. For instance, consider the bases $\Psi = \{P(\bot) \geq 0.1\}$, $\Delta = \{P(x_1) \geq 0.5, P(\neg x_1) \geq 0.6\}$ and $\Gamma = \{P(x_1 \vee x_2) \leq 0.5, P(\neg x_1 \vee x_2) \leq 0.5, P(x_2 \vee x_3) \geq 0.5, P(x_3) \leq 0.4\}$. Note that $P(x_3) \leq 0.4$ and $P(x_2 \vee x_3) \geq 0.5$ imply $P(x_2) \geq 0.1$ and that $P(x_1 \vee x_2) \leq 0.5$ and $P(\neg x_1 \vee x_2) \leq 0.5$ imply $P(x_2) = 0$; thus Γ is inconsistent. Since Ψ, Δ and Γ are minimal inconsistent sets, $\mathcal{I}_{MIS}(\Psi) = \mathcal{I}_{MIS}(\Delta) = \mathcal{I}_{MIS}(\Gamma) = 1$. Nevertheless, while in Γ four conditionals are required to produce the inconsistency, the two conditionals in Δ are already contradictory, and the single conditional in Ψ is unsatisfiable by itself. With this intuition, the inconsistency measure $\mathcal{I}_{MIS^C} = \mathbb{K}_c \to [0, \infty)$ is defined in a way that each MIS's contribution to the whole inconsistency is inversely proportional to its size (Hunter and Konieczny, 2008; Thimm, 2013):

[1]The empty set of conditionals is trivially consistent.

$$\mathcal{I}_{MIS^C}(\Gamma) = \sum \left\{ \frac{1}{|\Delta|} \mid \Delta \in \mathrm{MIS}(\Gamma) \right\} . \tag{3.3}$$

Recalling the probabilistic bases $\Psi = \{P(\bot) \geq 0.1\}$, $\Delta = \{P(x_1) \geq 0.5, P(\neg x_1) \geq 0.6\}$ and $\Gamma = \{P(x_1 \vee x_2) \leq 0.5, P(\neg x_1 \vee x_2) \leq 0.5, P(x_2 \vee x_3) \leq 0.5, P(x_3) \leq 0.4\}$, we have that $\mathcal{I}_{SIM^C}(\Psi) = 1 > \mathcal{I}_{SIM^C}(\Delta) = 1/2 > \mathcal{I}_{SIM^C}(\Gamma) = 1/4$.

Note that $\mathcal{I}_{MIS}$ and $\mathcal{I}_{MIS^C}$ are such that a conditional not belonging to any MIS does not contribute to the base's inconsistency. This property is captured by the following postulate for inconsistency measures (Hunter and Konieczny, 2006; Thimm, 2013):

Definition 3.1.4. A conditional α in a base $\Gamma \in \mathbb{K}$ is said to be *free* in Γ if, for all $\Delta \in \mathrm{MIS}(\Gamma)$, $\alpha \notin \Delta$.

Postulate 3.1.5 (Independence). *For any $\Gamma \in \mathbb{K}_c$ and $\alpha \in \Gamma$, if α is a free conditional in Γ, then $\mathcal{I}(\Gamma \setminus \{\alpha\}) = \mathcal{I}(\Gamma)$.*

Indeed, $\mathcal{I}_{MIS}$ and $\mathcal{I}_{MIS^C}$ enjoy a stronger property, according to which if the base can be split into two parts without "breaking" any MIS, the inconsistency of the whole base is the sum of the inconsistency of these parts. The following version is adapted from (Thimm, 2013), although the original version in the classical setting is due to Hunter and Konieczny (2006):

Postulate 3.1.6 (MIS-Separability). *For any $\Delta, \Psi, \Gamma \in \mathbb{K}_c$, if $\Gamma = \Delta \cup \Psi$, $\mathrm{MIS}(\Gamma) = \mathrm{MIS}(\Delta) \cup \mathrm{MIS}(\Psi)$ and $\Delta \cap \Psi = \varnothing$, then $\mathcal{I}(\Gamma) = \mathcal{I}(\Delta) + \mathcal{I}(\Psi)$.*

Given (Consistency), (MIS-Separability) implies (Independence)[2]:

Proposition 3.1.7. *If $\mathcal{I}$ satisfies (MIS-separability) and (Consistency), then $\mathcal{I}$ satisfies (Independence).*

The property imposed by (MIS-Separability) enables computing the inconsistency of a base Γ through a partition such that each MIS of Γ is a subset of a partition's element. The drastic measure does not satisfy such property, for instance.

We can summarise the properties of these inconsistency measures based on minimal inconsistent subsets:

Proposition 3.1.8. *$\mathcal{I}_{MIS}$ and $\mathcal{I}_{MIS^C}$ satisfy (Consistency), (Monotonicity), (Independence) and (MIS-separability).*

The computation of both $\mathcal{I}_{MIS}$ and $\mathcal{I}_{MIS^C}$ requires that all MISes in a base Γ be found, which is a computational task considerably harder than solving PSAT. The quantity of MISes in a base may be exponential in its size, and each MIS demands at least a satisfiability check to be found (Klinov, 2011).

Knight (2002) puts forward a different approach to measure inconsistency in classical propositional logic. The idea is to measure the consistency of a classical base B as the highest probability lower bound that can be consistently assigned to all formulas in B — and the inconsistency of a base can be defined as 1 minus its consistency. Formally, the base $B = \{\varphi_i \in \mathcal{L}_{X_n} | 1 \leq i \leq m\}$ is said to be *η-consistent* if the probabilistic base $\{P(\varphi_i) \geq \eta | 1 \leq i \leq m\}$ is consistent. If η is

[2]Proofs of technical results omitted from the main text are in Appendix A.

such that, for any $\eta' > \eta$, the base B is not η'-consistent, then B is *maximally η-consistent*. As a base is maximally 1-consistent iff it is consistent, we can define the inconsistency of a maximally η-consistent base as $1 - \eta$.

To adapt the approach of Knight (2002) to the probabilistic setting, we need to define when a set of probabilities assigned to probabilistic conditionals is consistent; a second-order PSAT, so to speak. We denote by Π_n be the set of all probabilistic interpretations $\pi : W_n \to [0, 1]$. Given probabilistic conditionals $\alpha_1, \ldots, \alpha_m$, we say the set $\{P(\alpha_i) \geq \eta_i | 1 \leq i \leq m\}$, with $\eta_1, \ldots, \eta_m \in [0, 1]$, is consistent iff there is a probability mass $\pi_2 : \Pi_n \to [0, 1]$ such that

$$P_{\pi_2}(\alpha_i) = \sum \{\pi_2(\pi_1) | \pi_1 \in \Pi_n \text{ satisfies } \alpha_i\} \geq \eta_i, \text{ for all } 1 \leq i \leq m \ .$$

In other words, π_2 induces a second order probability $P_{\pi_2}(\alpha_i)$, which is the sum of the probabilities assigned to the probabilistic interpretations that satisfy α_i. Now, the η-inconsistency measure $\mathcal{I}_\eta : \mathbb{K}_c \to [0, 1]$ can be defined as:

$$\begin{aligned} \mathcal{I}_\eta(\Gamma) &= 1 - \eta^\star, \text{ where} \\ \eta^\star &= \max\left\{\eta \,|\, \{P(\alpha_i) \geq \eta | \alpha_i \in \Gamma\} \text{ is consistent}\right\}. \end{aligned}$$

To decide the consistency of $\{P(\alpha_i) \geq \eta | \alpha_i \in \Gamma\}$, we can employ the same linear programming formulation as PSAT's, where each column would correspond to the indicator function I_π of a set $\{\alpha \in \mathcal{L}^P_{X_n} | \pi \text{ satisfies } \alpha\}$. To apply Simplex with column generation techniques to this problem, as done for PSAT (Jaumard *et al.*, 1991), we could use an oracle to the Generalised Satisfiability Problem (GPSAT) (De Bona *et al.*, 2013; De Bona *et al.*, 2015), which can handle the kind of nested probabilities $P(P(\varphi) \geq q) \geq \eta$ that would occur. When probability bounds are rational, GPSAT is NP-complete, thus no harder than PSAT in principle, but the disjunctive programming the former requires implies more computational effort in practice (De Bona *et al.*, 2015). Finding the greatest η leading to consistency is simply a modification of the sketched linear program, as in the optimisation version of PSAT — called OPSAT by Kavvadias and Papadimitriou (1990).

Since η is always between 0 and 1, $\mathcal{I}_\eta$ possess an interesting property:

Property 3.1.9 (Normalisation). *For any $\Gamma \in \mathbb{K}_c$, $\mathcal{I}(\Gamma) \in [0, 1]$.*

We state (Normalisation) as a property, and not a postulate, for it is known to be incompatible with the following property, as proved by Thimm (2013):

Property 3.1.10 (Super-additivity). *For any $\Gamma \cup \Delta \in \mathbb{K}_c$, if $\Gamma \cap \Delta = \varnothing$, then $\mathcal{I}(\Gamma \cup \Delta) \geq \mathcal{I}(\Gamma) + \mathcal{I}(\Delta)$.*

Proposition 3.1.11. *If $\mathcal{I}$ satisfies (Super-additivity), then $\mathcal{I}$ satisfies (Monotonicity).*

Note that (Super-additivity) is enjoyed by both $\mathcal{I}_{MIS}$ and $\mathcal{I}_{MIS^C}$, which naturally violate (Normalisation).

Now we can summarise the properties of $\mathcal{I}_\eta$:

Proposition 3.1.12. *$\mathcal{I}_\eta : \mathbb{K}_c \to [0, \infty)$ satisfies (Consistency), (Independence), (Monotonicity) and (Normalisation) and violates (MIS-Separability).*

3.2 Inconsistency Measures Tailored to the Probabilistic Logic

A problem common to the four metrics presented in Section 3.1, when applied to probabilistic bases, may be illustrated with the following example. Suppose $\Gamma_\varepsilon = \{P(\bot) \geq \varepsilon\}$, for some $\varepsilon \in [0,1]$. When $\varepsilon > 0$, each measure seen in the previous section returns the same inconsistency degree for all bases Γ_ε, regardless of the value of $\varepsilon \in (0,1]$. Nevertheless, intuition seems to indicate that $P(\bot) \geq 0.001$ is far less inconsistent that $P(\bot) \geq 1$. The introduced measures are all *qualitative*, depending only on the consistency of each subset of the base. Thimm (2013) points to the possibility of *quantitative* inconsistency measures for probabilistic bases, exploiting the natural gradation of probabilities.

A limitation of classical measures, derived from their qualitative character, is related to continuity. It seems intuitive that small changes in the probability bounds yield small changes in the inconsistency severity. For $\varepsilon = 0$, the base Γ_ε is consistent, and the (Consistency) postulate implies $\mathcal{I}(\Gamma_\varepsilon) = 0$. Nonetheless, as each measure we presented returns a fixed positive value for all $\varepsilon > 0$, there is a discrete leap in $\mathcal{I}(\Gamma_\varepsilon)$ when the parameter ε is minimally perturbed from 0.

To formalise this continuity intuition, we need a formal method to change the probability bounds within a base. Given a base $\Gamma \in \mathbb{K}_c$, we define its *characteristic function* $\Lambda_\Gamma : [0,1]^{|\Gamma|} \to \mathbb{K}_c$ in the following way: if $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq m\}$ is a canonical base and $q' = \langle q'_1, \ldots, q'_m \rangle$ is a vector in $[0,1]^m$, $\Lambda_\Gamma(q') = \{P(\varphi_i|\psi_i) \geq q'_i | 1 \leq i \leq m\}$. That is, $\Lambda_\Gamma(.)$ only changes the probability bounds in Γ. For the function Λ_Γ be unique and well-defined, Thimm (2013) imposes some order on the set Γ, considering bases as sequences instead of sets. For simplicity, we just suppose there is some order over the probabilistic conditionals used to uniquely specify Λ_Γ [3]. Henceforth, when applying a characteristic function, we assume the underlying order over the conditionals is that in which they were presented, or that of their indices. Now an extra postulate for inconsistency measures in the probabilistic context can be stated (Thimm, 2013):

Postulate 3.2.1 (Continuity). *For any $\Gamma \in \mathbb{K}_c$, the function $\mathcal{I} \circ \Lambda_\Gamma : [0,1]^{|\Gamma|} \to [0,\infty)$ is continuous.*

To satisfy (Continuity), an inconsistency measure has to take into account not only the satisfiability of the base's subsets but also the probability bounds' numeric values. The approaches in the literature with this feature can mainly be grouped into two classes: the inconsistency of a base can be measured via the distance between its probability bounds and consistent ones; or it can be gauged via the extent to which the conditionals are far from being jointly satisfied by a probabilistic interpretation. The former strategy is reviewed in Section 3.2.1, and the latter, in Sections 3.2.2 and 3.3.

3.2.1 Measuring Inconsistency as Distance between Probability Bounds

Two major informal intuitions behind the degree of inconsistency of a knowledge base are:

- the effort to consolidate the base — that is, to restore consistency;
- the distance from the base to consistency — that is, to a consistent base.

[3] Technically, we could use the lexicographic order over the pairs $(\varphi_i|\psi_i)$ to construct a function Lex taking each set Γ to the corresponding sequence $\Psi = Lex(\Gamma)$, uniquely specifying a function Λ'_Ψ that changes the probability bounds of the sequence Ψ. Then it could be defined $\Lambda_\Gamma(q) = Lex^{-1}(\Lambda'_\Psi(q))$.

Measuring inconsistency through distances between bases naturally capture these ideas, if the effort to repair the inconsistency is also understood as the distance to a consistent base. Generally, if $d : \mathbb{K}_c \times \mathbb{K}_c \to [0, \infty)$ is a distance between knowledge bases, it straightforwardly induces an inconsistency measure $\mathcal{I} : \mathbb{K}_c \to [0, \infty)$, which computes the distance to the closest consistent base:

$$\mathcal{I}(\Gamma) : \min\{d(\Gamma, \Psi) \mid \Psi \in \mathbb{K}_c \text{ is consistent}\}.$$

When two bases, say $\Gamma, \Psi \in \mathbb{K}_c$, assign probability lower bounds to the same pairs of propositions, $\Gamma = \Lambda_\Gamma(q)$ and $\Psi = \Lambda_\Gamma(q')$, an obvious way to quantify the distance between them is through the vectors q, q'. On the one hand, vector spaces are well-equipped with principled formalisms to measure distances. On the other hand, such measures are naturally continuous, yielding inconsistency measures that satisfy (Continuity). Hence, the commonest approaches to quantifying the inconsistency of a probabilistic knowledge base $\Gamma = \Lambda_\Gamma(q)$ are via its distance to a consistent base $\Psi = \Lambda_\Gamma(q')$, understood as the distance between the vectors q, q'. This is possible because for any knowledge base $\Gamma = \Lambda_\Gamma(q)$, there is a vector of probability bounds $q' \in [0, 1]^\Gamma$ such that $\Lambda_\Gamma(q')$ is consistent — in the "worst" case, $q' = \langle 0, \ldots, 0\rangle$.

To review the commonest set of distances applied to this end in the literature (Muiño, 2011; Thimm, 2011, 2013), we define the p-norm of a vector $\langle q_1, \ldots, q_m\rangle \in \mathbb{R}^m$, for any $p \in \mathbb{N}_{>0}$ (a positive natural p), as being:

$$\|q\|_p = \sqrt[p]{\sum_{i=1}^{m} |q_i|^p}.$$

Taking the limit $p \to \infty$, we also define $\|q\|_\infty = \max_i |q_i|$. Now the p-norm distance between two vectors $q, q' \in \mathbb{R}^m$ can be defined as:

$$d_p(q, q') = \|q - q'\|_p.$$

For instance, d_1 is the Manhattan (or absolute) distance, d_2 is the Euclidean distance and d_∞ is the Chebyshev distance.

Using this family of distances in vector spaces, we define, for any p in $\overline{\mathbb{N}}_{>0}$ (which denotes $\mathbb{N}_{>0} \cup \{\infty\}$), an inconsistency measure $\mathcal{I}_p : \mathbb{K}_c \to [0, \infty)$ such that, for any $\Gamma \in \mathbb{K}_c$:

$$\mathcal{I}_p(\Gamma) = \min_{q' \in [0,1]^{|\Gamma|}} \{d_p(q, q') \mid \Gamma = \Lambda_\Gamma(q) \text{ and } \Lambda_\Gamma(q') \text{ is consistent}\}.$$

Note that, as the distance d_p is never negative, neither is $\mathcal{I}$.

For instance, consider the bases $\Delta = \{P(\bot) \geq 0.001\}$ and $\Psi = \{P(\bot) \geq 1\}$. Whereas $\mathcal{I}(\Delta) = \mathcal{I}(\Psi) = 1$ for any inconsistency measure $\mathcal{I}$ among $\mathcal{I}_d$, $\mathcal{I}_{SIM}$, $\mathcal{I}_{SIM^C}$ and $\mathcal{I}_\eta$, we have that $\mathcal{I}_p(\Delta) = \sqrt[p]{0.001^p} = 0.001 < 1 = \sqrt[p]{1^p} = \mathcal{I}_p(\Psi)$, for any $p \in \overline{\mathbb{N}}_{>0}$. In the following, a more substantial illustration of how these inconsistency measures work is provided:

Example 3.2.2. Consider the following probabilistic bases:

$$\begin{aligned}\Gamma &= \{P(x_1) \geq 0.9, P(x_2) \geq 0.5, P(\neg x_1 \vee \neg x_2) \geq 0.75\};\\ \Delta &= \{P(x_1|x_2) \geq 0.6, P(\neg x_1|x_2) \geq 0.6, P(x_2) \geq 0.9\}.\end{aligned}$$

To see that Γ is inconsistent, recall that, by the probability axioms, any probability measure P_π

must obey the following identity:

$$\begin{aligned} P_\pi(x_1) + P_\pi(x_2) &= P_\pi(x_1 \vee x_2) + P_\pi(x_1 \wedge x_2) \; ; \\ P_\pi(x_1 \wedge x_2) &= P_\pi(x_1) + P_\pi(x_2) - P_\pi(x_1 \vee x_2) \; . \end{aligned} \tag{3.4}$$

Any π satisfying $P(x_1) \geq 0.9$ and $P(x_2) \geq 0.5$ implies $P_\pi(x_1)+P_\pi(x_2) \geq 1.4$. Since $P_\pi(x_1 \vee x_2) \leq 1$, by Equation (3.4), it follows that $P_\pi(x_1 \wedge x_2) \geq 0.4$. However, this contradicts $P(\neg x_1 \vee \neg x_2) \geq 0.75$, which is equivalent to $P(x_1 \wedge x_2) \leq 0.25$, so Γ is inconsistent.

The base Δ is inconsistent because any π satisfying $P(x_1|x_2) \geq 0.6$ and $P(x_2) \geq 0.9$ is such that $P_\pi(x_1 \wedge x_2) \geq 0.54$ and $P_\pi(\neg x_1 \wedge x_2) \leq 0.46$, so $P_\pi(\neg x_1 \wedge x_2) - 0.6P_\pi(x_2) \leq -0.08 < 0$, violating $P(\neg x_1|x_2) \geq 0.6$.

Let $q = \langle 0.9, 0.5, 0.75\rangle$ and $r = \langle 0.6, 0.6, 0.9\rangle$ be vectors in $[0,1]^3$, so that $\Gamma = \Lambda_\Gamma(q)$ and $\Delta = \Lambda_\Delta(r)$. For any $p > 1$, the consistent bases $\Lambda_\Gamma(s)$ and $\Lambda_\Delta(t)$ that minimise, respectively, $d_p(q,s)$ and $d_p(r,t)$ are:

$$\begin{aligned} \Gamma' = \Lambda_\Gamma(s) &= \{P(x_1) \geq 0.85, P(x_2) \geq 0.45, P(\neg x_1 \vee \neg x_2) \geq 0.7\} \; ; \\ \Delta' = \Lambda_\Delta(t) &= \{P(x_1) \geq 0.5, P(\neg x_1) \geq 0.5, , P(x_2) \geq 0.9\} \; . \end{aligned}$$

When $p = 1$, there are several closest consistent bases $\Lambda_\Gamma(s')$ and $\Lambda_\Delta(t')$, including Γ' and Δ'. In other words, for any p-norm distance, $r = \langle 0.85, 0.45, 0.7\rangle$ is a closest vector to $q = \langle 0.9, 0.5, 0.75\rangle$ such that $\Lambda_\Gamma(r)$ is consistent; and $s = \langle 0.5, 0.5, 0.9\rangle$ a closest vector to $r = \langle 0.6, 0.6, 0.9\rangle$ such that $\Lambda_\Delta(s)$ is consistent. To obtain the corresponding inconsistency measurements, we only need to compute the p-norm distance between these vectors, $d_p(q,s)$ and $d_p(r,t)$. For example, for $p \in \{1, 2, \infty\}$:

$$\begin{aligned} \mathcal{I}_1(\Gamma) &= \sqrt[1]{|0.85-0.9|^1 + |0.45-0.5|^1 + |0.7-0.75|^1} = 0.05 + 0.05 + 0.05 = 0.15 \; ; \\ \mathcal{I}_2(\Gamma) &= \sqrt[2]{|0.85-0.9|^2 + |0.45-0.5|^2 + |0.7-0.75|^2} = \sqrt[2]{3 \times 0.05^2} = 0.05\sqrt[2]{3} \approxeq 0.087 \; ; \\ \mathcal{I}_\infty(\Gamma) &= \max |0.85-0.9|^1 + |0.45-0.5|^1 + |0.7-0.75|^1 = \max\{0.05, 0.05.0.05\} = 0.05 \; ; \\ \mathcal{I}_1(\Delta) &= \sqrt[1]{|0.5-0.6|^1 + |0.5-0.6|^1 + |0.9-0.9|^1} = 0.1 + 0.1 = 0.2 \; ; \\ \mathcal{I}_2(\Delta) &= \sqrt[2]{|0.5-0.6|^2 + |0.5-0.6|^2 + |0.9-0.9|^2} = \sqrt[2]{0.1^2 + 0.1^2} = 0.1\sqrt[2]{2} \approxeq 0.141 \; ; \\ \mathcal{I}_\infty(\Delta) &= \max\{|0.5-0.6| + |0.5-0.6| + |0.9-0.9|\} = \max\{0.1.0.1, 0\} = 0.1 \; . \end{aligned}$$

□

This family of inconsistency measures was proposed by Thimm (2013) and by Muiño (2011) for precise probabilities, although the latter assumes a different semantics (see Section 2.2.3). Thimm has proved all members of this family satisfy certain properties, which we adapt for the imprecise probabilities setting. [4] Below, these properties are summarised (Muiño (2011) has similar results, though under his different semantics.):

Theorem 3.2.3. *For any $p \in \overline{\mathbb{N}}_{>0}$, $\mathcal{I}_p$ is well-defined and satisfies the postulates of (Consistency), (Continuity) and (Monotonicity), but not (Independence).*

[4] Although Thimm (2013) claimed to have proved that any $\mathcal{I}_p$ satisfies (Independence) and $\mathcal{I}_1$ additionally satisfies (MIS-separability), we shall see in Section 3.4 that it is not the case.

Furthermore, special members of the family enjoy additional properties:

Lemma 3.2.4. *$\mathcal{I}_p$ satisfies (Super-additivity) iff $p = 1$.*

If (Normalisation) is required, we can prove the following result, adapted from (Muiño, 2011), on our semantics:

Lemma 3.2.5. *$\mathcal{I}_p$ satisfies (Normalisation) iff $p = \infty$.*

To compute $\mathcal{I}_p(\Gamma)$, for a given canonical base $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq m\}$, we have to find a consistent $\Lambda_\Gamma(r)$ while minimizing $d_p(q, r)$. Consider the $(2^n \times 1)$-vector π and the $(m \times 2^n)$-matrix $A = [a_{ij}]$, with $a_{ij} = I_{w_j}(\varphi_i \wedge \psi_i) - q_i I_{w_j}(\psi_i)$, from the system (2.4)-(2.6), which decides the satisfiability of Γ, where I_w is the indicator function of the set $\{\varphi \in \mathcal{L}_{X_n} | \, w \models \varphi\}$. Replacing each number q_i for a variable r_i in A, forming A', and adding an objective function $\min d_p(q, r) = \sqrt[p]{\sum_i^m |r_i - q_i|^p}$ [5], we have a program whose solution — the minimum value for the objective function — is $\mathcal{I}_p(\Gamma)$:

$$\min \sqrt[p]{\sum_i^m |r_i - q_i|^p} \qquad \text{subject to:} \tag{3.5}$$

$$A'\pi \geq 0 \tag{3.6}$$

$$\sum \pi = 1 \tag{3.7}$$

$$\pi \geq 0 \, . \tag{3.8}$$

Note that within $A'\pi$ there are now quadratic terms of the form $r_i I_{w_j}(\psi_i)\pi_j$, for π_j and r_i are real-valued variables. This leads to a non-convex search space, making the computation rather hard.

The minimisation above can be rewritten to ease the presentation. We can avoid the possibility of a probability lower bound being raised, $r_i > q_i$, without alter the minimum feasible value of the objective function. If a feasible solution on r is such that $r_i > q_i$, for some i, taking $r_i = q_i$ yields a feasible solution with no greater value for the objective function — actually, the value of the objective function would decrease if $p < \infty$. Hence, we can assume $r_i \leq q_i$ and rewrite (3.6)-(3.8) using the substitutions $r_i = q_i - \varepsilon_i$, adding non-negative real-valued variables $\varepsilon_1, \ldots, \varepsilon_m$:

$$\min \sqrt[p]{\sum_i^m \varepsilon_i^p} \qquad \text{subject to:} \tag{3.9}$$

$$P_\pi(\varphi_i \wedge \psi_i) - (q_i - \varepsilon_i)P_\pi(\psi_i) \geq 0 \text{ for all } 1 \leq i \leq m \tag{3.10}$$

$$\sum \pi = 1 \tag{3.11}$$

$$\pi \geq 0 \tag{3.12}$$

$$\varepsilon_i \geq 0 \text{ for all } 1 \leq i \leq m \tag{3.13}$$

Now the quadratic terms come from $\varepsilon_i P_\pi(\psi_i)$, since $P_\pi(\psi_i) = \sum_{j=1}^m \pi_j I_{w_j}(\psi_i)$. Muiño (2011) points out that there is no effective method to solve these problems, even though there are approximate

[5] The pth-root within the objective function is a monotone function and may be ignored.

methods that can handle small instances; for details on these issues, see for instance (Bertsekas , 1999) and (Boyd and Vandenberghe, 2004). For $p = 1$, Batsell *et al.* (2002) and Thimm (2011) provide approximations via linear programs, but losing precision of course.

3.2.2 Measuring Inconsistency as the Distance to being Satisfied

Potyka (2014) emphasises the fact that the optimisation problems one need to solve in order to compute to compute $\mathcal{I}_p$ have (non-global) local optima, so convex optimisation techniques cannot be directly applied. Thus, he argues that computing $\mathcal{I}_p$ is typically less efficient than deciding PSAT, as empirical results indicate (Potyka, 2014). Focusing on computational efficiency, Potyka proposes a slight modification in (3.9)-(3.13) to quantify inconsistency as the extent to which the base is far from being satisfied by a probabilistic interpretation.

Potyka's idea is to make the restrictions in (3.10) linear, as Nau (1981) similarly had put forward. To accomplish this, "violation" variables $\varepsilon_i \geq 0$ are inserted in the right-hand side of $P_\pi(\varphi_i \wedge \psi_i) - q_i P_\pi(\psi_i) \geq 0$, yielding $P_\pi(\varphi_i \wedge \psi_i) - q_i P_\pi(\psi_i) \geq -\varepsilon_i$. Using the the vector π and the matrix A from (2.4), where $a_{ij} = I_{w_j}(\varphi_i \wedge \psi_i) - q_i I_{w_j}(\psi_i)$, and an $(m \times 1)$-vector $\varepsilon = [\varepsilon_1 \ldots \varepsilon_m]^T$, the following program returns a probabilistic interpretation (represented by π) that is the closest to satisfy $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq m\}$ in a certain sense:

$$\min \sqrt[p]{\sum_i^m \varepsilon_i^p} \qquad \text{subject to:} \tag{3.14}$$

$$A\pi \geq -\varepsilon \tag{3.15}$$

$$\sum \pi = 1 \tag{3.16}$$

$$\pi, \varepsilon \geq 0 \tag{3.17}$$

In a feasible solution π, ε for the optimisation problem (3.9)-(3.13), each ε_i corresponds to how much a probability bound q_i is being relaxed for the corresponding conditional $P(\varphi_i|\psi_i) \geq q_i$ to be satisfied by the probabilistic interpretation π. With Potyka's modification, in a feasible solution π, ε for (3.14)-(3.17), each ε_i corresponds to a tolerated error (or violation) in the satisfaction of that conditional by π. In other words, if π, ε is a feasible solution to (3.14)-(3.17), π satisfies $P(\varphi_i|\psi_i) \geq q_i$ with error, or violation, at most ε_i. The degree of inconsistency of a given base is then defined as minimum value of the objective function (3.14), which is the p-norm of these violations.

Formally, recall that Π_n is the set of all probabilistic interpretations $\pi : W_n \to [0, 1]$ and consider, for any $p \in \mathbb{N}_{>0} \cup \{\infty\}$, a function $d_\varepsilon^p : \mathbb{K}_c \times \Pi_n \to [0, \infty)$ defined as:

$$\begin{aligned} d_p^\varepsilon(\Gamma, \pi) &= \|\langle \varepsilon_1, \ldots, \varepsilon_m \rangle\|_p, \quad \text{where} \\ \varepsilon_i &= \max\{0, q_i P_\pi(\psi_i) - P_\pi(\varphi_i \wedge \psi_i)\} \text{ for all } 1 \leq i \leq m\,. \end{aligned}$$

In other words, $d_p^\varepsilon(\Gamma, \pi)$ is the p-norm of the violations by π of the restrictions corresponding to the conditionals in Γ. Alternatively, $d_p^\varepsilon(\Gamma, \pi)$ could be defined as the solution to (3.14)-(3.17) for a fixed probability mass π. We say $d_p^\varepsilon(\Gamma, \pi)$ is a *discrepancy* between a base Γ and a probabilistic interpretation π.

Note that the program (3.14)-(3.17) computes the probabilistic interpretation π that minimises

$d_p^\varepsilon(\Gamma,\pi)$ for a fixed Γ. Now Potyka's family of inconsistency measures, $\mathcal{I}_p^\varepsilon : \mathbb{K}_c \to [0,\infty)$, can be formally defined for all $p \in \mathbb{N}_{>0} \cup \{\infty\}$ and $\Gamma \in \mathbb{K}_c$:

$$\mathcal{I}_p^\varepsilon(\Gamma) = \min\{d_p^\varepsilon(\Gamma,\pi) | \pi \in \Pi_n\}\,.$$

Potyka originally presented these measures in the precise probabilities context, calling each member of this family a *minimal violation measure*. It is important to note that, for $1 < p < \infty$, any π minimising $d_p^\varepsilon(\Gamma,\pi)$ corresponds to the same violations vector (Potyka, 2014).

Example 3.2.6. Consider again the bases from Example 3.2.2:

$$\begin{aligned}\Gamma &= \{P(x_1) \geq 0.9, P(x_2) \geq 0.5, P(\neg x_1 \vee \neg x_2) \geq 0.75\}\,;\\ \Delta &= \{P(x_1|x_2) \geq 0.6, P(\neg x_1|x_2) \geq 0.6, P(x_2) \geq 0.9\}\,.\end{aligned}$$

For an arbitrary p, $\mathcal{I}_p^\varepsilon$ is obtained through a probabilistic interpretation π minimising the p-norm of the violations ε_i. To measure the inconsistency of Γ, we need to find those ε_i's. For each $2 \leq p < \infty$, and exceptionally also for $p = \infty$, any π minimising $d_p^\varepsilon(\Gamma,\pi)$ yield the same violations $\varepsilon_1, \varepsilon_2, \varepsilon_3$. Furthermore, in this case these violations vectors coincide for all $p \geq 2$:

$$\langle \varepsilon_1, \varepsilon_2, \varepsilon_3 \rangle = \langle 0.05, 0.05, 0.05 \rangle\,.$$

If $p = 1$, the minimisation of $d_1^\varepsilon(\Gamma,\pi)$ has multiple solutions on ε, all with the same 1-norm, including $\langle \varepsilon_1, \varepsilon_2, \varepsilon_3 \rangle = \langle 0.05, 0.05, 0.05 \rangle$. Hence $I_p^\varepsilon(\Gamma)$ can be expressed in the following general form:

$$\mathcal{I}_p^\varepsilon(\Gamma) \quad = \quad \sqrt[p]{0.05^p + 0.05^p + 0.05^p} = 0.05\sqrt[p]{3}\,.$$

For some values of p, we obtain for instance:

$$\begin{aligned}\mathcal{I}_1^\varepsilon(\Gamma) &= 0.15\,;\\ \mathcal{I}_2^\varepsilon(\Gamma) &= 0.087\,;\\ \mathcal{I}_3^\varepsilon(\Gamma) &= 0.072\,;\\ \mathcal{I}_\infty^\varepsilon(\Gamma) &= 0.05\,.\end{aligned}$$

While measuring the inconsistency of Δ, for a fixed p, any π minimising $d_p^\varepsilon(\Delta,\pi)$ has the same violation vector. However, each p yields different values for these violations. For instance, for $p \in \{1, 2, 3, \infty\}$:

$$\begin{aligned}p = 1: \quad \langle \varepsilon_1, \varepsilon_2, \varepsilon_3 \rangle &= \langle 0.09000, 0.09000, 0 \rangle\,;\\ p = 2: \quad \langle \varepsilon_1, \varepsilon_2, \varepsilon_3 \rangle &= \langle 0.08824, 0.08824, 0.01765 \rangle\,;\\ p = 3: \quad \langle \varepsilon_1, \varepsilon_2, \varepsilon_3 \rangle &= \langle 0.08615, 0.08615, 0.03853 \rangle\,;\\ p = \infty: \quad \langle \varepsilon_1, \varepsilon_2, \varepsilon_3 \rangle &= \langle 0.08182, 0.08182, 0.08182 \rangle\,.\end{aligned}$$

Taking the respective p-norm of the vectors above, we obtain the inconsistency measurements:

$$\begin{aligned}
\mathcal{I}_1^{\varepsilon}(\Delta) &= 0.180\,; \\
\mathcal{I}_2^{\varepsilon}(\Delta) &= 0.126\,; \\
\mathcal{I}_3^{\varepsilon}(\Delta) &= 0.110\,; \\
\mathcal{I}_{\infty}^{\varepsilon}(\Delta) &= 0.050\,.
\end{aligned}$$

□

Looking at the programs (3.9)-(3.13) and (3.14)-(3.17), which computes $\mathcal{I}_p$ and $\mathcal{I}_p^{\varepsilon}$, respectively, it is easy to see what happens when conditioning events (ψ_i) are tautological (implying $P_\pi(\psi_i) = 1$):

Proposition 3.2.7. *For any unconditional probabilistic knowledge base $\Gamma \in \mathbb{K}_c$ and any $p \in \overline{\mathbb{N}}_{>0}$, $\mathcal{I}_p^{\varepsilon}(\Gamma) = \mathcal{I}_p(\Gamma)$.*

Therefore, while measuring the inconsistency of unconditional bases, $\mathcal{I}_p^{\varepsilon}$ inherits the properties of $\mathcal{I}_p$. Indeed, we can generalise some properties proved by Potyka to the imprecise probabilities scenario, proving that $\mathcal{I}_p^{\varepsilon}$ satisfies the same basic properties as $\mathcal{I}_p$:

Theorem 3.2.8. *For any $p \in \overline{\mathbb{N}}_{>0}$, $\mathcal{I}_p^{\varepsilon} : \mathbb{K} \to [0, \infty)$ is well-defined and satisfies (Consistency), (Continuity) and (Monotonicity), but not (Independence). $\mathcal{I}_p^{\varepsilon}$ satisfies (Super-additivity) iff $p = 1$; and $\mathcal{I}_p^{\varepsilon}$ satisfies (Normalisation) iff $p = \infty$.*

Back to computational aspects, the improvement of having linear restrictions is clear. In general, programs like (3.14)-(3.17) correspond to convex optimisation problems, and various approved algorithms can tackle them (Boyd and Vandenberghe, 2004). For $p = 1$, the program (3.14)-(3.17) is linear, as the objective function becomes linear — and can be written as $\sum \varepsilon_i$. Moreover, when $p = \infty$, the program above is equivalent to a linear one with a single violation variable $\varepsilon_{max} \geq 0$, which represents the maximum violation. To rewrite the program (3.14)-(3.17) to the case $p = \infty$, let ε denote the vector $[\varepsilon_{max} \ldots \varepsilon_{max}]^T$:

$$\min \varepsilon_{max} \qquad \text{subject to:} \tag{3.18}$$

$$A\pi \geq -\varepsilon \tag{3.19}$$

$$\sum \pi = 1 \tag{3.20}$$

$$\pi, \varepsilon \geq 0 \tag{3.21}$$

When $p = 1$ or $p = \infty$, the resulting linear program can be solved with the same efficient techniques applied to decide PSAT, such as Simplex with column generation. In practice, these two particular cases are no harder than deciding PSAT (Potyka, 2014). Potyka (2014) argues in favour of these two inconsistency measures, for they can be called "feasible" in comparison to $\mathcal{I}_p$, even both computations being NP-hard.

Although the inconsistency measure $\mathcal{I}_p^{\varepsilon}$ may seem rather artificial at first, we show in Chapter 7 that it can, in fact, be given an interpretation with probabilities inducing prejudicial betting behaviour (via Dutch books) when $p = 1$ and $p = \infty$. Moreover, $\mathcal{I}_1^{\varepsilon}$ and $\mathcal{I}_{\infty}^{\varepsilon}$ were independently proposed in Bayesian statistics, under this cover of guaranteed losses in Dutch books, to measure

the incoherence of subjective probability assessments (Schervish *et al.*, 2002b), as we shall see in Chapter 7.

3.3 Related Work

In Statistics, a divergence is a function that quantifies the difference between two probability distributions. Distances, as the p-norm ones, can be seen as special cases of divergences, for the latter need not satisfy symmetry or triangle inequality. A widely-used divergence measure between two probability distributions q and q' is the Kullback-Leibler divergence (Kullback, 1997), that captures the expected information loss for using q when q' is the "actual" distribution — the one describing the objective chances, so to speak. In the probabilistic logic setting, this divergence can be defined over probabilistic interpretations $\pi : W_n \to [0, 1]$, which can be understood as probability distributions for the actual world. Firstly, we define the *entropy*, or the expected information, carried by a probability mass $\pi : W_n \to [0, 1]$ — for more on entropy, see (Shannon and Weaver, 1949) [6]:

$$H(\pi) = -\sum_{w \in W_n} \pi(w) \ln \pi(w) \,. \tag{3.22}$$

Given a probability mass π, the *information content* of realizing that the actual world is some particular $w \in W_n$ can be defined via the scoring rule $S_{inf}(\pi, w) = -\ln \pi(w)$ — the logarithm base is irrelevant, as a multiplicative constant. Hence, $H(\pi) = E_\pi(S_{inf}(\pi, w))$ is the expected value (according to π) of the amount of information carried by π. While moving from the probabilistic interpretation π_2 to π_1, one can measure the loss in the information content when realizing the actual world as some $w \in W_n$: $S_{inf}(\pi_2, w) - S_{inf}(\pi_1, w)$. Under π_1, one can compute the expected value of such information loss, which is exactly the Kullback-Leibler divergence of π_2 from π_1:

$$KL(\pi_1, \pi_2) = \sum_{w \in W_n} \pi_1(w) \ln \frac{\pi_1(w)}{\pi_2(w)} \,. \tag{3.23}$$

Kullback-Leibler divergence cannot be directly applied to vectors of probability bounds in order to measure the inconsistency of an arbitrary probabilistic knowledge base, due to the fact that these probability bounds do not form a probability mass. Nevertheless, it can be adapted to measure a "discrepancy" from a set of precise conditionals and a probabilistic interpretation. Capotorti *et al.* (2010) adapt this divergence to cope with precise conditional probability assessments under the coherence setting (see Section 2.2.3). Their main goal is to develop procedures to restore coherence, but it is done through minimising what can be understood as an inconsistency measure. In the following, we summarise the approach of Capotorti and Regoli (2008), also found in (Capotorti *et al.* , 2009, 2010), focusing on measuring inconsistency, and not on repairing incoherence.

Let $\Gamma = \{P(\varphi_i|\psi_i) = q_i | 1 \leq i \leq m\}$ be a precise probabilistic knowledge base, with $q = \langle q_1, \ldots, q_m \rangle$ in $(0, 1)^m$ [7] for $1 \leq i \leq m$. The approach of Capotorti *et al.* (2010) is based on the

[6] In such expressions, $0 \log 0 = 0$ by definition, respecting $\lim_{p \to 0^+} p \ln p = 0$.

[7] Capotorti *et al.* (2010) avoid extreme probabilities (0 or 1) for technical reasons.

following scoring rule, adapted from (Lad, 1996):

$$S_{CRV}(q,w) = \sum_{i=1}^{m} I_w(\varphi_i \wedge \psi_i) \ln q_i + \sum_{i=1}^{m} I_w(\neg\varphi_i \wedge \psi_i) \ln(1-q_i)\,. \tag{3.24}$$

Based on this scoring rule, the authors proceed to introduce a discrepancy between the vector of probabilities q, fixed a base Γ, and a probabilistic interpretation π such that $\pi(\bigvee_i \psi_i) = 1$. We choose instead to use the whole base Γ as an argument to the discrepancy, to avoid confusion, and drop that restriction on π, due to our semantics. To define such discrepancy for precise knowledge bases $\Gamma = \{P(\varphi_i|\psi_i) = q_i | 1 \leq i \leq m\}$, we use q^π to denote the vector $\langle \frac{P_\pi(\varphi_1 \wedge \psi_1)}{P_\pi(\psi_1)}, \ldots, \frac{P_\pi(\varphi_m \wedge \psi_m)}{P_\pi(\psi_m)} \rangle \in [0,1]^m$:

$$d_{CRV}(\Gamma,\pi) = E_\pi\Big(S_{CRV}(q^\pi,w) - S_{CRV}(q,w)\Big) = \sum_{1\leq i\leq m, P_\pi(\psi_i)>0} P_\pi(\psi_i)\Big(q_i^\pi \ln \frac{q_i^\pi}{q_i} + (1-q_i^\pi)\ln\frac{1-q_i^\pi}{1-q_i}\Big)\,. \tag{3.25}$$

This discrepancy directly yields an inconsistency measure for precise probabilistic knowledge bases:

$$\mathcal{I}^P_{CRV}(\Gamma) \quad = \quad \min\{d_{CRV}(\Gamma,\pi) | \pi \in \Pi_n\}\,.$$

To adapt this measure to general probabilistic knowledge base, we use ideas from Capotorti *et al.* (2009), who introduced a method to correct incoherent imprecise probabilities. In the same way as imprecise probabilities may be construed as sets of precise probabilities, we can view a conditional $P(\varphi|\psi) \geq q$ as a set of precise probability assessments, $\{P(\varphi|\psi) = r | r \in [q,1]\}$. Extending this analogy, one can think of a probabilistic base $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq m\}$ as a set of precise bases:

$$S_\Gamma = \big\{\{P(\varphi_i|\psi_i) = r_i | 1 \leq i \leq m\} | r_i \in [q_i, 1] \text{ for } 1 \leq i \leq m\big\}\,.$$

With this in mind, the inconsistency of a general base Γ can be viewed as the inconsistency of the set of precise bases S_Γ, taking its least inconsistent element. The rationale is that the "distance" from S_Γ to consistency is defined through its element that is the closest to consistency. To formalise this notion, let $\Lambda^P_\Gamma : [0,1]^{|\Gamma|} \to \mathbb{K}$ be a function that, besides changing the probabilities in a base, makes the assessments become precise; *i.e.*, if $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq m\}$, then $\Lambda^P_\Gamma(q') = \{P(\varphi_i|\psi_i) = q'_i | 1 \leq i \leq m\}$. Now the inconsistency measure $\mathcal{I}_{CRV} : \mathbb{K}_c \to [0,\infty)$ can be defined:[8]

$$\mathcal{I}_{CRV}(\Gamma) \quad = \quad \min\{\mathcal{I}^P_{CRV}(\Lambda^P_\Gamma(r)) | r \in [0,1]^{|\Gamma|}, \Gamma = \Lambda_\Gamma(q), r \geq q\}\,.$$

In other words, if Γ assigns probability lower bounds q, $\mathcal{I}_{CRV}(\Gamma)$ is defined through the base Ψ assigning precise probabilities r, with $r \geq q$, and minimising $\mathcal{I}^P_{CRV}(\Psi)$. Without using $\mathcal{I}^P_{CRV}$, $\mathcal{I}_{CRV}$ can be directly be defined as:

[8] We allow extreme probabilities ($q_i = 0$ and $q_i = 1$) due to the fact that, for the minimisation of d_{CRV} to reach a finite value, π must be such that $q_i^\pi = 0$ (or $q_i^\pi = 1$), yielding expressions $0 \ln 0$, defined as 0.

$$\mathcal{I}_{CRV}(\Gamma) = \min\{d_{CRV}(\Lambda_\Gamma^P(r), \pi) \mid \pi \in \Pi_n, r \in [0,1]^{|\Gamma|}, \Gamma = \Lambda_\Gamma(q), r \geq q\}\,.$$

Example 3.3.1. Consider again the bases from Example 3.2.2:

$$\begin{aligned}\Gamma &= \{P(x_1) \geq 0.9, P(x_2) \geq 0.5, P(\neg x_1 \vee \neg x_2) \geq 0.75\}\,.\\ \Delta &= \{P(x_1|x_2) \geq 0.6, P(\neg x_1|x_2) \geq 0.6, P(x_2) \geq 0.9\}\,.\end{aligned}$$

Solving the optimisation problem that computes $\mathcal{I}_{CRV}$, we obtain:

$$\begin{aligned}\mathcal{I}_{CRV}(\Gamma) &= 0.021\,;\\ \mathcal{I}_{CRV}(\Delta) &= 0.037\,.\end{aligned}$$

□

Adapting results presented in (Capotorti and Regoli, 2008; Capotorti *et al.*, 2010) for d_{CRV}, we can prove some properties of $\mathcal{I}_{CRV}$:

Theorem 3.3.2. *$\mathcal{I}_{CRV}$ is well-defined and satisfies the postulates of (Consistency), (Continuity), (Monotonicity), but not (Independence).*

To compute $\mathcal{I}_{CRV}$, one has to solve a program with linear constraints, yielding a convex search space (Capotorti *et al.*, 2010), but with logarithms within the objective function:

$$\min d_{CRV}(\Lambda_\Gamma^P(r), \pi) \quad \text{subject to:} \tag{3.26}$$

$$r \geq q \tag{3.27}$$

$$\pi \geq 0 \tag{3.28}$$

$$\sum_{j=1}^{2^n} \pi_j = 1 \tag{3.29}$$

The approach above can be extended to any divergence generated by a proper scoring rule, always producing inconsistency measures satisfying (Consistency). Here we focus on d_{CRV} following the work of Capotorti and Regoli (2008), where the reader can found different divergences based on proper scoring rules — for divergences and scoring rules, see for instance (Gneiting and Raftery , 2007).

3.4 The Incompatibility of Postulates

Thimm (2013) has done a foundational work, extending Hunter and Konieczny's postulates for inconsistency measures to the probabilistic case, adding (Continuity), and proposing a whole family of distance-based measures to satisfy them. However, it happens that no inconsistency measure presented in Section 3.2 satisfies (Independence). In fact, this is not just a coincidence: assuming (Consistency), the postulate of (Continuity) precludes (Independence) from being satisfied (De Bona and Finger, 2015):

Theorem 3.4.1. *There is no inconsistency measure $\mathcal{I} : \mathbb{K}_c \to [0, \infty)$ that satisfies (Consistency), (Independence) and (Continuity).*

Proof. To prove by contradiction, suppose there is a measure $\mathcal{I}$ satisfying consistency, independence and continuity. Consider the following knowledge bases:

$$\Gamma = \{P(x_1 \wedge x_2) \geq 0.5 + \varepsilon, P(x_1 \wedge \neg x_2) \geq 0.5\} \text{ for some } 0 < \varepsilon \leq 0.1 \tag{3.30}$$

$$\Delta = \Gamma \cup \{\alpha\}, \alpha = P(\neg x_1) \geq 0.2 \tag{3.31}$$

We are going to use $\mathcal{I}$ to measure the inconsistency of Δ when $\varepsilon \to 0$. To apply (Independence), we are going to show that α is free in Δ; we prove that Γ is the only MIS in Δ. Note that $\{P(x_1 \wedge x_2) \geq 0.5 + \varepsilon, P(\neg x_1) \geq 0.2\}$ is consistent for any $\varepsilon \in (0, 0.1]$, for such set is satisfied by the probability measure induced by the following probability mass: $\pi_1(x_1 \wedge x_2) = 0.5 + \varepsilon, \pi_1(x_1 \wedge \neg x_2) = 0.3 - \varepsilon, \pi_1(\neg x_1 \wedge x_2) = \pi_1(\neg x_1 \wedge \neg x_2) = 0.1$. To prove that $\{P(x_1 \wedge \neg x_2) \geq 0.5, P(\neg x_1) \geq 0.2\}$ is consistent, consider the following probability mass: $\pi_2(x_1 \wedge x_2) = 0.3, \pi_2(x_1 \wedge \neg x_2) = 0.5, \pi_2(\neg x_1 \wedge x_2) = \pi_2(\neg x_1 \wedge \neg x_2) = 0.1$. Hence, all MISes of Δ must contain $\Gamma = \{P(x_1 \wedge x_2) \geq 0.5 + \varepsilon, P(x_1 \wedge \neg x_2) \geq 0.5\}$, for the other subsets are all consistent. Furthermore, note that Γ is inconsistent and minimally so, therefore it is a MIS. We can conclude that Γ is the only MIS in Δ, for any value of $0 < \varepsilon \leq 0.1$. As α is a free probabilistic conditional in Δ, we can employ (Independence):

$$\mathcal{I}(\Delta) = \mathcal{I}(\Gamma)\,,$$

for any $0 < \varepsilon \leq 0.1$.

To exploit the fact that $\mathcal{I}$ satisfies (Continuity), we need the characteristic function of Δ, $\Lambda_\Delta : [0, 1]^3 \to \mathbb{K}_c$, to be well-defined; so, we need an order over the probabilistic conditionals. Suppose that Γ and Δ are ordered as they were defined in (3.30) and (3.31). Let q^* be the vector $\langle 0.5, 0.5, 0.2\rangle$. It follows that $\Lambda_\Delta(q^*)$ differs from Δ only in its first conditional, which becomes $P(x_1 \wedge x_2) \geq 0.5$. Now we prove that $\Lambda_\Delta(q^*)$ is inconsistent. For any probability measure P_π, $P_\pi(x_1 \wedge x_2) \geq 0.5$ and $P_\pi(x_1 \wedge \neg x_2) \geq 0.5$ imply $P_\pi(x_1) = 1$, contradicting $\alpha = P(\neg x_1) \geq 0.2 = P(x_1) \leq 0.8$. As $\mathcal{I}$ satisfies (Consistency),

$$\mathcal{I} \circ \Lambda_\Delta(q^*) > 0. \tag{3.32}$$

By (Continuity), the function $\mathcal{I} \circ \Lambda_\Delta : [0, 1]^3 \to [0, \infty)$ must be continuous, so there must be a limit at the point q^*, and such limit must be unique for any path approaching q^*:

$$\lim_{q \to q^*} \mathcal{I} \circ \Lambda_\Delta(q) = \lim_{\varepsilon \to 0^+} \mathcal{I} \circ \Lambda_\Delta(\langle 0.5 + \varepsilon, 0.5, 0.2\rangle) = \lim_{\varepsilon \to 0^+} \mathcal{I}(\Delta)\,.$$

By (Independence), we also have:

$$\lim_{\varepsilon \to 0^+} \mathcal{I}(\Delta) = \lim_{\varepsilon \to 0^+} \mathcal{I}(\Gamma)\,.$$

As $\mathcal{I}$ satisfies (Continuity) and $\{P(x_1 \wedge x_2) \geq 0.5, P(x_1 \wedge \neg x_2) \geq 0.5\}$ is satisfiable, (Consistency) implies

$$\lim_{\varepsilon \to 0^+} \mathcal{I}(\Gamma) = \mathcal{I}(\{P(x_1 \wedge x_2) \geq 0.5, P(x_1 \wedge \neg x_2) \geq 0.5\}) = 0 = \lim_{q \to q^*} \mathcal{I} \circ \Lambda_\Delta(q)\,. \tag{3.33}$$

The continuity of $\mathcal{I}$ requires that $\mathcal{I} \circ \Lambda_\Gamma(q^*) = \lim_{q \to q^*} \mathcal{I} \circ \Lambda_\Gamma(q)$, which by (3.32) and (3.33) is a contradiction, finishing the proof. □

Corollary 3.4.2. *There is no inconsistency measure $\mathcal{I} : \mathbb{K}_c \to [0, \infty)$ that satisfies (Consistency), (MIS-separability) and (Continuity).*

In Thimm's work, the compatibility of these postulates is implicitly stated when it is proved that the whole family of inconsistency measures $\mathcal{I}_p$ satisfies them; and another family is proved to enjoy (MIS-Separability). Actually, Thimm (2013) investigated the particular case of precise probabilistic bases, with a slightly different version of (Continuity). Nevertheless, it is easy to adapt the proof of Theorem 3.4.1 to the precise probabilities framework, as it is done in (De Bona and Finger, 2015).

These incompatibility results suggest that in order to drive the rational choice of an inconsistency measure for probabilistic knowledge bases, we must abandon at least one postulate among (Consistency), (Independence) and (Continuity). In this work, we look for a weakening of the desired properties that could restore their compatibility and investigate paths for achieving that goal. Firstly, we need to decide which postulate to tackle.

The (Consistency) postulate seems to be indisputable since the least one can expect from an inconsistency measure is that it separates inconsistent from consistent cases, or some inconsistency from none. The answer to the question of which property we should relax to restore compatibility is thus reduced to either (Independence) or (Continuity). Hunter and Konieczny (2010) have already noted problems with (Independence) in knowledge bases over classical logic, proposing a weaker version, as we shall see in Section 6.2. Intuition shall be inclined towards keeping (Continuity), for it reflects the particular quantitative nature of probabilistic reasoning. A pragmatic reason to give up (Independence) — and so (MIS-separability) — is simply to keep (Continuity), given (Consistency), to save the (continuous) inconsistency measures seen in Section 3.2. Furthermore, it seems that the definition of free conditional, and so (Independence), can be refined to be suitable for analysing continuous measures, while (Continuity) is a harder definition to be contrived to be compatible with (Independence). In fact, the withdrawal of (Independence) within probabilistic logic can be argued for in a more compelling way, by analysing the relation between consolidation procedures and the characterisation of primitive conflicts (as MISes), what is done in the next chapters.

Chapter 4

Consolidating Probabilistic Knowledge Bases

While presenting, in the previous chapter, the inconsistency measures found in the literature and the rationality postulates that drive their formulation, we came across an incompatibility result. As (Consistency), (Independence) and (Continuity) are not jointly satisfiable, we intend to find positive reasons to abandon (Independence), tracing back its link to minimal inconsistent sets and searching for paths to consistently rewriting it. The analysis of the classical method of restoring consistency (or *consolidating*), in which are grounded the inconsistency measures for classical logic, shall reveal why (Independence) is stated in terms of minimal inconsistent sets and free conditionals. Nonetheless, we show that the continuous inconsistency measures proper to probabilistic bases are underpinned by their own consolidation procedures, tailored to probabilities as well. Consequently, in Chapter 6, (Independence) via minimal inconsistent sets will be shown to be unjustifiable in the probabilistic context, and the probabilistic consolidation strategies shall guide the contrivance of new versions of this postulate.

Presenting the consolidation methods also brings other benefits to understanding inconsistency measuring. A major intuition behind the severity of the inconsistency in a given knowledge base is the "effort" required to restore its consistency — to *consolidate* it. This effort can be understood as a quantification of the changes in the base needed to achieve consistency. In this sense, to measure inconsistency as the extent to which a base has to be modified in order to be consolidated, we must assume an *underlying consolidation procedure*. Alternatively, the inconsistency of a base can be viewed as how "far" it is from being satisfied by an interpretation, and the search for this "closest" interpretation yields consolidation methods as well.

This chapter is organised as follows: Section 4.1 brings a general introduction on consolidations methods for probabilistic knowledge bases and presents Sections 4.2, 4.3 and 4.4, which details those methods; related work regarding coherence checking is reviewed in Section 4.5; Section 4.6 sums up this chapter's conclusions and indicates our path towards reconciling the postulates for inconsistency measures.

4.1 Consolidation Operators and Strategies

Informally, a consolidation operator is a function that takes possibly inconsistent probabilistic knowledge bases and returns consistent ones. Formally, a *consolidation operator* is a function $\mathcal{C} : \mathbb{K}_c \to \mathbb{K}_c$ satisfying two rationality requirements, adapted from Potyka and Thimm (2014):

Postulate 4.1.1 (**Success**). *For any* $\Gamma \in \mathbb{K}_c$, $\mathcal{C}(\Gamma)$ *is consistent.*

Postulate 4.1.2 (**Vacuity**). *For any consistent* $\Gamma \in \mathbb{K}_c$, $\mathcal{C}(\Gamma) = \Gamma$.

The postulate of (Vacuity), called "Consistency" by Potyka and Thimm (2014), demands that the consistency restoration not modify a consistent base. This is a special case of the *minimum change* desiderata, which will be formalised later for consolidation via probabilities changing.

Given an inconsistent canonical base $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq m\}$, the strategies to consolidate it, within the language $\mathcal{L}^P_{X_n}$, can be classified into three groups:

- **Removing Formulas**: $\mathcal{C}(\Gamma)$ is a subset of Γ;
- **Changing Probability Bounds**: $\mathcal{C}(\Gamma) = \{P(\varphi_i|\psi_i) \geq q'_i | 1 \leq i \leq m\}$;
- **Other Methods**: $\mathcal{C}(\Gamma) = \{P(\varphi'_i|\psi'_i) \geq q'_i | 1 \leq i \leq m'\}$.

The first method applies to any logic and is the reason for minimal inconsistent sets to be usually viewed as the pure causes of inconsistencies — a position we argue against in Chapter 6. Hence, to understand the reasons behind the postulate of (Independence), we investigate this approach in Section 4.2, under the well-established framework of belief revision.

The second method, discussed in Section 4.3, is the commonest in a probabilistic setting and underlies the inconsistency measure $\mathcal{I}_p$. Besides that, such approach to repair inconsistency interests us for several other reasons. Firstly, removing a conditional $P(\varphi|\psi) \geq q$ from a base is equivalent to changing it to $P(\varphi|\psi) \geq 0$, a tautology, and consolidation via formulas discarding is just an extreme case of consolidation via probability bounds changing. Secondly, it is arguably more intuitive for a natural agent who holds inconsistent probabilistic statements to fix them by modifying the numeric values than by changing the propositions the probabilities are assigned to. Finally, changes in probability bounds can be measured using distances in vector spaces, yielding well-founded implementations of the minimum change notion. We also refer to this method of consolidation, and a base resulting from it, as *quantitative consolidation*, due to its numeric nature.

Fixed the language $\mathcal{L}^P_{X_n}$, the remaining methods can be grouped in two classes: methods that merge the first two approaches; and those in which the consolidated base $\mathcal{C}(\Gamma)$ contains conditionals $P(\varphi|\psi) \geq q$ such that there is no $P(\varphi|\psi) \geq q'$ in Γ. The latter are rarely (if ever) seen, probably for the second and third reasons listed above. An example of the former can be found in (Finthammer *et al.*, 2007), but, being a pragmatic, heuristic approach, it is out of the scope of our work.

If we allow the consolidated base to be in a more expressive language than the original one, new strategies for restoring consistency can be devised. Although it seems not sensible at first to change the underlying language while consolidating, in Section 4.4 we discuss such an approach, showing its correspondence to Potyka's minimal violation measures $\mathcal{I}^\varepsilon_p$. Departing from a $\mathcal{I}^\varepsilon_p$-based consolidation procedure from the literature, we show how an extension of the logical language enables us to found that method on the AGM approach, in Chapter 5, and to derive a corresponding definition of primitive (or atomic) inconsistent set, in Chapter 6.

4.2 Removing Formulas

The classical way of handling inconsistency in bases through ruling out formulas was proposed, in Computer Science, by Reiter (1987) in his diagnosis problem and is also the basis for consolidation in the AGM paradigm (named after Alchourrón, Gärdenfors, and Makinson (1985)) of belief revision (see Hansson (1999) for a general view of the AGM paradigm for belief bases). Reiter's *hitting sets* technique views the consolidation of a base as the discarding of at least one element from each minimal inconsistent set, finding these sets while computing the consolidation. In the AGM theory, kernel consolidation also is based on discarding formulas from minimal inconsistent sets, but without requiring that the withdrawal be minimal, being more general — thus we focus on the AGM approach.

The AGM paradigm is based on a duality between rationality postulates and operator constructions. The standard consolidation, via discarding formulas, is briefly presented in this chapter as a special case of base contraction, while the next chapter reviews in detail and extends the AGM paradigm aiming at founding other forms of consolidation.

Hansson (1997) proposed two general consolidation methods for belief bases, which are special cases of the corresponding contraction methods: partial meet and kernel consolidation. Each method has a correspondence with a set of rationality postulates that capture the desired properties of a consolidation operation. We begin with these desiderata, then link them to the constructions, adapting Hansson's work to our probabilistic setting.

In Hansson's notation, $\Gamma!$ is the result of consolidating a base Γ. We say $\Gamma! \in \mathbb{K}_c$ is a *consolidation operation* for $\Gamma \in \mathbb{K}_c$ if it satisfies the following postulates:

Postulate 4.2.1 (Inclusion). $\Gamma! \subseteq \Gamma$.

Postulate 4.2.2 (Success). $\Gamma!$ *is consistent.*

To not overload the term "consolidation" and avoid confusion in some parts, we also refer to an AGM-like consolidation operation, which must satisfy the pair of postulates above, as an *abrupt consolidation*. These postulates ignores the minimum change notion; indeed, $\Gamma! = \varnothing$ is a consolidation operation for any $\Gamma \in \mathbb{K}_c$. The following two postulates capture the minimality of changes to different degrees:

Postulate 4.2.3 (Relevance). *If* $\beta \in \Gamma \setminus \Gamma!$*, there is a* Ψ *such that* $\Gamma! \subseteq \Psi \subseteq \Gamma$*,* Ψ *is consistent and* $\Psi \cup \{\beta\}$ *is not.*

Postulate 4.2.4 (Core-retainment). *If* $\beta \in \Gamma \setminus \Gamma!$*, there is a* Ψ *such that* $\Psi \subseteq \Gamma$*,* Ψ *is consistent and* $\Psi \cup \{\beta\}$ *is not.*

It is easy to see that (Relevance) implies (Core-retainment). While the latter demands that any discarded formula β be in a minimal inconsistent subset of Γ, the former requires more: that $\Gamma! \subsetneq \Gamma$ be consistent and that the addition of β to $\Gamma!$, when consistent, preclude it from being consistently expanded in some other way. For instance, consider the base $\Gamma = \{P(x_1) \geq 1, P(\neg x_1) \geq 1, P(\neg x_1 \vee \bot) \geq 0.5\}$ – where $\neg x_1 \vee \bot$ is used to keep the base canonical –, which is clearly inconsistent. While $\Gamma! = \{P(\neg x_1) \geq 1, P(\neg x_1 \vee \bot) \geq 0.5\}$ satisfies (Relevance), $\Gamma! = \{P(\neg x_1) \geq 1\}$ satisfies only (Core-retainment), because $P(\neg x_1 \vee \bot) \geq 0.5$ could be consistently added to it without precluding some other consistent expansion.

To start linking these postulates to constructions, we need some definitions:

Definition 4.2.5. Given a base $\Gamma \in \mathbb{K}_c$, the *remainder set* $\Gamma\perp^{\perp}$ is such that $\Psi \in \Gamma\perp^{\perp}$ if:

- $\Psi \subseteq \Gamma$;
- Ψ is consistent;
- any set Ψ' such that $\Psi \subsetneq \Psi' \subseteq \Gamma$ is inconsistent.

In other words, $\Gamma\perp^{\perp}$ is the set of *maximal consistent subsets* of Γ.

Definition 4.2.6. A function $\gamma : \{\Gamma\perp^{\perp}\} \to 2^{\mathbb{K}_c}$ is a *selection function* for a base $\Gamma \in \mathbb{K}_c$ if:

- $\Gamma\perp^{\perp} \neq \varnothing$ implies $\varnothing \neq \gamma(\Gamma\perp^{\perp}) \subseteq \Gamma\perp^{\perp}$;
- $\Gamma\perp^{\perp} = \varnothing$ implies $\gamma(\Gamma\perp^{\perp}) = \{\Gamma\}$.

Definition 4.2.7. The operation $\Gamma!$ is a *partial meet consolidation* for a base $\Gamma \in \mathbb{K}_c$ if $\Gamma! = \bigcap \gamma(\Gamma\perp^{\perp})$ for some selection function γ.

Now the first of Hanssen's representation results for base consolidation can be adapted to our case:

Proposition 4.2.8. $\Gamma!$ *satisfies (Success), (Inclusion) and (Relevance) iff* $\Gamma!$ *is a partial meet consolidation.*

Together with (Inclusion) and (Success), (Relevance) characterises a consolidation operation $\Gamma!$ as the intersection of some maximal consistent subsets of Γ. If we use (Core-retainment) instead of (Relevance), we need other definitions to specify the corresponding construction:

Definition 4.2.9. Given a base $\Gamma \in \mathbb{K}_c$, the *kernel set* $\Gamma\perp\!\!\!\perp^{\perp}$ is such that $\Psi \in \Gamma\perp\!\!\!\perp^{\perp}$ if:

- $\Psi \subseteq \Gamma$;
- Ψ is inconsistent;
- any set Ψ' such that $\Psi' \subsetneq \Psi \subseteq \Gamma$ is consistent.

Each $\Psi \subseteq \Gamma$ is in $\Gamma\perp\!\!\!\perp^{\perp}$ iff it is a minimal inconsistent set, also called a $\perp$-*kernel*.

Definition 4.2.10. A function $\delta : \{\Gamma\perp\!\!\!\perp^{\perp}\} \to \mathbb{K}_c$ is a *incision function* for a base $\Gamma \in \mathbb{K}_c$ if:

- $\delta(\Gamma\perp\!\!\!\perp^{\perp}) \subseteq \bigcup \Gamma\perp\!\!\!\perp^{\perp}$;
- for any $\Psi \in \Gamma\perp\!\!\!\perp^{\perp}$, $\Psi \cap \delta(\Gamma\perp\!\!\!\perp^{\perp}) \neq \varnothing$.

The set $\delta(\Gamma\perp\!\!\!\perp^{\perp})$ is said to be a *hitting set* of $\Gamma\perp\!\!\!\perp^{\perp}$; *i.e.*, a set containing at least one element from each minimal inconsistent subset of Γ.

Definition 4.2.11. The operation $\Gamma!$ is a *kernel consolidation* for a base $\Gamma \in \mathbb{K}_c$ if $\Gamma! = \Gamma \setminus \delta(\Gamma\perp^{\perp})$ for some incision function δ.

In other words, a kernel consolidation of a base Γ is the result of removing at least one element from each minimal inconsistent subset of Γ.

With (Core-retainment) instead of (Relevance), another representation result from Hansson (1997) can be adapted to probabilistic logic:

Proposition 4.2.12. $\Gamma!$ *satisfies (Success), (Inclusion), (Core-retainment) iff* $\Gamma!$ *is a kernel consolidation.*

As (Relevance) implies (Core-retainment), every partial meet consolidation is also a kernel consolidation. For example, recall the inconsistent base $\Gamma = \{P(x_1) \geq 1, P(\neg x_1) \geq 1, P(\neg x_1 \vee \bot) \geq 0.5\}$. There are two maximal consistent subsets of Γ, $\Psi_1 = \{P(x_1) \geq 1\}$ and $\Psi_2 = \{P(\neg x_1) \geq 1, P(\neg x_1 \vee \bot) \geq 0.5\}$, and two minimal inconsistent subsets, $\Delta_1 = \{P(x_1) \geq 1, P(\neg x_1) \geq 1\}$ and $\Delta_1 = \{P(x_1) \geq 1, P(\neg x_1 \vee \bot) \geq 0.5\}$. The consolidation $\Gamma! = \{P(x_1) \geq 1\}$ is both a partial meet consolidation, $\Gamma! = \bigcap \gamma(\{\Psi_1, \Psi_2\})$ with $\gamma(\{\Psi_1, \Psi_2\}) = \{\Psi_1\}$, and a kernel consolidation, $\Gamma! = \Gamma \setminus \delta(\{\Delta_1, \Delta_2\})$ with $\delta(\{\Delta_1, \Delta_2\}) = \Psi_2$.

The observation above implies that, while performing either a partial meet or a kernel consolidation, only conditionals from the minimal inconsistent subsets are removed. Thus, a conditional that does not belong to any minimal inconsistent subset (that is, a free conditional) can be ignored when one is choosing what to discard. This is a consequence of (Core-retainment) — and, logically, of (Relevance) —, which precludes free conditionals from being ruled out during a consolidation. For these reasons, Hunter and Konieczny (2006, 2008) say that minimal inconsistent sets are the causes of the inconsistency and are its purest form.

A consolidation operator $\mathcal{C} : \mathbb{K}_c \to \mathbb{K}_c$ can be defined in such a way that $\mathcal{C}(\Gamma)$ is a kernel (or a partial meet) consolidation for each $\Gamma \in \mathbb{K}_c$. We say that a consolidation operator $\mathcal{C}(\Gamma)$ satisfying (Inclusion) for all $\Gamma \in \mathbb{K}_c$ — $\mathcal{C}(\Gamma) \subseteq \Gamma$ — is an *abrupt consolidation operator*. For an abrupt consolidation operator via kernel or partial meet consolidation to be well-defined, one needs an incision function δ (or a selection function γ) for each base $\Gamma \in \mathbb{K}_c$. This can trivially achieved for instance by defining $\delta(\Gamma \perp\!\!\!\perp^{\perp}) = \bigcup \Gamma \perp\!\!\!\perp^{\perp}$ ($\gamma(\Gamma \perp^{\perp}) = \Gamma \perp\!\!\!\perp^{\perp}$), in what is called *full meet* consolidation. In practice, one can define functions δ or γ through a computer program that compute them, returning for instance a first-found hitting set $\delta(\Gamma \perp\!\!\!\perp^{\perp})$ or the largest maximal consistent subset of Γ in $\Gamma \perp^{\perp}$.

To implement these operators, one needs to compute maximal consistent subsets or minimal inconsistent subsets. To compute a single minimal inconsistent subset of a base in an arbitrary logic, the standard methods are the constructive (van Maaren and Wieringa, 2008) and the destructive approaches (Marques-Silva, 2010). These methods can implement an oracle in algorithms that find the set of all minimal inconsistent subsets together with its minimal hitting sets, as Reiter (1987) proposed. Furthermore, if Ψ is a minimal hitting set of $\Gamma \perp\!\!\!\perp^{\perp}$, $\Gamma \setminus \Psi$ is a maximal consistent subset of Γ, which reinforces the role of minimal inconsistent sets in these types of consolidation. Klinov (2011) has adapted these techniques to probabilistic logic, enhancing Reiter's algorithm, exploiting linear programming features. It is also possible to compute maximal consistent subsets directly, without computing minimal inconsistent subsets, as Liffiton and Sakallah (2005) show for classical propositional logic. The idea is to solve a sequence of MAXSAT (Maximum Satisfiability) instances, and Hansen *et al.* (1998) put forward a mixed integer linear programming (MILP) approach to the probabilistic version of MAXSAT.

The detailed discussion of methods to compute minimal inconsistent and maximal consistent subsets is out of the scope of this work. The main goal of this section is to point out the central role of minimal inconsistent sets when bases are consolidated via ruling formulas out. Due to such centrality, minimal inconsistent sets have been the basis for devising both inconsistency measures, such as $\mathcal{I}_{MIS}$ and $\mathcal{I}_{MIS^C}$, and their rationality postulates, such as (Independence) and (MIS-

Separability).

4.3 Changing Probability Bounds

When we move from classical to probabilistic logic, there is a natural way to relax a conditional without completely losing its information. Note that ruling out a probabilistic conditional $P(\varphi|\psi) \geq q$ is semantically equivalent to changing it to $P(\varphi|\psi) \geq 0$, so it is a particular (and extreme) case of relaxing the probability bound. If we need to give up the belief on $P(\varphi|\psi) \geq q$ to restore consistency, perhaps there is some positive $q' < q$ such that $P(\varphi|\psi) \geq q'$ can still be consistently believed. As it will be discussed, it is this more general sort of consolidation that is behind the consolidation measures $\mathcal{I}_p$. This motivates the definition of consolidation operators that only change the probability bounds, which are characterised by the following property, adapted from (Thimm, 2011):

Postulate 4.3.1 (**Structural Preservation**). *For any* $\Gamma \in \mathbb{K}_c$, $\mathcal{C}(\Gamma) = \Lambda_\Gamma(q)$, *for some vector* $q \in [0,1]^{|\Gamma|}$.

A *quantitative consolidation operator* is a function $\mathcal{C} : \mathbb{K}_c \to \mathbb{K}_c$ satisfying (Success), (Vacuity) and (Structural Preservation); in this case, we say that $\mathcal{C}(\Gamma)$ is a *quantitative consolidation* of Γ. This postulates triad does not rule out a trivial consolidation operator defined as $\mathcal{C}(\Gamma) = \Lambda_\Gamma(\langle 0, 0, \ldots, 0\rangle)$ for all Γ; such $\mathcal{C}$ replaces all probability lower bounds by 0, being equivalent to the abrupt consolidation operator that removes all formulas from the base. To avoid this, we need a minimal change requirement, demanding that the probability bounds changes be minimal somehow. This is captured by the definition of maximal consolidation:

Definition 4.3.2. Let $\Gamma = \Lambda_\Gamma(q)$ be an arbitrary canonical base in $\mathbb{K}_c$ with size $|\Gamma| = m$. We say a consistent $\Lambda_\Gamma(q')$ is a *maximal consolidation* of Γ if there is no $q'' \in [0,1]^m$ such that $|q''_i - q_i| \leq |q'_i - q_i|$ for all $1 \leq i \leq m$, $|q''_i - q_i| < |q'_i - q_i|$ for some $1 \leq i \leq m$ and $\Lambda_\Gamma(q'')$ is consistent.

Using maximal consolidations, we can rewrite a postulate from Potyka and Thimm (2014):

Postulate 4.3.3 (**Pareto-Optimality**). *If* $\mathcal{C}$ *is a quantitative consolidation operator, then, for any base* $\Gamma \in \mathbb{K}_c$, $\mathcal{C}(\Gamma)$ *is a maximal consolidation of* Γ.

We say a consolidation operator satisfying *Pareto-Optimality* is a *maximal consolidation operator*. The postulate of (Pareto-Optimality) says that $\mathcal{C}$ is such that, if any probability bound in $\mathcal{C}(\Gamma)$ were closer to the corresponding bound in Γ, the resulting base would not be consistent. This desired property implies that no lower bound in a conditional $\alpha = P(\varphi|\psi) \geq q$ in Γ can be raised to form $\beta = P(\varphi|\psi) \geq q'$, with $q' > q$, in $\mathcal{C}(\Gamma)$, for the base $(\mathcal{C}(\Gamma) \setminus \{\alpha\}) \cup \{\beta\}$ would still be consistent with a strictly smaller adjustment in the probability bounds. This consequence seems intuitive, since it ensures the consolidation will only weaken, but not strengthen formulas in a base, and we can state it separately:

Definition 4.3.4. Let $\Gamma = \Lambda_\Gamma(q)$ be an arbitrary canonical base in $\mathbb{K}_c$ with size $|\Gamma| = m$. We say a $\Lambda_\Gamma(q')$ is a *weakening* of Γ if $q'_i \leq q_i$ for all $1 \leq i \leq |\Gamma|$. When a weakening of Γ is consistent, we say it is a *natural consolidation* of Γ.

Postulate 4.3.5 (Non-Strengthening). *If $\mathcal{C}$ is a quantitative consolidation operator, then, for any base $\Gamma \in \mathbb{K}_c$, $\mathcal{C}(\Gamma)$ is a weakening of Γ.*

We can say a maximal consolidation is a natural consolidation such that, if some probability lower bound were less relaxed, the base would still be inconsistent. For instance, consider the example below:

Example 4.3.6. Consider the canonical base $\Gamma = \{P(\varphi) \geq 0.6, P(\varphi) \leq 0.3\}$. An adjustment in the probability bounds that makes the base consistent is $\Gamma' = \{P(\varphi) \geq 0.3, P(\varphi) \leq 0.6\}$ [1]. Nonetheless, another adjustment conserves strictly more information, $\Gamma'' = \{P(\varphi) \geq 0.3, P(\varphi) \leq 0.3\}$, so that Γ' is not a maximal consolidation of Γ. Γ'' is a maximal consolidation, for the probability assessment $P(\varphi) \geq 0.3$ could not be consistently less relaxed. □

Typically, an inconsistent canonical base has several (often infinite) maximal consolidations, corresponding to a Pareto frontier. In the example above, any $\Psi = \{P(\varphi) \geq q_1, P(\varphi) \leq q_2\}$ such that $q_1 \in [0.3, 0.6]$ and $q_1 = q_2$ would be a maximal consolidation of Γ. Nevertheless, some maximal consolidations of Γ can be said to be closer to Γ than others, if one considers probability bounds in a vector space. To construct a consolidation operator that returns some closer maximal consolidation, one can employ methods that minimise distances between these vectors of probability bounds, as done for $\mathcal{I}_p$, plus some criterion to guarantee uniqueness. We discuss this method in Section 4.3.1. Other approaches to modify probability bounds minimise "discrepancies" between the base and a probabilistic interpretation, similarly to $\mathcal{I}_p^\varepsilon$. Then, the probabilistic interpretation is used to recover consistent probability bounds, but (Pareto-Optimality) is not assured, as we show in Section 4.3.2.

4.3.1 Minimising Distances between Probabilities

The most direct approach to consolidate a probabilistic knowledge base via probability bound changing is to minimise a distance between the original, possibly inconsistent probability bounds in the base and the probability bounds in a maximal consolidation. This approach to consistency restoration is clearly inspired by the inconsistency measures $\mathcal{I}_p$; or vice-versa. In fact, Muiño (2011) introduced both the consolidation operator and inconsistency measure together.

Given a canonical base $\Gamma = \Lambda_\Gamma(q)$, each p-norm distance defines a set of consistent probability lower bound vectors that are closest to q — each defining a "closest" consolidation. Let $D_p : \mathbb{K}_c \to 2^{\mathbb{K}_c}$ be a function that returns the the closest consolidations according to a p-norm distance:

$$D_p(\Gamma) = \{\Lambda_\Gamma(q') | \Gamma = \Lambda_\Gamma(q) \ , \ \Lambda_\Gamma(q') \text{ is consistent and } d_p(q, q') \text{ is minimum}\}$$

If Γ has only unconditional probability assessments, any finite $p > 1$ yields a unique closest vector of probability bounds and $D_p(\Gamma)$ is a singleton, but in the general case this does not hold, due to non-convexity:

Example 4.3.7. Consider the base $\Gamma = \{P(x_1|x_2) \geq 1, P(x_2) \geq 1, P(\neg x_1) \geq 1\}$. As the third conditional is equivalent to $P(x_1) \leq 0$ and the first two imply $P(x_1) \geq 1$, Γ is inconsistent. For the

[1] Technically, $P(\varphi) \leq 0.3$ abbreviates $P(\neg\varphi) \geq 1 - 0.3$, and decreasing the lower bound in the latter is increasing the upper bound in the former.

1-norm (Manhattan) distance, we have:

$$D_1(\Gamma) = \big\{\{P(x_1|x_2) \geq 0, P(x_2) \geq 1, P(\neg x_1) \geq 1\}, \{P(x_1|x_2) \geq 1, P(x_2) \geq 0, P(\neg x_1) \geq 1\}, \\ \{P(x_1|x_2) \geq 1, P(x_2) \geq 1, P(\neg x_1) \geq 0\}\big\}.$$

Let $q = \langle 1, 1, 1\rangle$ be the vector of the lower bounds in Γ. There are three vectors q' of consistent lower bounds for Γ minimising $d_1(q, q')$: $r = \langle 0, 1, 1\rangle$, $s = \langle 1, 0, 1\rangle$ and $t = \langle 1, 1, 0\rangle$. Nonetheless, note that the vector $u = 0.5r + 0.5s = \langle 0.5, 0.5, 1\rangle$ yields an inconsistent base:

$$\Lambda_\Gamma(u) = \{P(x_1|x_2) \geq 0.5, P(x_2) \geq 0.5, P(\neg x_1) \geq 1\},$$

due to the fact that $P(x_1|x_2) \geq 0.5$ and $P(x_2) \geq 0.5$ imply $P(x_1) \geq 0.25$, contradicting $P(\neg x_1) \geq 1$. Hence, the set of q''s such that $\Lambda_\Gamma(q')$ is consistent is not convex. In the general case, this non-convexity entails the non-uniqueness of the closest consistent probability bounds vector even for finite $p > 1$. For instance, consider the inconsistent base:

$$\Psi = \{P(x_1 \wedge \neg x_1|x_2) \geq 1, P(x_2) \geq 1\}.$$

In order not to imply a positive probability for $x_1 \wedge \neg x_1$, any consolidation $\Lambda_\Psi(\langle q_1, q_2\rangle)$ must be such that $q_1 q_2 = 0$. Therefore, for any $p \in \mathbb{N}_{>0}$ there are only two vectors q such that $\Lambda_\Psi(q)$ is consistent and $d_p(\langle 1, 1\rangle, q)$ is minimised: $r = \langle 1, 0\rangle$ and $s = \langle 0, 1\rangle$. That is because either $q_1 = 0$ or $q_2 = 0$ and there is no need to change the other lower bound. For any positive integer p, it follows that:

$$D_p = \big\{\{P(x_1 \wedge \neg x_1|x_2) \geq 0, P(x_2) \geq 1\}, \{P(x_1 \wedge \neg x_1|x_2) \geq 1, P(x_2) \geq 0\}\big\}$$

For $p = \infty$, any other base with $q_1 = 0$ or $q_2 = 0$ yields a closest consolidation, and D_p still contains more than one base. □

As the closest consistent vector of probability lower bounds may be not unique, further constraints must be used to choose a single maximal consolidation in order to specify a consolidation operator as a function. For instance, one can consider the closest base with maximum entropy, the closest base that is the most preferred according to some relation, or even the first closest base found by some implementation that computes the distance minimisation. For any $p \in \overline{\mathbb{N}}_{>0}$, we define the consolidation operator $\mathcal{C}_p : \mathbb{K}_c \to \mathbb{K}_c$ as a function that takes an arbitrary base $\Gamma = \Lambda_\Gamma(q)$ and returns a natural consolidation $\Lambda_\Gamma(q')$ that minimises $d_p(q, q')$, employing some criteria to select a single q'. That is, for each $\Gamma \in \mathbb{K}_c$:

$$\mathcal{C}_p(\Gamma) \in D_p(\Gamma)$$

We point out that a myriad of different consolidation operators can implement a $\mathcal{C}_p$, depending on how a single natural consolidation is chosen when there are several closest ones. We use $\mathcal{C}_p$ to refer to an arbitrary consolidation operator satisfying the definition above, and the properties proved for an arbitrary $\mathcal{C}_p$ hold for all such operators.

Example 4.3.8. Consider the following probabilistic bases:

$$\begin{aligned} \Gamma &= \{P(x_1) \geq 0.9, P(x_2) \geq 0.5, P(\neg x_1 \vee \neg x_2) \geq 0.75\}; \\ \Delta &= \{P(x_1|x_2) \geq 0.6, P(\neg x_1|x_2) \geq 0.6, P(x_2) \geq 0.9\}. \end{aligned}$$

As we argued in Example 3.2.2, both bases are inconsistent.

For any finite $p \geq 2$, $D_p(\Gamma)$ and $D_p(\Delta)$ are singletons, and applying the operator $\mathcal{C}_p$ returns:

$$\begin{aligned} \mathcal{C}_p(\Gamma) &= \{P(x_1) \geq 0.85, P(x_2) \geq 0.45, P(\neg x_1 \vee \neg x_2) \geq 0.7\}\,; \\ \mathcal{C}_p(\Delta) &= \{P(x_1|x_2) \geq 0.5, P(\neg x_1|x_2) \geq 0.5, P(x_2) \geq 0.9\}\,. \end{aligned}$$

When $p = \infty$, the minimisation of d_∞ in the definition of both $\mathcal{C}_\infty(\Delta)$ and $\mathcal{C}_\infty(\Gamma)$ exceptionally gives a unique solution, corresponding to those consolidations above. Nevertheless, if $p = 1$, the minimisation d_1 has multiple solutions in both cases. Depending on the selection criteria, applying $\mathcal{C}_1$ may return any base in the form below:

$$\begin{aligned} \mathcal{C}_1(\Gamma) &= \{P(x_1) \geq r_1, P(x_2) \geq r_2, P(\neg x_1 \vee \neg x_2) \geq r_3\},\ r \leq \langle 0.9, 0.5, 0.75\rangle, r_1 + r_2 = 2 - r_3\,; \\ \mathcal{C}_1(\Delta) &= \{P(x_1|x_2) \geq q_1, P(\neg x_1|x_2) \geq q_2, P(x_2) \geq 0.9\},\ q \leq \langle 0.6, 0.6\rangle,\ q_1 + q_2 = 1. \end{aligned}$$

□

A priori, one cannot guarantee that the consolidation operator $\mathcal{C}_p$ always returns a maximal consolidation, but the following result assures this is indeed the case:

Proposition 4.3.9. *For any canonical base $\Gamma \in \mathbb{K}_c$ and $p \in \overline{\mathbb{N}}_{>0}$, $\mathcal{C}_p(\Gamma)$ is a maximal consolidation of Γ.*

Lemma 4.3.10. *For any $p \in \overline{\mathbb{N}}_{>0}$, $\mathcal{C}_p$ is well-defined and satisfies (Success), (Vacuity), (Structural Preservation) and (Pareto-Optimality).*

Muiño (2011) proposed similar consolidation operators for precise probabilistic knowledge bases. Jaumard *et al.* (1991) applied this method to the special case of unconditional probabilities, with $p = 1$, in the problem they called Restore Satisfiability (RSAT). Baioletti and Capotorti explore this same particular case for the Coherence of Probability Assessment problem, which is equivalent to PSAT (Cozman and di Ianni, 2013). In Psychology, Batsell *et al.* (2002) also uses $\mathcal{C}_1$ to eliminate incoherence from subjective precise probability assessments in human reasoning.

To compute $\mathcal{C}_p$ is to solve the optimisation problem (3.6)-(3.8) (or, equivalently, (3.9)-(3.13)) seen in Section 3.2. The consistent, consolidated bounds returned by $\mathcal{C}_p$ are the solution on r to that problem. Consequently, the computation of $\mathcal{C}_p$ has the same drawbacks as that of $\mathcal{I}_p$, due to the fact that the optimisation problem has a non-convex search space.

It might be clear that the objective function being minimised in the program that computes $\mathcal{C}_p(\Gamma)$ is exactly the p-norm distance to the closest consistent probability bounds, that is, $\mathcal{I}_p(\Gamma)$. In this sense, $\mathcal{I}_p(\Gamma)$ can be viewed as gauging the effort to consolidate Γ via $\mathcal{C}_p$, measured through the changes made in each probability bound.

4.3.2 Minimising a Discrepancy to a Probabilistic Interpretation

In this section, we present a consolidation operator related to the minimal violation measure $\mathcal{I}_p^\varepsilon$. Since this inconsistency measure looks for a probabilistic interpretation π that minimises $d_p^\varepsilon(\Gamma, \pi)$, we can employ this probability mass to recover consistent probability bounds for Γ. Recall that $d_p^\varepsilon(\Gamma, \pi)$ measures how much a probabilistic interpretation π violates the restrictions corresponding to the conditionals in Γ, so that its minimum yields an "approximated model" for Γ. With this

model in hands, it is straightforward to find the adjustments needed in the lower bounds for each conditional to be satisfied.

Potyka and Thimm (2014) put forward a consolidation procedure for precise probabilistic bases based on $\mathcal{I}_p^\varepsilon$ with these ideas. As there may be several π minimising $d_p^\varepsilon(\Gamma, \pi)$, the authors select the one that maximises entropy, as defined in Section 3.3, then compute the conditional probability it induces. If π is such that the conditioning formula has null probability, that conditional is trivially satisfied, and no change in the probabilities is needed. To formalise this consolidation operator, let $D_p^\varepsilon : \mathbb{K}_c \to 2^{\Pi_n}$ be a function that takes bases $\Gamma \in \mathbb{K}_c$ and returns the set of probabilistic interpretations π minimising $d_p^\varepsilon(\Gamma, \pi)$. For any $p \in \overline{\mathbb{N}}_{>0}$ and precise base $\Gamma = \{P(\varphi_i|\psi_i) = q_i | 1 \leq i \leq m\}$, Potyka and Thimm (2014) define the consolidation operator $\mathcal{C}_p^{ME}$ via:

$$\begin{aligned} \mathcal{C}_p^{ME}(\Gamma) &= \{P(\varphi_i|\psi_i) = q_i' | 1 \leq i \leq m\}, \text{ where} \\ \pi &= \underset{\pi \in D_p^\varepsilon(\Gamma)}{\arg\max} \Big\{ - \sum_{w \in W_n} \pi(w) \ln(\pi(w)) \Big\} \\ q_i' &= \begin{cases} \frac{P_\pi(\varphi_i \wedge \psi_i)}{P_\pi(\psi_i)} & , \quad \text{if } P_\pi(\psi_i) > 0; \\ q_i & , \quad \text{otherwise.} \end{cases} \end{aligned} \tag{4.1}$$

Their consolidation operator applies only to precise knowledge bases and clearly does no respect (Non-strengthening), since lower bounds can be raised (or upper bounds, decreased). Once the discrepancy d_p^ε is already generalised for imprecise probabilistic knowledge bases, it is simple to generalise $\mathcal{C}_p^{ME}$. Nevertheless, in order to satisfy (Non-strengthening), we propose an extra modification, precluding lower bounds from being raised. Furthermore, as the measure $\mathcal{I}_p^\varepsilon$ is given by any π minimising d_p^ε, not only by the one with maximum entropy, we abandon this criterion. That is, to consolidate the base Γ, we consider that a $\pi \in D_p^\varepsilon(\Gamma)$ is singled out with some arbitrary criteria, for the consolidation operator to be well-defined as a function. For a base $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq m\}$ in $\mathbb{K}_c$ and any $p \in \overline{\mathbb{N}}_{>0}$, the consolidation operator $\mathcal{C}_p^\varepsilon : \mathbb{K}_c \to \mathbb{K}_c$ is defined as:

$$\begin{aligned} \mathcal{C}_p^\varepsilon(\Gamma) &= \{P(\varphi_i|\psi_i) \geq q_i' | 1 \leq i \leq m\}, \text{ where} \\ \pi &\in D_p^\varepsilon(\Gamma) \\ q_i' &= \begin{cases} \min\left\{\frac{P_\pi(\varphi_i \wedge \psi_i)}{P_\pi(\psi_i)}, q_i\right\} & , \quad \text{if } P_\pi(\psi_i) > 0; \\ q_i & , \quad \text{otherwise.} \end{cases} \end{aligned} \tag{4.2}$$

Again, as $\mathcal{C}_p$, there are several consolidation operators that can implement $\mathcal{C}_p^\varepsilon$, depending on how it is chosen the probabilistic interpretation π that minimises d_p^ε. Entropy maximisation is one option, endorsed by Potyka and Thimm (2014), but other, less costly criteria may be employed. For instance, one can pick the first found π, based on a computer program that minimises d_p^ε. Every result proved for $\mathcal{C}_p^\varepsilon$ is meant to hold for every such consolidation operator based on a probabilistic interpretation π minimising the discrepancy d_p^ε, no matter how such π is chosen.

Another way to see $\mathcal{C}_p^\varepsilon$ is through the violations corresponding to the probabilistic interpretation π. When a π minimising $d_p^\varepsilon(\Gamma, \pi)$ does not satisfy a conditional $P(\varphi_i|\psi_i) \geq q_i \in \Gamma$, the corresponding violation is:

$$\varepsilon_i = q_i P_\pi(\psi_i) - P_\pi(\varphi_i \wedge \psi_i)\,.$$

As the conditional is being violated, $P(\psi_i) > 0$, and with little algebraic manipulation we have:

$$\frac{P_\pi(\varphi_i \wedge \psi_i)}{P_\pi(\psi_i)} = q_i - \frac{\varepsilon_i}{P_\pi(\psi_i)}\,.$$

If $\varepsilon_i > 0$, the expression above gives the consolidated lower bound corresponding to π. Recall that, when π satisfies the conditional (including $P_\pi(\psi_i) = 0$), we have $\varepsilon_i = 0$, and the corresponding lower bound is not modified by $\mathcal{C}_p^\varepsilon$. In other words, given the violations vector $\langle \varepsilon_1, \ldots, \varepsilon_m \rangle$ corresponding to a π minimising $d_p^\varepsilon(\Gamma, \pi)$, the probabilities of the conditioning formulas according to π (*i.e.*, $P_\pi(\psi_i)$) define the consolidated lower bounds. In unconditional bases, where all conditioning formulas have probability 1, the violations by themselves fully determine the consolidated probability bounds. Furthermore, for any fixed $1 < p < \infty$, the minimal violations are independent from π (Potyka, 2014).

Example 4.3.11. Consider again the probabilistic bases from Example 3.2.2, 3.2.6 and 4.3.8:

$$\begin{aligned}\Gamma &= \{P(x_1) \geq 0.9, P(x_2) \geq 0.5, P(\neg x_1 \vee \neg x_2) \geq 0.75\}\,;\\ \Delta &= \{P(x_1|x_2) \geq 0.6, P(\neg x_1|x_2) \geq 0.6, P(x_2) \geq 0.9\}\,.\end{aligned}$$

For any finite $p \geq 2$, applying the operator $\mathcal{C}_p^\varepsilon$ to Γ returns the same base, for the violations $\varepsilon_1, \varepsilon_2, \varepsilon_3$ resulting from minimising $d_p^\varepsilon(\Gamma, \pi)$ are the same, as seen in Example 3.2.6, and Γ is unconditional:

$$\begin{aligned}\langle \varepsilon_1, \varepsilon_2, \varepsilon_3 \rangle &= \langle 0.05, 0.05, 0.05 \rangle\,;\\ \mathcal{C}_p^\varepsilon(\Gamma) &= \{P(x_1) \geq 0.85, P(x_2) \geq 0.45, P(\neg x_1 \vee \neg x_2) \geq 0.7\}\,.\end{aligned}$$

Note that $\mathcal{C}_p^\varepsilon(\Gamma) = \mathcal{C}_p(\Gamma)$ for any finite $p \geq 2$. This is not just a coincidence, as Proposition 4.3.12 shows.

When $p = \infty$, the minimisation of $d_\infty^\varepsilon(\Gamma, \pi)$ in the definition of $\mathcal{C}_\infty^\varepsilon(\Gamma)$ gives a unique solution on ε, corresponding to the consolidation above. Nevertheless, if $p = 1$, the minimisation of $d_1^\varepsilon(\Gamma, \pi)$ has multiple solutions. Depending on the selection criteria, applying $\mathcal{C}_1^\varepsilon$ to Γ returns:

$$\mathcal{C}_1^\varepsilon(\Gamma) = \{P(x_1) \geq r_1, P(x_2) \geq r_2, P(\neg x_1 \vee \neg x_2) \geq r_3\},\ r \leq \langle 0.9, 0.5, 0.75 \rangle, r_1 + r_2 = 2 - r_3\,.$$

To consolidate Δ, $\mathcal{C}_p^\varepsilon$ returns different bases for different values of p. Consider the values of $\varepsilon_1, \varepsilon_2, \varepsilon_3$ corresponding to the minimisation of $d_p^\varepsilon(\Delta, \pi)$ for some values of p, as shown in Example 3.2.6:

$$\begin{aligned}p = 1:\quad \langle \varepsilon_1, \varepsilon_2, \varepsilon_3 \rangle &= \langle 0.09000, 0.09000, 0 \rangle\,;\\ p = 2:\quad \langle \varepsilon_1, \varepsilon_2, \varepsilon_3 \rangle &= \langle 0.08824, 0.08824, 0.01765 \rangle\,;\\ p = 3:\quad \langle \varepsilon_1, \varepsilon_2, \varepsilon_3 \rangle &= \langle 0.08615, 0.08615, 0.03853 \rangle\,;\\ p = \infty:\quad \langle \varepsilon_1, \varepsilon_2, \varepsilon_3 \rangle &= \langle 0.08182, 0.08182, 0.08182 \rangle\,.\end{aligned}$$

Note that the assessment $P(x_2) \geq 0.9$ is unconditional, therefore the value of ε_3 determines the value of the corresponding consolidated lower bound: $0.9 - \varepsilon_3$. Hence, if $\varepsilon_3 > 0$, any pair $\langle \pi, \varepsilon \rangle$ minimising $d_p^\varepsilon(\Delta, \pi)$ will be such that $P_\pi(x_2) = 0.9 - \varepsilon_3$. Consequently, the consolidated bounds for

$P(x_1|x_2) \geq 0.6$ and $P(\neg x_1|x_2) \geq 0.6$ depends only on the values of the violations as well:

$$\frac{P_\pi(x_1 \wedge x_2)}{P_\pi(x_2)} = 0.6 - \frac{\varepsilon_1}{0.9 - \varepsilon_3} \text{ and } \frac{P_\pi(\neg x_1 \wedge x_2)}{P_\pi(x_2)} = 0.6 - \frac{\varepsilon_2}{0.9 - \varepsilon_3}\,.$$

Finally, for each value of p, we can compute the consolidated bases:

$$\begin{aligned}
\mathcal{C}_1^\varepsilon(\Delta) &= \{P(x_1|x_2) \geq 0.5, P(\neg x_1|x_2) \geq 0.5, P(x_2) \geq 0.9\}\,;\\
\mathcal{C}_2^\varepsilon(\Delta) &= \{P(x_1|x_2) \geq 0.5, P(\neg x_1|x_2) \geq 0.5, P(x_2) \geq 0.88235\}\,;\\
\mathcal{C}_3^\varepsilon(\Delta) &= \{P(x_1|x_2) \geq 0.5, P(\neg x_1|x_2) \geq 0.5, P(x_2) \geq 0.86147\}\,;\\
\mathcal{C}_\infty^\varepsilon(\Delta) &= \{P(x_1|x_2) \geq 0.5, P(\neg x_1|x_2) \geq 0.5, P(x_2) \geq 0.81818\}\,.
\end{aligned}$$

□

Before exploring the postulates satisfied by $\mathcal{C}_p^\varepsilon$, we remark an equivalence within the unconditional probabilities setting:

Proposition 4.3.12. *For any unconditional probabilistic knowledge base $\Gamma \in \mathbb{K}_c$ and any integer $p > 1$, $\mathcal{C}_p^\varepsilon(\Gamma) = \mathcal{C}_p(\Gamma)$.*

For this result to additionally hold for $p = 1$ and $p = \infty$, the criteria to select a $\pi \in D_p^\varepsilon(\Gamma)$ in $\mathcal{C}_p^\varepsilon$ must yield the base selected out of $D_p(\Gamma)$ in $\mathcal{C}_p$. As a consequence of Propositition 4.3.12, within unconditional bases, $\mathcal{C}_p^\varepsilon$ inherits the properties of $\mathcal{C}_p$. It is easy to see that (Success), (Vacuity) and (Structural Preservation) keep being held by $\mathcal{C}_p^\varepsilon$ in the general case, but (Pareto-Optimality) fails:

Proposition 4.3.13. *For any $p \in \mathbb{N}_{>0} \cup \{\infty\}$, $\mathcal{C}_p^\varepsilon$ satisfies (Success), (Vacuity), (Structural Preservation) and (Non-Strengthening).*

Proposition 4.3.14. *For any $p \in \mathbb{N}_{>0} \cup \{\infty\}$, $\mathcal{C}_p^\varepsilon$ violates (Pareto-Optimality).*

To assess the computational aspects of $\mathcal{C}_p^\varepsilon$, note that, once given a π, computing the vector q' of probability bounds can be done in polynomial time. Hence, the complexity of computing $\mathcal{C}_p^\varepsilon(\Gamma)$ is that of finding the probabilistic interpretation π that minimises the violations. Without imposing extra criteria (e.g. maximum entropy), the search for such a π can be performed through solving the optimisation problem (3.14)-(3.17), which minimises the discrepancy $d_p^\varepsilon(\Gamma, \pi)$ and computes $\mathcal{I}_p^\varepsilon(\Gamma)$. That program has linear constraints, corresponding to a convex search space. Furthermore, Potyka (2014) has already noted that the objective function is linear for $p = 1$ and $p = \infty$ (see Section 3.2.2). Consequently, linear programming techniques also apply to these cases, and computing $\mathcal{C}_1^\varepsilon$ and $\mathcal{C}_\infty^\varepsilon$ can be as hard as solving PSAT in practice.

4.4 Adding Violations

We remark that the main focus of Potyka and Thimm's procedure is not to find consistent probabilities or probability bounds. Their method returns a probabilistic interpretation that can be defined as a generalised model for a possibly inconsistent base. In fact, when probability bounds are recovered from a probability mass π in such a way that (Non-Strengthening) is obeyed — as it is done by $\mathcal{C}_p^\varepsilon$ —, the resulting base may in principle be satisfied by other probabilistic interpretations.

When a π minimising $d_p^\varepsilon(\Gamma,\pi)$ is used to consolidate Γ, the resulting $\mathcal{C}_p^\varepsilon(\Gamma)$ might be satisfied by a π' such that $d_p^\varepsilon(\Gamma,\pi') > d_p^\varepsilon(\Gamma,\pi)$; that is, the p-norm of the violations vector corresponding to π' is greater than the p-norm of such vector yielded by π. Furthermore, it is not clear how $\mathcal{C}_p^\varepsilon$ relates to $\mathcal{I}_p^\varepsilon$, in the sense of measuring inconsistency as the "effort" to consolidate or the distance to consistency. To address these issues and devise a consolidation operator that reflects $\mathcal{I}_p^\varepsilon$ somehow, we can extend the logical language to include the violations ε_i.

An *ε-conditional* is a formula of the form $P(\varphi|\psi) \geq_\varepsilon q$ where $P(\varphi|\psi) \geq q \in \mathcal{L}_{X_n}^P$ is a conditional and $\varepsilon \in [0,1]$ is a real number. The intended meaning is that a probabilistic interpretation π satisfies $P(\varphi|\psi) \geq_\varepsilon q$ if $P_\pi(\varphi \wedge \psi) - qP_\pi(\psi) \geq -\varepsilon$. When $P_\pi(\psi) > 0$, we can rewrite this restriction as $\frac{P_\pi(\varphi\wedge\psi)}{P_\pi(\psi)} \geq q - \frac{\varepsilon}{P_\pi(\psi)}$; that is, an ε-conditional $P(\varphi|\psi) \geq_\varepsilon q$ means that the value of the probability of φ being true given that ψ is true is at least q minus ε over the probability of ψ being true. When $\varepsilon = 0$, the ε-conditional $P(\varphi|\psi) \geq_\varepsilon q$ is denoted simply by $P(\varphi|\psi) \geq q$, as it is equivalent to this regular conditional; and every conditional denotes an ε-conditional with $\varepsilon = 0$. We use $\alpha \bowtie_\varepsilon \beta$, with $\bowtie \in \{\leq, \geq, <, >\}$, to denote that $\alpha = P(\varphi|\varphi) \geq_\varepsilon q$, $\beta = P(\varphi|\varphi) \geq_{\varepsilon'} q$ and $\varepsilon \bowtie \varepsilon'$, for some $\varphi, \psi \in \mathcal{L}_{X_n}$ and $q \in [0,1]$.

We denote by $\mathcal{L}_{X_n}^\varepsilon$ the set of all ε-conditionals. An *ε-base* is a set of ε-conditionals, and $\mathbb{K}^\varepsilon$ denotes the set of all such bases. The consistency of an ε-base is defined as the joint satisfiability of its ε-conditionals by a probabilistic interpretation $\pi : W_n \to [0,1]$. Analogously, the consequence operation Cn_{Pr} can be extended to ε-bases through probabilistic interpretations and the satisfiability definition. We use the term canonical from the regular bases also for ε-bases, with the analogous meaning, and $\mathbb{K}_c^\varepsilon$ is the set of all canonical ε-bases. Within the context of ε-bases, $\mathbb{K}$ and $\mathbb{K}_c$ are understood as the sets of ε-bases that are denoted by the bases in $\mathbb{K}$ and $\mathbb{K}_c$, respectively.

With this extended language in mind, another type of consolidation methods arises: changing the ε's. Even though the consolidation operator $\mathcal{C}_p^\varepsilon$ returns a regular base, the resulting probability bounds are just derived from the probabilistic interpretations minimising $d_p^\varepsilon(\Gamma,\pi)$, and these π's are the intended semantic result of Potyka and Thimm's method. To obtain a consolidated base that is satisfied only by these intended π's, we can use an ε-base to encode exactly those violations that yield the minimum $d_p^\varepsilon(\Gamma,\pi)$. Recall that, for $1 < p < \infty$, any π minimising $d_p^\varepsilon(\Gamma,\pi)$ corresponds to the same violations vector (Potyka, 2014); and any π with this vector of violations minimises $d_p^\varepsilon(\Gamma,\pi)$. Hence, by inserting these violations to form ε-conditionals, one has a consistent base Γ' satisfied only by π's that minimise $d_p^\varepsilon(\Gamma,\pi)$; and if $1 < p < \infty$, Γ' is satisfied by all those π's.

Formally, consider a base $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq m\}$ in $\mathbb{K}_c$. We say a consistent $\Psi \in \mathbb{K}_c^\varepsilon$ is an *ε-consolidation* of Γ if $\Psi = \{P(\varphi_i|\psi_i) \geq_{\varepsilon_i} q_i | 1 \leq i \leq m\}$. Note that, by adopting $\varepsilon_i = 1$, an ε-consolidation can emulate conditionals withdrawal, due to the trivial satisfiability of $P(\varphi_i|\psi_i) \geq_1 q_i$. Hence, ε-consolidations can also be viewed as generalisations of abrupt consolidations.

An *ε-consolidation operator* is a function $\mathcal{V} : \mathbb{K}_c \to \mathbb{K}_c^\varepsilon$ satisfying (Success) and (Vacuity) that takes bases and returns their ε-consolidations. To ensure some minimality of changes, the concept of maximal consolidation can be adapted to ε-bases:

Definition 4.4.1. Let $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq m\}$ be an arbitrary canonical base in $\mathbb{K}_c$. We say a consistent $\Psi = \{P(\varphi_i|\psi_i) \geq_{\varepsilon_i} q_i | 1 \leq i \leq m\}$ is a *maximal ε-consolidation* of Γ if there is no $\varepsilon' \in [0,1]^m$ such that $\varepsilon'_i \leq \varepsilon_i$ for all $1 \leq i \leq m$, $\varepsilon'_i < \varepsilon_i$ for some $1 \leq i \leq m$ and $\{P(\varphi_i|\psi_i) \geq_{\varepsilon'_i} q_i | 1 \leq i \leq m\}$ is consistent.

Analogously, the postulate of (Pareto-Optimality) can be adapted to ε-consolidations:

Postulate 4.4.2 (**Pareto-ε-Optimality**). *If $\mathcal{V}$ is a ε-consolidation operator, then, for any base $\Gamma \in \mathbb{K}_c$, $\mathcal{V}(\Gamma)$ is a maximal ε-consolidation of Γ.*

We say an ε-consolidation operator satisfying (Pareto-Optimality) is a *maximal ε-consolidation operator*. We can also define a version of (Non-Strengthening) for ε-consolidation operators:

Definition 4.4.3. *Let $\Gamma = \{P(\varphi_i|\psi_i) \geq_{\varepsilon_i} q_i | 1 \leq i \leq m\}$ be an arbitrary canonical base in $\mathbb{K}_c^\varepsilon$. We say a $\Gamma = \{P(\varphi_i|\psi_i) \geq_{\varepsilon'_i} q_i | 1 \leq i \leq m\}$ is a ε-*weakening *of Γ if $\varepsilon'_i \geq \varepsilon_i$ for all $1 \leq i \leq |\Gamma|$.*

Postulate 4.4.4 (**Non-ε-Strengthening**). *If $\mathcal{V}$ is a ε-consolidation operator, then, for any base $\Gamma \in \mathbb{K}_c^\varepsilon$, $\mathcal{V}(\Gamma)$ is a ε-weakening of Γ.*

The postulate of (Non-ε-Strengthening) is implied by (Pareto-ε-Optimality), for there is never the need to reduce a violation ε_i in order to consolidate a base.

Through minimising the violations vector p-norm, we define the ε-consolidation operator $\mathcal{V}_p : \mathbb{K}_c \to \mathbb{K}_c^\varepsilon$ for an arbitrary $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq\}$ and for any $1 < p < \infty$:

$$\begin{aligned} \mathcal{V}_p(\Gamma) &= \{P(\varphi_i|\psi_i) \geq_{\varepsilon_i} q_i | 1 \leq i \leq m\} \text{ , where} \\ \langle \varepsilon_1, \varepsilon_2, \ldots, \varepsilon_m \rangle &= \varepsilon = \arg\min_{\varepsilon} \{\|\varepsilon\|_p | \{P(\varphi_i|\psi_i) \geq_{\varepsilon_i} q_i | 1 \leq i \leq m\} \text{ is consitent}\} . \end{aligned}$$

Again, when $p = 1$ or $p = \infty$, there may exist several vectors $\varepsilon = \langle \varepsilon_1, \varepsilon_2, \ldots, \varepsilon_m \rangle$ of violations by a π minimising $d_p^\varepsilon(\Gamma, \pi)$ while yielding an ε-consolidation, thus the definition above of the function $\mathcal{V}_p$ would fail for non-uniqueness. In these cases, an ε-consolidation operator can be defined using any selection criteria to single out a unique vector ε. Potyka and Thimm (2014) favour the maximum entropy principle application, and a more practical decision is to attach this selection to the solution returned by the computer program that solves the optimisation problem (3.14)-(3.17). Analogously to $\mathcal{C}_1^\varepsilon$ and $\mathcal{C}_\infty^\varepsilon$, we assume $\mathcal{V}_1(\Gamma)$ and $\mathcal{V}_\infty(\Gamma)$ return the ε-consolidation corresponding to a vector $\varepsilon = \langle \varepsilon_1, \varepsilon_2, \ldots, \varepsilon_m \rangle$ yielded by a π that minimises $d_1^\varepsilon(\Gamma, \pi)$ or $d_\infty^\varepsilon(\Gamma, \pi)$, respectively, regardless of the selection criteria.

The close relation between $\mathcal{V}_p$ and $\mathcal{I}_p^\varepsilon$ is brought to light if one considers the informal intuitions behind the inconsistency measure. On the one hand, one could conceive $\mathcal{I}_p^\varepsilon(\Gamma)$ as the "effort" required to consolidate Γ by employing $\mathcal{V}_p$, when such effort is measured as the extent to which the (initially null) violations were modified. On the other hand, $\mathcal{I}_p^\varepsilon(\Gamma)$ would be the distance from Γ to its closest consistent ε-consolidation if this distance is quantified through the violations vector p-norm.

Example 4.4.5. Consider again the probabilistic bases from Examples 3.2.6 and 4.3.11:

$$\begin{aligned} \Gamma &= \{P(x_1) \geq 0.9, P(x_2) \geq 0.5, P(\neg x_1 \vee \neg x_2) \geq 0.75\} ; \\ \Delta &= \{P(x_1|x_2) \geq 0.6, P(\neg x_1|x_2) \geq 0.6, P(x_2) \geq 0.9\} . \end{aligned}$$

For any $p \geq 2$, applying the operator $\mathcal{V}_p$ to Γ returns the same base, for the violations $\varepsilon_1, \varepsilon_2, \varepsilon_3$ resulting from minimising $d_p^\varepsilon(\Gamma, \pi)$ are the same, as we seen in Examples 3.2.6 and 4.3.11:

$$\begin{aligned} \langle \varepsilon_1, \varepsilon_2, \varepsilon_3 \rangle &= \langle 0.05, 0.05.0.05 \rangle ; \\ \mathcal{C}_p^\varepsilon(\Gamma) &= \{P(x_1) \geq_{0.05} 0.9, P(x_2) \geq_{0.05} 0.5, P(\neg x_1 \vee \neg x_2) \geq_{0.05} 0.75\} . \end{aligned}$$

When $p = 1$, the minimisation of $d_1^\varepsilon(\Gamma, \pi)$ has multiple solutions. Depending on the selection criteria, applying $\mathcal{V}_1$ to Γ may return any base in the form below:

$$\mathcal{C}_1^\varepsilon(\Gamma) \quad = \quad \{P(x_1) \geq_{\varepsilon_1} 0.9, P(x_2) \geq_{\varepsilon_2} 0.5, P(\neg x_1 \vee \neg x_2) \geq_{\varepsilon_3} 0.75\},\ \varepsilon \geq 0, \varepsilon_1 + \varepsilon_2 + \varepsilon_3 = 0.15\,.$$

To consolidate Δ, $\mathcal{V}_p$ returns different bases for different values of p, since the values $\varepsilon_1, \varepsilon_2, \varepsilon_3$ corresponding to the minimisation of $d_p^\varepsilon(\Delta, \pi)$ varies with p, as shown in Examples 3.2.6 and 4.3.11:

$$\begin{aligned}
p = 1: \quad \langle \varepsilon_1, \varepsilon_2, \varepsilon_3 \rangle &= \langle 0.09000, 0.09000, 0 \rangle\,; \\
p = 2: \quad \langle \varepsilon_1, \varepsilon_2, \varepsilon_3 \rangle &= \langle 0.08824, 0.08824, 0.01765 \rangle\,; \\
p = 3: \quad \langle \varepsilon_1, \varepsilon_2, \varepsilon_3 \rangle &= \langle 0.08615, 0.08615, 0.03853 \rangle\,; \\
p = \infty: \quad \langle \varepsilon_1, \varepsilon_2, \varepsilon_3 \rangle &= \langle 0.08182, 0.08182, 0.08182 \rangle\,.
\end{aligned}$$

From these violations, the corresponding ε-consolidated bases are directly obtained (rounding the values):

$$\begin{aligned}
\mathcal{V}_1(\Delta) &= \{P(x_1|x_2) \geq_{0.090} 0.5, P(\neg x_1|x_2) \geq_{0.090} 0.5, P(x_2) \geq_0 0.9\}\,; \\
\mathcal{V}_2(\Delta) &= \{P(x_1|x_2) \geq_{0.088} 0.5, P(\neg x_1|x_2) \geq_{0.088} 0.5, P(x_2) \geq_{0.018} 0.9\}\,, \\
\mathcal{V}_3(\Delta) &= \{P(x_1|x_2) \geq_{0.086} 0.5, P(\neg x_1|x_2) \geq_{0.086} 0.5, P(x_2) \geq_{0.039} 0.9\}\,, \\
\mathcal{V}_\infty(\Delta) &= \{P(x_1|x_2) \geq_{0.082} 0.5, P(\neg x_1|x_2) \geq_{0.082} 0.5, P(x_2) \geq_{0.082} 0.9\}\,.
\end{aligned}$$

□

The proposition below summarises the properties of $\mathcal{V}_p$ (when $p = 1$ and $p = \infty$, these properties hold for any consolidation operator implementing $\mathcal{V}_p$):

Proposition 4.4.6. *For any $p \in \mathbb{N}_{>0} \cup \{\infty\}$, $\mathcal{C}_p^\varepsilon$ satisfies (Success), (Vacuity) and (Pareto-ε-Optimality).*

The consolidation operator $\mathcal{V}_p$ can be computed through solving program (3.14)-(3.17), taking the solution on ε. Hence, when $p = 1$ or $p = \infty$, such program is linear, and computing $\mathcal{V}_p(\Gamma)$ is as hard as computing $\mathcal{I}_p^\varepsilon(\Gamma)$ — which in practice is no harder than deciding the satisfiability of Γ.

4.5 Related Work

To consolidate inconsistent probabilities, under the coherence setting, Capotorti *et al.* (2010) propose a stepwise correction procedure based on the discrepancy d_{CRV} defined in Section 3.3. In the first step, one looks for a π, with $P_\pi(\bigvee_i \psi_i) = 1$, that minimises $d_{CRV}(\Gamma, \pi)$. If $P_\pi(\psi_i) > 0$ for all $1 \leq i \leq m$, π induces a vector of probabilities $q^\pi \in [0,1]^m$ such that $q_i^\pi = \frac{P_\pi(\varphi_i \wedge \psi_i)}{P_\pi(\psi_i)}$, and the correction procedure halts. The base $\Gamma = \{P(\varphi_i|\psi_i) = q_i^\pi | 1 \leq i \leq m\}$ is the returned consolidation. While $P_\pi(\psi_i) = 0$ for some conditioning event ψ_i, their correction procedure iterates until a solution is found. According to the semantics we use, when $P_\pi(\psi_i) = 0$, we can simply take $q_i^\pi = q_i$, for the conditional is trivially satisfied, defining a consolidation operator for precise knowledge bases $\Gamma = \{P(\varphi_i|\psi_i) = q_i | 1 \leq i \leq m\}$:

$$\begin{aligned}
\mathcal{C}^1_{CRV}(\Gamma) &= \{P(\varphi_i|\psi_i) = q_i^\pi | 1 \leq i \leq m\}, \text{ where} \\
\pi &= \arg\min_{\pi \in \Pi_n} \{d_{CRV}(\Gamma, \pi)\} \\
q_i^\pi &= \begin{cases} \frac{P_\pi(\varphi_i \wedge \psi_i)}{P_\pi(\psi_i)} & , \quad \text{if } P_\pi(\psi_i) > 0; \\ q_i & , \quad \text{otherwise.} \end{cases}
\end{aligned}$$

Note that such $\mathcal{C}^1_{CRV}$ violates (Pareto-Optimality), for pairs of lower and upper bounds are modified together. To fix this, we recall that the precise base Γ can be written as $\{P(\varphi_i|\psi_i) \geq q_i, P(\varphi_i|\psi_i) \leq q_i | 1 \leq i \leq m\}$ and define:

$$\begin{aligned}
\mathcal{C}^2_{CRV}(\Gamma) &= \{P(\varphi_i|\psi_i) \geq \underline{q}_i, P(\varphi_i|\psi_i) \leq \bar{q}_i | 1 \leq i \leq m\}, \text{ where} \\
\pi &= \arg\min_{\pi \in \Pi_n} \{d_{CRV}(\Gamma, \pi)\} \\
q_i^\pi &= \begin{cases} \frac{P_\pi(\varphi_i \wedge \psi_i)}{P_\pi(\psi_i)} & , \quad \text{if } P - \pi(\psi_i) > 0; \\ q_i & , \quad \text{otherwise.} \end{cases} \\
\underline{q}_i &= \min\{q_i, q_i^\pi\}, \ \bar{q}_i = \max\{q_i, q_i^\pi\}.
\end{aligned}$$

If $P(\varphi_i|\psi_i) = q_i^\pi \in \mathcal{C}^1_{CRV}(\Gamma)$ for some $q_i^\pi > q_i$, then $P(\varphi_i|\psi_i) \geq q_i, P(\varphi_i|\psi_i) \leq q_i^\pi \in \mathcal{C}^2_{CRV}(\Gamma)$, correcting a precise probability to an interval, without moving the lower bound up. Similarly, if $q_i^\pi < q_i$, $\mathcal{C}^2_{CRV}$ does not decrease the upper bound.

The scope of both consolidation operators defined above is restricted to precise probabilistic knowledge bases. To deal with general bases, the discrepancy must be tailored to probability bounds. Capotorti *et al.* (2009) introduced a method to correct imprecise probability assessments using d_{CRV} under the coherence framework. In such theory, a set of probability bounds is coherent only if each probability bound is attainable while the others are respected. To compute their correction procedure for an input $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq m\}$, we start by solving, for each $1 \leq f \leq m$, a program. Recall that $\Lambda^P_\Gamma(q') = \{P(\varphi_i|\psi_i) = q'_i | 1 \leq i \leq m\}$. Consider the real variables vectors $\pi = \langle \pi_1, \ldots, \pi_{2^n} \rangle$ and $r = \langle r_1, \ldots, r_m \rangle$, corresponding to a probability mass over W_n and to precise probabilities, respectively:

$$\begin{aligned}
\min d_{CRV}(\Lambda^P_\Gamma(r), \pi) \quad & \text{subject to:} \\
r_f &= q_f \\
r_i &\geq q_i \quad 1 \leq i \leq m \\
\pi &\geq 0 \\
\sum_{j=1}^{2^n} \pi_j &= 1 \\
\sum_{j=1}^{2^n} \pi_j I_{w_j}(\bigvee_i \psi_i) &= 1
\end{aligned}$$

For a fixed f, the program above computes the precise probabilistic base $\Gamma_f = \{P(\varphi_i|\psi_i) = r_i | 1 \leq i \leq m\}$, with $r \geq q$ and $r_f = q_f$, such that the discrepancy d_{CRV} from this base to some

probability mass π is minimised. If $P_\pi(\psi_i) > 0$ for all i, then one can construct a q^π corresponding to a coherent set of conditionals; else, the process can be iterated, as before. Repeating this procedure for all $1 \leq f \leq m$, one has a set of m probability vectors that correspond to coherent probabilistic bases, and Capotorti *et al.* (2009) define the consolidation of Γ by taking the lower envelope of these vectors. The resulting set of assessments is coherent, in the sense that each lower bound is attainable.

Under the semantics we adopt, this attainability restriction is not necessary, as we identify consistency with satisfiability. Furthermore, null probabilities for all conditioning events is not an issue in our setting. Hence, looking for a probabilistic interpretation π in order to consolidate a base Γ, we can simply solve the program (3.26)–(3.29), which computes $\mathcal{I}_{CRV}$, defined in Section 3.3. That is, if $\Gamma = \Lambda_\Gamma(q)$, we minimise the discrepancy between a precise probabilistic base $\Gamma = \Lambda^P_\Gamma(r)$, with $r \geq q$, and a probabilistic interpretation π. Given a solution π, we can construct a vector q' of consistent lower bounds as done for $\mathcal{C}^\varepsilon_p$.

Finally, a consolidation operator $\mathcal{C}_{CRV} : \mathbb{K}_c \to \mathbb{K}_c$ for general (imprecise) probabilistic bases can be defined adapting the method from Capotorti *et al.* (2009), where $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq m\}$ is a base in $\mathbb{K}_c$ and $r, q \in [0,1]^m$ are vectors:

$$\begin{aligned}
\mathcal{C}_{CRV}(\Gamma) &= \{P(\varphi_i|\psi_i) \geq q'_i | 1 \leq i \leq m\}, \text{ where} \\
\pi &= \arg\min_{\pi \in \Pi_n} \{d_{CRV}(\Lambda^P_\Gamma(r), \pi) | r \geq q\} \ , \\
q'_i &= \begin{cases} \min\{\frac{P_\pi(\varphi_i \wedge \psi_i)}{P_\pi(\psi_i)}, q_i\} & , \quad \text{if } P_\pi(\psi_i) > 0; \\ q_i & , \quad \text{otherwise.} \end{cases}
\end{aligned} \tag{4.3}$$

As does the inconsistency measure $\mathcal{I}_{CRV}$, the consolidation operator $\mathcal{C}_{CRV}$ minimises the discrepancy between assigning some precise probabilities $r = \langle r_1, \ldots, r_m \rangle$ that are at least the lower bounds $q = \langle q_1, \ldots, q_m \rangle$ and a probabilistic interpretation π; in some sense, r corresponds to the precise assignment that is closer to coherence while respecting the probability bounds. When $P_\pi(\psi_i) > 0$, such π defines precise conditional probabilities $\frac{P_\pi(\varphi_i \wedge \psi_i)}{P_\pi(\psi_i)}$, which the probability lower bounds are eventually decreased to allow for. Note that, according to our semantic choice, no changes in the probability bounds are needed when the conditioning formula has probability 0 according to π.

Example 4.5.1. Consider again the probabilistic bases from Example 4.3.8:

$$\begin{aligned}
\Gamma &= \{P(x_1) \geq 0.9, P(x_2) \geq 0.5, P(\neg x_1 \vee \neg x_2) \geq 0.75\}\,; \\
\Delta &= \{P(x_1|x_2) \geq 0.6, P(\neg x_1|x_2) \geq 0.6, P(x_2) \geq 0.9\}\,.
\end{aligned}$$

Solving the numeric optimisation, applying the operator $\mathcal{C}_{CRV}$ to these bases returns (probability bounds are truncated):

$$\begin{aligned}
\mathcal{C}_{CRV}(\Gamma) &= \{P(x_1) \geq 0.8726, P(x_2) \geq 0.4321, P(\neg x_1 \vee \neg x_2) \geq 0.6953\}\,; \\
\mathcal{C}_{CRV}(\Delta) &= \{P(x_1|x_2) \geq 0.5, P(\neg x_1|x_2) \geq 0.5, P(x_2) \geq 0.89527\}\,.
\end{aligned}$$

It is important to note that the probabilistic interpretation π that yields $\mathcal{C}_{CRV}(\Gamma)$ is the same used for computing the minimum discrepancy $d_{CRV}(\Lambda^P_\Gamma(r), \pi)$, which is by definition equal to $\mathcal{I}_{CRV}(\Gamma)$,

in Example 3.3.1 — and this also the case with Δ. □

Even though a probability lower bound is relaxed only when the assignment induced by π violates that bound, consolidation operators based on entropy may not satisfy (Pareto-Optimality):

Proposition 4.5.2. *The consolidation operator $\mathcal{C}_{CRV} : \mathbb{K}_c \to \mathbb{K}_c$ satisfies (Success), (Vacuity), (Structural Preservation) and (Non-Strengthening) but not (Pareto-Optimality).*

The proof of the proposition above points out that this may happen in the precise case as well, for the original correction procedure proposed by Capotorti and Regoli (2008).

To compute $\mathcal{C}_{CRV}$, one has to solve the program (3.26)-(3.29), which computes $\mathcal{I}_{CRV}(\Gamma)$, with a logarithmic objective function, despite the convex search space (Capotorti *et al.*, 2010).

4.6 Back to the Postulates Incompatibility

We have seen in this chapter some consolidation operators for probabilistic knowledge bases. The fact that (Consistency), (Independence) and (Continuity) are not jointly satisfiable, as found in the previous chapter, calls for the modification of at least one of these postulates. It was hinted in Section 3.4 that (Independence) could be the one to be abandoned/relaxed, but we are looking for more concrete reasons to do so, hopefully with paths to fix the incompatibility. The postulate of (Independence) is based on the idea that minimal inconsistent sets are the primitive, atomic form of inconsistency — and that free conditionals are "harmless". We showed in Section 4.2 that this view is due to the classical strategy to consolidate bases: formulas withdrawal (abrupt consolidation). In Sections 4.3 and 4.4, consolidation operators tailored to probabilistic logic were investigated. In particular, we showed how $\mathcal{C}_p$ and $\mathcal{V}_p$ generalise abrupt consolidation and can be understood as the consolidation procedures underlying the inconsistency measures $\mathcal{I}_p$ and $\mathcal{I}_p^\varepsilon$.

Even though it shall be now clear the link between minimal inconsistent sets, free conditionals, and the classical approach to consolidating knowledge bases, we still do not have positive reasons to abandon (Independence). The fact that the consolidation methods underlying the inconsistency measures proper to probabilistic bases are different from (although generalisations of) the standard approach of ruling formulas out does not by itself undermines the (Independence) postulate. *A priori*, it may well be the case that minimal inconsistent sets are a characterisation of the inconsistency causes that suits the consolidation operators $\mathcal{C}_p$ and $\mathcal{V}_p$. Only if minimal inconsistent sets somehow fail to capture the inconsistencies in a base — and free conditionals, its harmless elements —, when considering these essentially probabilistic consolidation methods, there will be sufficient evidence towards weakening or dropping (Independence).

In Chapter 6, we investigate the exact link between free conditionals, minimal inconsistent sets and abrupt consolidations and transpose it to quantitative and ε-consolidations, finding new forms of inconsistency characterisation. Before doing so, we take a detour through Chapter 5 to show how maximal consolidation operators, like $\mathcal{C}_p$, may be embedded into the AGM theory of belief revision.

Chapter 5

A Brief Review on Belief Revision: Quantitative Consolidation within the AGM Paradigm

Before returning to the main topic of this work, inconsistency measuring, and its dependencies, we put forward a method to encompass consolidation via probability bounds adjustment within the standard AGM framework of belief revision.

The task of restoring consistency among propositions within classical logic has been formally tackled within the AGM theory, introduced by Alchourrón, Gärdenfors, and Makinson (1985), as a contraction by the contradiction (Hansson, 1997). That is, to consolidate a set of beliefs is to give up — *i.e.*, contract by — the belief in the contradiction. This means that one has to withdraw formulas from the base in a way that the contradiction is no longer implied. Nevertheless, whereas consolidation in classical logic is formalised within the AGM paradigm as discarding formulas (for general AGM theory, see (Hansson, 1999)), in probabilistic logic it is usually the case that probabilities (or probability bounds) are changed instead of being ruled out, as we have seen in Chapter 4. This difference in approach, discarding formulas versus changing them, apparently makes the AGM paradigm unsuitable for the probabilistic consolidation through probability bounds adjustment to be grounded in. The main difficulty is that AGM contractions of a set of beliefs (a *base*) are constrained by rationality postulates that require, among other things, that the contracted base be a subset of the original one; but when a probability (bound) is changed, this requirement is violated.

Indeed, the AGM paradigm has room for weakening formulas while giving up beliefs, in what Hansson (1999) called *pseudo-contractions*. The idea is that a base being contracted by a proposition — like the contradiction — may receive new beliefs implied by the old ones. Nonetheless, until recent there was little (if any) clue about how to construct this pseudo-contractions; *i.e.*, which consequences of the original formulas to include in the contracted (or consolidated) base. When contracting by the contradiction, this issue becomes even more evident, since the classical consequences of an inconsistent set of formulas are the whole logical language. It follows that any consistent base can be a pseudo-contraction by the contradiction of any inconsistent base. Santos, Ribeiro, and Wassermann (2015) have recently put forward a formal method to construct pseudo-contractions that can avoid this problem by using an arbitrary subclassical consequence operation. With such weaker consequence, it is not the case that any classical consequence of the

original base may appear in the contraction, but just some of them.

In this chapter, we propose a way of founding the probabilistic consolidation via probability bounds adjustment on the AGM theory, departing from Santos, Ribeiro, and Wassermann's approach. The concept of *liftable contraction* is introduced as a generalisation of both contraction and pseudo-contraction for a general logic; then, assuming (Pareto-Optimality), the probabilistic consolidation via probability bounds adjustment can be characterised as a liftable contraction by the contradiction. As consolidation is seen only as a particular case of contraction within the AGM paradigm, such characterisation also yields a method for probabilistic contraction. Even though the probabilistic consolidation methods via distance minimisation, presented in Section 4.3, are well-known in the literature (Muiño, 2011; Thimm, 2011), they are shown to be particular cases of general probabilistic contraction methods, which are themselves original to the best of our knowledge.

This chapter is organised as follows: in Section 5.1, we review the AGM theory of base contraction, focusing on partial meet contractions; in Section 5.2, we show how consolidation via probability adjustment can be related to the work of Santos *et al.* (2015); such connection is formalised and generalised in Section 5.3, where liftable contractions are introduced; under such formalism, we introduce a probabilistic contraction in Section 5.4; finally, we sketch how the same concepts can be used to underpin the consolidation via violations, using the extended language seen in Section 4.4.

5.1 Consolidation and Contraction in AGM Theory

The AGM paradigm is a theory that models changes in the belief state of a rational agent (see, for instance, (Hansson, 1999)). A belief base is a way to model a belief state via a set of formulas from some logical language, in which the agent believes in the logical consequences of this set. Probabilistic knowledge bases can be understood as belief bases if we consider a hypothetical underlying agent, which might, for instance, be implemented in an autonomous system. The three main belief change operations in the AGM paradigm are expansion, contraction and revision. The second one handles the giving up of beliefs, through which consolidation is defined as the contraction by the contradiction; *i.e.*, to consolidate a belief state is to give up the belief on the contradiction. To see how belief contraction relates to the consolidation operators we presented, we quickly review the AGM approach to giving up beliefs on bases.

The whole AGM theory is grounded in a set of rationality postulates for belief change operations, which have a correspondence with procedures that implement these operations. Even though the original rationality postulates were introduced by Alchourrón *et al.* (1985) for (logically closed) belief sets, Hansson (1992) generalised them to belief bases, which are sets of logical formulas. In this section, we focus on partial meet contraction for belief bases, which will be later used to found consolidation via probability bounds adjustment. To formalise and characterise this contraction operation, we generalise the postulates and concepts seen in Section 4.2 for partial meet consolidation.

Consider a general language $\mathcal{L}$ containing $\perp$ and a consequence operation $Cn(.)$ over it. A *belief base* is any set $B \subseteq \mathcal{L}$. A contraction operator for a base B is formalised as a function $- : \mathcal{L} \to 2^{\mathcal{L}}$ that takes a formula $\alpha \in \mathcal{L}$ as argument and returns a base $-(\alpha) = B - \alpha \in 2^{\mathcal{L}}$ — the contraction of B by α. For a function to be called a *contraction* operator for B, it must satisfy the following two rationality postulates (generalisations of those seen in Section 4.2), for any $\alpha \in \mathcal{L}$:

Postulate 5.1.1 (Inclusion). *$B - \alpha \subseteq B$.*

Postulate 5.1.2 (Success). *If $\alpha \notin Cn(\varnothing)$, $\alpha \notin Cn(B - \alpha)$.*

In words, the contraction of B by α should introduce no new beliefs in the base, and the resulting base should not imply α, as long as α is not a tautology. A weaker form of inclusion is (Hansson, 1999):

Postulate 5.1.3 (Logical Inclusion). *$Cn(B - \alpha) \subseteq Cn(B)$.*

An operator satisfying (Success) and (Logical Inclusion) is called a *pseudo-contraction* (Hansson , 1999).

The next postulate deals with the notion of minimal changes or the idea of maximal information preservation. While contracting a base by a belief α, it is sensible not to discard beliefs that are not involved in the derivation of α, which would yield a version of (Core Retainment) for contraction. Usually, a stronger postulate is required, which is the contraction-based form of the (Relevance) postulate seen in Section 4.2:

Postulate 5.1.4 (Relevance). *If $\beta \in B \setminus B - \alpha$, there is a B' such that $B - \alpha \subseteq B' \subseteq B$, $\alpha \notin Cn(B')$ and $\alpha \in Cn(B' \cup \{\beta\})$.*

A final postulate states that contracting by formulas that are equivalent (modulo the base B) should yield the same result:

Postulate 5.1.5 (Uniformity). *For all $B' \subseteq B$, if it is the case that $\alpha \in Cn(B')$ iff $\beta \in Cn(B')$, then $B - \alpha = B - \beta$.*

Hansson (1999) proved, for classical propositional logic, that any contraction on a belief base satisfying (Success), (Inclusion), (Relevance) and (Uniformity) can be expressed in a specific form. To present it, we need to generalise some definitions presented in in Section 4.2:

Definition 5.1.6. Given a base $B \subseteq \mathcal{L}$ and a formula $\alpha \in \mathcal{L}$, the *remainder set* $B \perp \alpha$ is such that $X \in B \perp \alpha$ iff:

- $X \subseteq B$;
- $\alpha \notin Cn(X)$;
- for any set Y, if $X \subsetneq Y \subseteq B$, $\alpha \in Cn(Y)$.

Definition 5.1.7. A function $\gamma : \{B \perp \alpha | \alpha \in \mathcal{L}\} \to 2^{\mathbb{K}_c}$ is a *selection function* for a base $B \subseteq \mathcal{L}$ iff:

- if $B \perp \alpha \neq \varnothing$, $\varnothing \neq \gamma(B \perp \alpha) \subseteq B \perp \alpha$;
- otherwise $\gamma(B \perp \alpha) = \{B\}$.

Definition 5.1.8. The operator $- : \mathcal{L} \to 2^{\mathcal{L}}$ is a *partial meet contraction* for a base $B \subseteq \mathcal{L}$ if, for any formula $\alpha \in \mathcal{L}$, $B - \alpha = \bigcap \gamma(B \perp \alpha)$ for some selection function γ. If $\gamma(B \perp \alpha)$ is a singleton for any $\alpha \in \mathcal{L}$, $-$ is called a *maxichoice contraction.*

For classical propositional logic, Hansson linked the rationality postulates to the partial meet construction (Hansson, 1999):

Theorem 5.1.9. *Consider the propositional language $\mathcal{L}_{X_n}$ and the classical consequence operation Cn_{Cl}. The operator $- : \mathcal{L}_{X_n} \to 2^{\mathcal{L}}_{X_n}$ for a base $B \subseteq 2^{\mathcal{L}}_{X_n}$ satisfies (Success), (Inclusion), (Relevance) and (Uniformity) iff $-$ is a partial meet contraction.*

To obtain results for a general language $\mathcal{L}$ with an arbitrary consequence operation $Cn(.)$, we have to impose some restriction on the latter. Below, we define some useful properties a consequence operation may hold:

Definition 5.1.10. Consider a general language $\mathcal{L}$, arbitrary bases $B, B', B'' \subseteq \mathcal{L}$, an arbitrary formula $\alpha \in \mathcal{L}$ and a consequence operation $Cn'(.)$. A consequence operation $Cn(.)$ satisfies

- *Monotonicity* if $B \subseteq B'$ implies $Cn(B) \subseteq Cn(B')$;
- *Idempotence* if $Cn(Cn(B)) \subseteq Cn(B)$;
- *Inclusion* if $B \subseteq Cn(B)$;
- *Compactness* if $\alpha \in Cn(B)$ implies that there is some finite $B' \subseteq B$ such that $\alpha \in Cn'(B)$;
- the *Upper Bound Property* if, for every $B' \subseteq B$ with $\alpha \notin Cn(B')$, there is a $B'' \supseteq B'$ such that $B'' \in B \perp \alpha$;
- *Cn'-Dominance* if $Cn(B) \subseteq Cn'(B)$;
- *Subclassicality*, when $\mathcal{L} = \mathcal{L}_{X_n}$, if and Cn satisfies Cn_{Cl}-Dominance.

If Cn satisfies monotonicity, inclusion and idempotence, we say it is *Tarskian*.

For instance, Cn_{Cl} for the propositional language and Cn_{Pr} for probabilistic conditionals are Tarskian consequence operations, due to their definition through models, but only the former is compact, as we shall see.

Now, Hansson's representation theorem can be generalised (Wassermann, 2000):

Theorem 5.1.11. *Let Cn be a consequence operation satisfying monotonicity and the upper bound property. The operator $- : \mathcal{L} \to 2^{\mathcal{L}}$ for a base $B \subseteq 2^{\mathcal{L}}$ satisfies (Success), (Inclusion), (Relevance) and (Uniformity) iff $-$ is a partial meet contraction.*

To axiomatise maxichoice contraction, (Relevance) can be strengthened (Hansson, 1999):

Postulate 5.1.12 (Fullness). *If $\beta \in B$ and $\beta \notin B - \alpha$, then $\alpha \notin Cn(B - \alpha)$ and $\alpha \in Cn((B - \alpha) \cup \{\beta\})$.*

Hansson (1999) proved the result below for classical logic, and Wassermann (2000) generalised it:

Theorem 5.1.13. *Let Cn be a consequence operation satisfying monotonicity and the upper bound property. The operator $- : \mathcal{L} \to 2^{\mathcal{L}}$ for a base $B \subseteq 2^{\mathcal{L}}$ satisfies (Success), (Inclusion), (Fullness) and (Uniformity) iff $-$ is a maxichoice contraction.*

Hansson (1997) defines a *partial meet consolidation* as a partial meet contraction by the contradiction ($\perp$), which is here adapted:

Definition 5.1.14. The operation $B!$ is a *partial meet consolidation* for a base $B \subseteq \mathcal{L}$ if $B! = \bigcap \gamma(B\bot^{\bot})$ for some selection function γ. If $\gamma(B\bot^{\bot})$ is a singleton, ! is called a *maxichoice consolidation.*

Except for (Uniformity), the postulates for contraction apply to consolidation ($B!$), as contraction by the contradiction, just taking $\alpha = \bot$; and (Success) for consolidation is also called (Consistency). Hansson (1999) proved the following representation result, generalised by Wassermann (2000):

Theorem 5.1.15. *Let Cn be a consequence operation satisfying monotonicity, the upper bound property and such that $\bot \notin Cn(\varnothing)$. $B!$ satisfies (Success), (Inclusion), (Relevance) iff $B!$ is a partial meet consolidation.*

Corollary 5.1.16. *Let Cn be a consequence operation satisfying monotonicity, the upper bound property and such that $\bot \notin Cn(\varnothing)$. $B!$ satisfies (Success), (Inclusion), (Fullness) iff $B!$ is a maxichoice consolidation.*

Taking canonical probabilistic knowledge bases as belief bases over the language $\mathcal{L}^P_{X_n}$ and considering the probabilistic consequence operation $Cn_{Pr}(.)$, we can evaluate some consolidation operators introduced in Chapter 4 with respect to the AGM rationality postulates. We focus on those operators satisfying (Pareto-Optimality), as $\mathcal{C}_p$, for reasons that shall be clear soon.

Consider a quantitative consolidation operator $\mathcal{C}$. For a given canonical base $\Gamma \in \mathbb{K}_c$, $\mathcal{C}(\Gamma)$ is not a contraction by the contradiction, for formulas are not being discarded, but changed, and (Inclusion) — for base contraction — is violated. The postulate of (Logical Inclusion) is satisfied, so $\mathcal{C}(\Gamma)$ could be viewed as a pseudo-contraction by the contradiction. Nevertheless, this postulate is vacuous when Γ is inconsistent — which is the only meaningful application of consolidation —, for $Cn_{Pr}(\Gamma) = \mathcal{L}^P_{X_n}$. In the following, we look for ways to embed quantitative consolidation operators, like $\mathcal{C}_p$, into the AGM paradigm in a more suitable way.

5.2 Changing Probability Bounds versus Discarding Formulas

When probability lower bounds in a canonical probabilistic knowledge base are decreased to consolidate it, this operation satisfies (Logical Inclusion) in a particular way. Not any consequence of a base Γ can appear in its consolidation, but only consequences of each single probabilistic conditional, just with the probability bound modified. Santos *et al.* (2015) investigated contractions on propositional bases considering, besides the classical consequence Cn_{Cl}, an arbitrary Tarskian, subclassical consequence operation $Cn^{\star}$. Based on their ideas, we explore three ways of connecting probabilistic consolidation via changing lower bounds with consolidation via partial meet. First, we define an element-wise consequence operation Cn_{ew} for the probabilistic language $\mathcal{L}^P_{X_n}$.

Definition 5.2.1. The function $Cn_{ew} : \mathbb{K} \to \mathbb{K}$ is a consequence operation such that:

- For all $\Gamma \in \mathbb{K}$, $Cn_{ew}(\Gamma) = \bigcup\{Cn_{ew}(\{\alpha\})|\alpha \in \Gamma\}$;
- For all $\alpha = P(\varphi|\psi) \geq q \in \mathcal{L}^P_{X_n}$, $Cn_{ew}(\{\alpha\}) = \bigcup\{P(\varphi|\psi) \geq q'|q' \in [0,1], q' \leq q\}$.

Proposition 5.2.2. *Cn_{ew} is Tarskian and satisfies Cn_{Pr}-dominance.*

Given a canonical probabilistic base $\Gamma \in \mathbb{K}_c$ and a maximal consolidation operator $\mathcal{C}$, we have $\mathcal{C}(\Gamma) \subseteq Cn_{ew}(\Gamma)$ — equivalently, by monotonicity and idempotency ($\rightarrow$) and inclusion ($\leftarrow$), $Cn_{ew}(\mathcal{C}(\Gamma)) \subseteq Cn_{ew}(\Gamma)$ —, a special kind of inclusion. Furthermore, when Γ is inconsistent, it is not necessarily the case that $Cn_{ew}(\Gamma) = \mathcal{L}^P_{X_n}$, thus this inclusion is not trivially satisfied, as it happens with (Logical Inclusion) using Cn_{Pr}. This motivates the following postulate, proposed by Santos *et al.* (2015), for a general logical language $\mathcal{L}$ and an arbitrary consequence operation $Cn^\star$:

Postulate 5.2.3 (**Inclusion***). $B - \alpha \subseteq Cn^\star(B)$.

If $Cn^\star = Cn_{ew}$, (Inclusion*) is the kind of inclusion any maximal consolidation operator respects, when viewed as a contraction by the contradiction. In other words, decreasing probability lower bounds in a canonical base Γ is the same as picking a canonical subset of $Cn_{ew}(\Gamma)$.

In a similar way, Santos *et al.* (2015) propose starred versions for (Relevance) and (Uniformity).

Postulate 5.2.4 (**Relevance***). *If $\beta \in Cn^\star(B) \setminus B - \alpha$, there is a B' such that $B - \alpha \subseteq B' \subseteq Cn^\star(B)$, $\alpha \notin Cn(B')$ and $\alpha \in Cn(B' \cup \{\beta\})$.*

Postulate 5.2.5 (**Uniformity***). *For all $B' \subseteq Cn^\star(B)$, if it is the case that $\alpha \in Cn(B')$ iff $\beta \in Cn(B')$, then $B - \alpha = B - \beta$.*

The postulate of (Relevance*), together with (Inclusion*), forces a contraction — and a consolidation — to be closed under $Cn^\star$, what is called (Enforced Closure*) (Santos *et al.*, 2015):

Postulate 5.2.6 (**Enforced Closure***). $B - \alpha = Cn^\star(B - \alpha)$.

As maximal consolidation operators return canonical bases, which are not closed under Cn_{ew}, (Relevance*) is violated. Nonetheless, if we consider the closure of the resulting base under Cn_{ew}, this postulate is recovered:

Proposition 5.2.7. *Let $\mathcal{C} : \mathbb{K}_c \rightarrow \mathbb{K}_c$ be consolidation operator satisfying (Pareto-Optimatility) and consider $Cn^\star = Cn_{ew}$. For any $\Gamma \in \mathbb{K}_c$, the consolidation operation ! for Γ defined as $\Gamma! = Cn_{ew}(\mathcal{C}(\Gamma))$ satisfies (Relevance*).*

In the classical propositional case, Santos *et al.* (2015) proved that, for a compact, Tarskian, subclassical $Cn^\star$, a contraction operator $-$ for a base B satisfies (Success), (Inclusion*), (Relevance*) and (Uniformity*) iff $B - \alpha = \bigcap \gamma(Cn^\star(B) \perp \alpha)$ for some selection function γ. Consequently, if $B! = B - \perp$ for a $-$ satisfying these four postulates, ! is a partial meet contraction by the contradiction of the starred closure $Cn^\star(B)$. To prove the analogous result for probabilistic logic, where $\perp$ can denote $P(\perp) \geq 1$, we need an intermediate result, for the probabilistic consequence operation Cn_{Pr} is not compact:

Lemma 5.2.8 (Upper bound property for probabilistic consolidation). *Let $\Gamma \in \mathbb{K}_c$ be a canonical base. For every consistent $\Delta \subseteq Cn_{ew}(\Gamma)$, there is a Δ' such that $\Delta \subseteq \Delta' \subseteq Cn_{ew}(\Gamma)$ and $\Delta' \in Cn_{ew}(\Gamma) \perp^\perp$.*

The Lemma below states that besides the remainder set $Cn_{ew}(\Gamma) \perp^\perp$ being well-defined, it is characterised by the maximal consolidations. This result reveals a clear connection between partial meet constructions, via remainder sets, and maximal consolidation operators — *i.e.*, satisfying (Pareto-Optimality).

Lemma 5.2.9. *Let $\Gamma \in \mathbb{K}_c$ be a canonical base. $\Psi \in \mathbb{K}_c$ is a maximal consolidation of Γ iff $Cn_{ew}(\Psi) \in Cn_{ew}(\Gamma)\perp^{\perp}$.*

Now a representation theorem can be proved:

Theorem 5.2.10. *Consider a base $\Gamma \in \mathbb{K}_c$. An operation $\Gamma!$ satisfies (Success), (Inclusion*) and (Relevance*), for $Cn^\star = Cn_{ew}$, iff $\Gamma! = \bigcap \gamma(Cn_{ew}(\Gamma)\perp^{\perp})$, for some selection function γ.*

And then it follows the first result connecting maximal consolidation operators and the AGM theory, in which $\circ$ denotes function composition:

Corollary 5.2.11. *Consider the consequence operation $Cn_{ew} = Cn^\star$, a maximal consolidation operator $\mathcal{C}$ and a consolidation operation ! for each $\Gamma \in \mathbb{K}_c$ defined as $\Gamma! = Cn_{ew} \circ \mathcal{C}(\Gamma)$. For each canonical base $\Gamma \in \mathbb{K}_c$, $\Gamma!$ satisfies (Success), (Inclusion*) and (Relevance*).*

This result relates a consolidation operation ! that returns the closure of a maximal consolidation under Cn_{ew} to a partial meet construction. From Lemma 5.2.9, it is easy to see that any consolidation operation that returns the Cn_{ew}-closure of a maximal consolidation can be seen as a partial meet construction that selects a single element out of the remainder set — a maxichoice contraction. Hence, we define the starred version of (Fullness):

Postulate 5.2.12 (Fullness*). *If $\beta \in Cn^\star(B)$ and $\beta \notin B - \alpha$, then $\alpha \notin Cn(B - \alpha)$ and $\alpha \in Cn((B - \alpha) \cup \{\beta\})$.*

Corollary 5.2.13. *Consider the consequence operation $Cn_{ew} = Cn^\star$. The consolidation operation $\Gamma!$, for any $\Gamma \in \mathbb{K}_c$, satisfies (Success), (Inclusion*) and (Fullness*) iff there is a maximal consolidation operator $\mathcal{C}$ such that $\Gamma! = Cn_{ew} \circ \mathcal{C}(\Gamma)$ for all $\Gamma \in \mathbb{K}_c$.*

If the input base is already closed under Cn_{ew} (*i.e.*, $Cn_{ew}(\Gamma) = \Gamma$)), we can link any maximal consolidation operator $\mathcal{C}$ to the original (not starred) postulates, through a canonical base Ψ such that $Cn_{ew}(\Psi) = \Gamma$. To do that, we firstly need a way to get such canonical base Ψ:

Definition 5.2.14. For any probabilistic knowledge base Γ such that $Cn_{ew}(\Gamma) = Cn_{ew}(\Psi)$ for a canonical $\Psi \in \mathbb{K}_c$[1]:

$$Cn_{ew}^{-1}(\Gamma) = \{P(\varphi|\psi) \geq q \in \Gamma | \text{ there is no } P(\varphi|\psi) \geq q' \in \Gamma \text{ with } q' > q\}$$

Now the second way to relate a maximal consolidation operator to the partial meet construction follows:

Corollary 5.2.15. *Consider a base $\Gamma = Cn_{ew}(\Gamma)$ in $\mathbb{K}$. The consolidation operation $\Gamma!$, for any $\Gamma = Cn_{ew}(\Gamma)$ in $\mathbb{K}$, satisfies (Success), (Inclusion) and (Fullness) iff there is a maximal consolidation operator $\mathcal{C}$ such that $\Gamma! = Cn_{ew} \circ \mathcal{C} \circ Cn_{ew}^{-1}(\Gamma)$ for all $\Gamma = Cn_{ew}(\Gamma)$ in $\mathbb{K}$.*

From Theorem 5.2.10, it follows that any consolidation operation $\Gamma! = Cn_{ew} \circ \mathcal{C} \circ Cn_{ew}^{-1}(\Gamma)$, for a Cn_{ew}-closed Γ, is a partial meet consolidation $\Gamma! = \bigcap \gamma(\Gamma\perp^{\perp})$. Indeed, $\Gamma! = Cn_{ew}(\Psi)$, for some maximal consolidation Ψ of $Cn_{ew}^{-1}(\Gamma)$.

Until now, we have founded on the AGM paradigm two kinds of consolidation operations based on a maximal consolidation operator $\mathcal{C}$:

[1] This avoid ill-defined cases, such as $\Gamma = \{P(\varphi) \geq q | q < 1\}$.

- $Cn_{ew} \circ \mathcal{C}(\Gamma)$ for each canonical base $\Gamma \in \mathbb{K}_c$ returns a base closed under Cn_{ew}, satisfying (Success), (Inclusion*) and (Fullness*);
- $Cn_{ew} \circ \mathcal{C} \circ Cn_{ew}^{-1}(\Gamma)$ each base $\Gamma \in \mathbb{K}$ closed under Cn_{ew} returns a base closed under Cn_{ew}, satisfying (Success), (Inclusion) and (Fullness).

However, we have still not characterised maximal consolidation operators themselves with postulates or partial meet constructions. Corollary 5.2.15 points a way to do that via the starred closure of the contraction operator, which motivates a generalisation of the contraction concept.

5.3 Liftable Contractions

Given a contraction operator $-$ for a belief base B over a general language $\mathcal{L}$ and a consequence operation $Cn^\star$, we define the *starred closure* of $-$ as the operator $\overset{\star}{-}$ for $Cn^\star(B)$ such that $Cn^\star(B) \overset{\star}{-} \alpha = Cn^\star(B - \alpha)$. That is, if $-$ takes a formula α and returns $B' = B - \alpha$, the operator $\overset{\star}{-}$ takes α and returns $Cn^\star(B')$. The starred closure $!^\star$ of a consolidation operation $!$ can be defined through $Cn^\star(B)!^\star = Cn^\star(B!)$.

The (Logical Inclusion) postulate can be seen as a restriction on the Cn-closure of a contraction operator. That is, taking $Cn^\star = Cn$, (Logical Inclusion) requires that the starred closure of the operator satisfy (Inclusion): $Cn^\star(B) \overset{\star}{-} \alpha = Cn^\star(B - \alpha) \subseteq Cn^\star(B)$. Similarly, if $Cn^\star$ were the identity function, (Inclusion) can be viewed as a restriction on the starred operator as well, for $Cn^\star(B) = B$. For a general consequence operation $Cn^\star$, we can "lift" in this way the (Inclusion) postulate to the starred closure of the operator:

Postulate 5.3.1 (***Lifted Inclusion**). *The operator $-$ for a base B is such that $\overset{\star}{-}$ for $Cn^\star(B)$ satisfies (Inclusion).*

Any postulate can be similarly lifted, but note that (Success) is equivalent to (*Lifted Success), given some assumptions:

Proposition 5.3.2. *If Cn satisfies monotonicity and idempotency, and $Cn^\star$, inclusion an Cn-dominance, (Success) is equivalent to (*Lifted Success).*

We call an operator satisfying (*Lifted Inclusion) and (*Lifted Success) a *liftable contraction.* As we pointed out, contractions and pseudo-contractions are particular cases of liftable contractions, where $Cn^\star(B) = B$ and $Cn^\star(B) = Cn(B)$, respectively. Similarly, we can lift (Relevance), (Uniformity) and (Fullness):

Postulate 5.3.3 (***Lifted Relevance**). *The operator $-$ for a base B is such that $\overset{\star}{-}$ for $Cn^\star(B)$ satisfies (Relevance).*

Postulate 5.3.4 (***Lifted Uniformity**). *The operator $-$ for a base B is such that $\overset{\star}{-}$ for $Cn^\star(B)$ satisfies (Uniformity).*

Postulate 5.3.5 (***Lifted Fullness**). *The operator $-$ for a base B is such that $\overset{\star}{-}$ for $Cn^\star(B)$ satisfies (Fullness).*

The following result shows how the lifted postulates relate to the starred ones:

Proposition 5.3.6. *Consider a consequence operation Cn that satisfies monotonicity and idempotence and a Tarskian consequence operation $Cn^\star$ satisfying Cn-dominance. The postulate of (Inclusion*) is equivalent to (*Lifted Inclusion). The postulate of (Relevance*) implies (*Lifted Relevance). The postulate of (Uniformity*) implies (*Lifted Uniformity). The postulate of (Fullness*) implies (*Lifted Fullness).*

To guarantee equivalence between the lifted and the starred postulates, it suffices to require the closure of the resulting base under $Cn^\star$:

Proposition 5.3.7. *Consider a consequence operation Cn that satisfies monotonicity and idempotence and a Tarskian consequence operation $Cn^\star$ satisfying Cn-dominance. If (Enforced Closure*) is satisfied, then (*Lifted Relevance) is equivalent to (Relevance*), (*Lifted Uniformity) is equivalent to (Uniformity*) and (*Lifted Fullness) is equivalent to (Fullness*).*

To link the lifted postulates to partial meet constructions for the starred closure, we define:

Definition 5.3.8. We say a liftable contraction is a *liftable partial meet contraction (consolidation)* if its starred closure is a partial meet contraction (consolidation); we say it is a *liftable maxichoice contraction (consolidation)* if its starred closure is a maxichoice contraction (consolidation).

The implementation of a liftable partial meet contraction is only partially determined through its starred closure, unless we require the closure of the resulting base under $Cn^\star$:

Proposition 5.3.9. *Consider a monotonic consequence operation Cn that satisfies the upper bound property and a Tarskian consequence operation $Cn^\star$ satisfying Cn-dominance. A contraction operator $- : \mathcal{L} \to 2^{\mathcal{L}}$ for a base $B \subseteq 2^{\mathcal{L}}$ satisfies the *lifted versions of the postulates of success, inclusion, relevance (fullness) and uniformity iff it is a liftable partial meet contraction (liftable maxichoice contraction). Furthermore, $-$ additionally satisfies (Enforced Closure*) iff $B - \alpha = \bigcap \gamma(Cn^\star(B) \perp \alpha)$ (and $\gamma(Cn^\star(B) \perp \alpha)$ is a singleton) for all $\alpha \in \mathcal{L}$ and for some selection function γ.*

For consolidations, it is always the case that $\alpha = \perp$, and one can drop (*Lifted Uniformity) from the result above:

Proposition 5.3.10. *Consider a monotonic consequence operation Cn that satisfies the upper bound property and a Tarskian consequence operation $Cn^\star$ satisfying Cn-dominance. A consolidation operation $B!$ satisfies the *lifted versions of the postulates of success, inclusion, and relevance (fullness) iff it is a liftable partial meet consolidation (liftable maxichoice consolidation). Furthermore, if $B!$ satisfies (Enforced Closure*), then $B! = \bigcap \gamma(Cn^\star(B) \perp^\perp)$ for some selection function γ ($B! \in (Cn^\star(B) \perp^\perp)$).*

Any maximal consolidation operator satisfies the *lifted postulates, including (*Lifted Fullness), but not (Enforced Closure*). Hence, from Proposition 5.3.10, maximal consolidations are liftable maxichoice consolidations. To uniquely characterise maximal consolidations, we need to require that output bases be canonical. Given a general consequence operation $Cn^\star$ for a language $\mathcal{L}$, we say a *canonical form* is a function $f_\star : 2^{\mathcal{L}} \to 2^{\mathcal{L}}$ such that $Cn^\star(B) = Cn^\star(f_\star(B))$, and $Cn^\star(B) = Cn^\star(B')$ implies $f_\star(B) = f_\star(B')$. Note that these conditions also imply $f_\star(f_\star(B)) = f_\star(B)$. For any base B, we say $f_\star(B)$ is its canonical form (regarding $f_\star$). Given a canonical form $f_\star$ for a consequence operation $Cn^\star$, we can define:

Postulate 5.3.11 ($f_\star$**-Canonicity**). $B - \alpha = f_\star(B - \alpha)$.

With this extra postulate, the partial meet construction for the liftable contraction is uniquely determined by the selection function γ, and (*Lifted Fullness) forces γ to return a singleton:

Proposition 5.3.12. *Consider a consequence operation Cn that satisfies monotonicity, idempotence and the upper bound property and a Tarskian consequence operation $Cn^\star$ satisfying Cn-dominance. An operator $- : \mathcal{L} \to 2^{\mathcal{L}}$ for a base B is a liftable partial meet contraction (liftable maxichoice contraction) and satisfies ($f_\star$-Canonicity) iff, for any formula $\alpha \in \mathcal{L}$, $B - \alpha = f_\star(\bigcap \gamma(Cn^\star(B)\bot\alpha))$ for some selection function γ (and $\gamma(Cn^\star(B)\bot\alpha)$ is a singleton).*

Corollary 5.3.13. *Consider a consequence operation Cn that satisfies monotonicity, idempotence and the upper bound property and a Tarskian consequence operation $Cn^\star$ satisfying Cn-dominance. An operation $B!$ is a liftable partial meet consolidation (liftable maxichoice consolidation) and satisfies ($f_\star$-Canonicity) iff, $B! = f_\star(\bigcap \gamma(Cn^\star(B)\bot^\perp))$ for some selection function γ (and $\gamma(Cn^\star(B)\bot^\perp)$ is a singleton).*

As the function Cn_{ew}^{-1} satisfies the properties of a canonical form $f_\star$ for the consequence operation Cn_{ew}, we can say that maximal consolidations satisfy the corresponding ($f_\star$-Canonicity) — in fact, any $\Gamma \in \mathbb{K}_c$ is such that $Cn_{ew}^{-1}(\Gamma) = \Gamma$. Even though the probabilistic logic is not compact, using Lemma 5.2.8 we can prove some representation results:

Lemma 5.3.14. *Consider $Cn^\star = Cn_{ew}$ and $f_\star = Cn_{ew}^{-1}$. The operation $\Gamma!$ satisfies (*Lifted Success), (*Lifted Inclusion), (*Lifted Fullness) and ($f_\star$-Canonicity) for all $\Gamma \in \mathbb{K}_c$ iff $\mathcal{C} : \mathbb{K}_c \to \mathbb{K}_c$, with $\mathcal{C}(\Gamma) = \Gamma!$ for all $\Gamma \in \mathbb{K}_c$, is a maximal consolidation operator.*

If only (*Lifted Relevance) is required, and not (*Lifted Fullness), we obtain constructions that select the infimum of the lower bounds in a set of maximal consolidations.

Definition 5.3.15. Consider a base $\Gamma \in \mathbb{K}_c$ and a set of vectors $Q \subseteq [0,1]^{|\Gamma|}$. We define $\inf\{\Lambda_\Gamma(q) | q \in Q\} = \Lambda_\Gamma(q')$, where $q' \in [0,1]^{|\Gamma|}$ is such that $q'_i = \inf\{q_i | q \in Q\}$ for every $1 \leq i \leq |\Gamma|$.

Lemma 5.3.16. *Consider $Cn^\star = Cn_{ew}$, $f_\star = Cn_{ew}^{-1}$, a $\Gamma \in \mathbb{K}_c$. $\Gamma!$ satisfies (*Lifted Success), (*Lifted Inclusion), (*Lifted Relevance) and ($f_\star$-Canonicity) iff $\Gamma! = \inf M$, for a set M of maximal consolidations of Γ.*

The set M can be defined through the remainder set, as the elements of the former (maximal consolidations) are canonical forms of elements of the latter: $M = \{f_\star(\Psi) | \Psi \in \gamma(Cn^\star(\Gamma)\bot^\perp)\}$. The *full meet* criterion is met when M is the set of all maximal consolidations.

By leaving the maxichoice contraction, allowing M to be a set of maximal consolidations, we do not need methods that specify a single one, but only a set of preferred ones. We said in section 4.3 that minimizing the p-norm distance does not necessarily yield a unique maximal consolidation; that is, $\mathcal{C}_p$ requires additional selection criteria to be a well-defined function from bases to bases. Nevertheless, minimising these distances (d_p) leads to a set D_p of preferred, closest maximal consolidations, inducing a selection function γ. To perform a consolidation, we can employ these distances to construct selection functions γ for the partial meet construction, which corresponds to take the minimum lower bounds of the corresponding maximal consolidations.

Definition 5.3.17. *Consider a canonical base* $\Gamma = \Lambda_\Gamma(q)$*. For any* $p \in \overline{\mathbb{N}}_{>0}$*,*

$$\gamma_p(Cn_{ew}(\Gamma)\perp^\perp) = \{Cn_{ew}(\Lambda_\Gamma(q)) \in Cn_{ew}(\Gamma)\perp^\perp | d_p(q, q') \text{ is minimum}\}\,.$$

Example 5.3.18. Consider that we want to consolidate the canonical base $\Gamma = \{P(\varphi) \leq 0.3, P(\varphi) \geq 0.6\}$ through a liftable partial meet contraction by the contradiction, returning a canonical base (in practice, we just want to change the probability bounds). Firstly, we compute the closure of Γ under Cn_{ew}:

$$Cn_{ew}(\Gamma) = Cn_{ew}(\{P(\varphi) \leq 0.3\}) \cup Cn_{ew}(\{P(\varphi) \geq 0.6\})\,.$$

Then the remainder set $Cn_{ew}(\Gamma)\perp^\perp$ can be defined through the (Cn_{ew})-closure of the maximal consolidations:

$$Cn_{ew}(\Gamma)\perp^\perp = \{Cn_{ew}(\{P(\varphi) \leq q_1, P(\varphi) \geq q_2\}) | q_1 = q_2, q_1 \in [0.3, 0.6]\}\,.$$

Now we need a selection function γ. Consider for instance γ_p for $p = 1$, $p = 2$ and $p = \infty$:

$$\gamma_1(Cn_{ew}(\Gamma)\perp^\perp) = Cn_{ew}(\Gamma)\perp^\perp\,,$$

$$\gamma_2(Cn_{ew}(\Gamma)\perp^\perp) = Cn_{ew}(\{P(\varphi) \leq 0.45, P(\varphi) \geq 0.45\})\,,$$

$$\gamma_\infty(Cn_{ew}(\Gamma)\perp^\perp) = Cn_{ew}(\{P(\varphi) \leq 0.45, P(\varphi) \geq 0.45\})\,.$$

When $p = 1$, any $q_1, q_2 \in [0.3, 0.6]$ with $q_1 = q_2$ is such that $d_1(\langle q_1, q_2\rangle, \langle 0.3, 0.6\rangle) = |q_1 - 0.3| + |q_1 - 0.6| = 0.3$ [2]. Hence, γ_1 selects all elements in the remainder set. Nonetheless, $d_2(\langle q_1, q_2\rangle, \langle 0.3, 0.6\rangle)$ has a single minimum with $q_1 = q_2$ at $q_1 = q_2 = 0.45$, and γ_2 selects only the (Cn_{ew}) closure of the corresponding maximal consolidation. The same happens to γ_∞.

To consolidate $Cn_{ew}(\Gamma)$, we take the intersection of the selected bases for each selection function:

$$Cn_{ew}(\Gamma)!_1^\star = \bigcap \gamma_1(Cn_{ew}(\Gamma)\perp^\perp) = \bigcap Cn_{ew}(\Gamma)\perp^\perp = Cn_{ew}(\{P(\varphi) \leq 0.6, P(\varphi) \geq 0.3\})$$

$$Cn_{ew}(\Gamma)!_2^\star = \bigcap \gamma_2(Cn_{ew}(\Gamma)\perp^\perp) = \bigcap Cn_{ew}(\{P(\varphi) \leq 0.45, P(\varphi) \geq 0.45\})$$

$$Cn_{ew}(\Gamma)!_2^\star = Cn_{ew}(\{P(\varphi) \leq 0.45, P(\varphi) \geq 0.45\})$$

$$Cn_{ew}(\Gamma)!_\infty^\star = Cn_{ew}(\Gamma)!_2^\star$$

Note that, in this particular case, $Cn_{ew}(\Gamma)!_2^\star$ and $Cn_{ew}(\Gamma)!_\infty^\star$ are maxichoice consolidations, while $Cn_{ew}(\Gamma)!_1^\star$ is a full meet consolidation.

As $Cn_{ew}(\Gamma)!_p^\star$, for any $p \in \{1, 2, \infty\}$, is a partial meet consolidation, any $\Gamma!_p$ such that $Cn_{ew}(\Gamma!_p) = Cn_{ew}(\Gamma)!_p^\star$ will be a liftable partial meet consolidation. To satisfy ($f_\star$-canonicity), we take the canonical form:

$$\Gamma!_1 = Cn_{ew}^{-1}(\Gamma!_1^\star) = \{P(\varphi) \leq 0.6, P(\varphi) \geq 0.3\}$$

$$\Gamma!_2 = Cn_{ew}^{-1}(\Gamma!_2^\star) = \{P(\varphi) = 0.45\}$$

$$\Gamma!_\infty = Cn_{ew}^{-1}(\Gamma!_\infty^\star) = \{P(\varphi) = 0.45\}$$

[2] Technically, $P(\varphi) \leq 0.3$ and $P(\varphi) \leq q_1$ abbreviate $P(\neg\varphi) \geq 1 - 0.3$ and $P(\neg\varphi) \geq 1 - q_1$, respectively, but note that $|1 - q_1 - (1 - 0.3)| = |0.3 - q_1| = |q_1 - 0.3|$.

Due to their starred closures, $\Gamma!_2$ and $\Gamma!_\infty$ are liftable maxichoice consolidations, whereas $\Gamma!_1$ is a liftable full meet consolidation.

Recall that $P(\varphi) \leq q_1$ abbreviates $P(\neg\varphi) \geq 1 - q_1$. Thus, $\Gamma!_1$ is taking the minimum lower bounds from the set of maximal consolidations: $1 - q_1 = 1 - 0.6 = 0.4$ and $q_2 = 0.3$. The operations $\Gamma!_2$ and $\Gamma!_\infty$ return a single preferred maximal consolidation. In the end, by consolidating the closure and taking the canonical form, we are just changing the probability bounds. □

Once we have fully characterised the consolidation of canonical probabilistic knowledge bases, by employing — and extending — the AGM approach, a compelling question is how to generalise these results to probabilistic contraction, if possible, which we tackle in the next section.

5.4 Probabilistic Contraction

In order to capture the properties of a maximal consolidation operator, we lifted the postulates for contractions in a general logic, pointing out that they characterise a partial meet contraction if the corresponding consequence operation satisfies the upper bound property. Using Lemma 5.2.8, we derived results to characterise (liftable) contractions of canonical probabilistic bases by the contradiction ($\Gamma - \bot$) through their starred closure ($Cn_{ew}(\Gamma) \overset{\star}{-} \bot$), but not by arbitrary formulas $\alpha \in \mathcal{L}^P_{X_n}$. In fact, contracting a probabilistic knowledge base $\Gamma = Cn_{ew}(\Gamma)$ by an arbitrary conditional $\alpha \in \mathcal{L}^P_{X_n}$ cannot satisfy (Relevance) without violating (Success). To better understand this, it helps to recall the non-compactness result for the probabilistic logic, adapted from (Ognjanovic and Raškovic, 2000):

Proposition 5.4.1. *Consider the base $\Gamma \in \mathbb{K}$ such that $\Gamma = \{P(\varphi) \geq q | q \in [0,1]$ and $q < 1\}$, for a non-tautological, consistent φ. $P(\varphi) \geq 1 \in Cn_{Pr}(\Gamma)$ but $P(\varphi) \geq 1 \notin Cn_{Pr}(\Psi)$ for every finite $\Psi \subseteq \Gamma$.*

Although proving non-compactness, this result by itself does not undermine contraction via partial meet, but shows how (Success) conflicts with (Relevance). Suppose a contraction of the Γ above by $\alpha = P(\varphi) \geq 1$, $\Gamma - \alpha$, satisfies both (Success) and (Relevance). As $\alpha \in Cn_{Pr}(\Gamma)$, (Success) entails $\Gamma \not\subseteq \Gamma - \alpha$, for Cn_{Pr} is monotonic, and there is some $\beta \in \Gamma \setminus (\Gamma - \alpha)$. By (Relevance), there must be a Ψ such that $\Gamma - \alpha \subseteq \Psi \subseteq \Gamma$, $\alpha \notin Cn_{Pr}(\Psi)$ and $\alpha \in Cn_{Pr}(\Psi \cup \{\beta\})$. As $\alpha \notin Cn_{Pr}(\Psi)$, there is a π satisfying Ψ such that $P_\pi(\varphi) = q < 1$. As $\beta \in \Gamma$, $\beta = P(\varphi) \geq q'$, for some $q' < 1$. Let $q_{max} < 1$ be the maximum between $q < 1$ and $q' < 1$. Any π such that $P_\pi(\varphi) = q_{max}$ (which is feasible, for φ is consistent) satisfies Ψ and β, but not α. Hence $\alpha \notin Cn_{Pr}(\Psi \cup \{\beta\})$, a contradiction. By generalizing this argument, it seems that any contraction (satisfying the success postulate) of a probabilistic base shall violate (Relevance).

To prove that partial meet does not work for contracting probabilistic bases by arbitrary conditionals, we need a stronger result, showing that the upper bound property in general does not hold for Cn_{Pr}:

Proposition 5.4.2. *Let $\Gamma = Cn_{ew}\{P(\varphi) \geq q\}$, for some $q \in (0,1]$ and a consistent φ. There is no $\Psi \subseteq \Gamma$ such that $P(\varphi) \geq q \notin Cn_{Pr}(\Psi)$ and $P(\varphi) \geq q \in Cn_{Pr}(\Psi')$ for any Ψ' such that $\Psi \subsetneq \Psi' \subseteq \Gamma$.*

As a consequence of the result above, the remainder set $\Gamma \bot P(\varphi) \geq q$ would be empty, and the corresponding partial meet contraction, ill-defined: $\gamma(\Gamma \bot P(\varphi) \geq q) = \Gamma$ and (Success) is violated.

To understand how probabilistic consolidation works fine via partial meet, respecting the AGM postulates, while contraction in general does not, we need to extend the probabilistic language and its semantics. Let $\bar{\mathcal{L}}^P_{X_n}$ be the language containing the *negated conditionals*: $\bar{\mathcal{L}}^P_{X_n} = \{\neg\alpha | \alpha \in \mathcal{L}^P_{X_n}\}$. The semantics for the language $\bar{\mathcal{L}}^P_{X_n} \cup \mathcal{L}^P_{X_n}$ is given by the semantics for $\mathcal{L}^P_{X_n}$ with an extra rule for negation: a probability interpretation π satisfies $\neg\alpha \in \bar{\mathcal{L}}^P_{X_n}$ if π does not satisfy $\alpha \in \mathcal{L}^P_{X_n}$. The consequence operation Cn_{Pr} can be extended accordingly. We use $P(\varphi|\psi) > q$ and $P(\varphi|\psi) < q$ to denote $\neg P(\varphi|\psi) \leq q$ and $\neg P(\varphi|\psi) \geq q$, respectively.

With this extended logic in mind, we can now generalise Lemma 5.2.8:

Theorem 5.4.3 (Upper bound)**.** *Let $\Gamma = Cn_{ew}(\Gamma)$ be a probabilistic knowledge base in $\mathbb{K}$, and $\alpha \in \bar{\mathcal{L}}^P_{X_n}$, a negated conditional. For every $\Delta \subseteq \Gamma$ such that $\alpha \notin Cn_{Pr}(\Delta)$, there is a $\Psi \in (\Gamma \perp \alpha)$ such that $\Delta \subseteq \Psi$.*

From which follows that the scope of the well-behaved probabilistic contraction is $\bar{\mathcal{L}}^P_{X_n}$ [3]. With this in mind, we relax in this section the definition of a contraction operator $-$ to allow a domain that is strictly smaller than the language, $\bar{\mathcal{L}}^P_{X_n} \subsetneq \bar{\mathcal{L}}^P_{X_n} \cup \mathcal{L}^P_{X_n}$, in order to obtain representation results:

Corollary 5.4.4. *The operator $- : \bar{\mathcal{L}}^P_{X_n} \to \mathbb{K}$ for a $\Gamma = Cn_{ew}(\Gamma)$ in $\mathbb{K}_c$ satisfies (Success), (Inclusion), (Relevance) and (Uniformity) iff $\Gamma - \alpha = \bigcap \gamma(\Gamma \perp \alpha)$ for all $\alpha \in \bar{\mathcal{L}}^P_{X_n}$ and for some selection function γ.*

To see that this result holds for consolidation, just take $\alpha = P(\perp) > 0 \in \bar{\mathcal{L}}^P_{X_n}$. That is, we can contract by $\perp = P(\perp) \geq 1$ just because it is equivalent to $\perp' = P(\perp) > 0$. However, it is not the case that any $P(\varphi|\psi) \geq q$ has an equivalent $P(\varphi'|\psi') > q'$.

Since the contraction of a probabilistic base $\Gamma = Cn_{ew}(\Gamma)$ by negated conditionals is well-defined, so is the corresponding liftable contraction. Consequently, Propositions 5.3.9 and 5.3.12 also hold in this case. Similarly, Lemma 5.3.16 can be generalised to contractions by arbitrary negated conditionals:

Theorem 5.4.5. *The operator $- : \bar{\mathcal{L}}^P_{X_n} \to \mathbb{K}$ for a $\Gamma \in \mathbb{K}_c$ satisfies the *lifted versions of success, inclusion, relevance, uniformity and ($f_\star$-Canonicity), for $Cn^\star = Cn_{ew}$ and $f_\star = Cn^{-1}_{ew}$, iff, for all $\alpha \in \bar{\mathcal{L}}^P_{X_n}$, $\Gamma - \alpha = \inf M$, where $M = \{\Psi \in \mathbb{K}_c | Cn_{ew}(\Psi) \in \gamma(Cn_{ew}(\Gamma) \perp \alpha)\}$ for some selection function γ.*

The set M in the result above can be defined by some distance between probability bounds vector, as we did for consolidation. That is, we can generalise the selection function γ_p to contractions by an arbitrary negated conditional:

Definition 5.4.6. *Consider a canonical base $\Gamma = \Lambda_\Gamma(q)$ and a $\alpha \in \bar{\mathcal{L}}^P_{X_n}$. For any integer $p \geq 1$ and $p = \infty$, if $\alpha \in Cn_{Pr}(\varnothing)$, then $\gamma_p(Cn_{ew}(\Gamma) \perp \alpha) = Cn_{ew}(\Gamma) \perp \alpha$, else*

$$\gamma_p(Cn_{ew}(\Gamma) \perp \alpha) = \{Cn_{ew}(\Lambda_\Gamma(q)) \in Cn_{ew}(\Gamma) \perp \alpha | d_p(q, q') \text{ is minimum}\}.$$

Example 5.4.7. Consider the canonical base $\Gamma = \{P(\varphi) \geq 0.3, P(\psi) \geq 0.6, P(\neg\varphi \vee \neg\psi) \geq 1\}$. Suppose we want to contract it by $\alpha = \neg P(\varphi \vee \psi) \leq 0.6$ (equivalently, $\alpha = P(\varphi \vee \psi) > 0.6$) using

[3] Note that this contraction is regarding the standard semantics (Cn_{Pr}), and does not apply to the maximum entropy entailment.

a liftable partial meet contraction and returning a canonical base. Once again, we want to change the lower bounds in Γ in order not to imply α, but we will show how to perform this via AGM's partial meet construction on the closure $Cn_{ew}(\Gamma)$.

Firstly, note that $P(\neg\varphi \vee \neg\psi) \geq 1 \in \Gamma$ is equivalent to $P(\varphi \wedge \psi) \leq 0$, meaning that φ and ψ are disjoint. Thus, any π satisfying Γ must be such that

$$P_\pi(\varphi \vee \psi) = P_\pi(\varphi) + P_\pi(\psi) \geq 0.3 + 0.6 = 0.9 > 0.6\,.$$

Consequently, $\alpha \in Cn_{Pr}(\Gamma)$ and $\alpha \in Cn_{Pr}(Cn_{ew}(\Gamma))$. To do a partial meet on the closure $Cn_{ew}(\Gamma)$, we compute its maximal subsets not implying (via Cn_{Pr}) α, the remainder set:

$$Cn_{ew}(\Gamma)\bot\alpha = \{Cn_{ew}(\{P(\varphi) \geq q_1, P(\psi) \geq q_2, P(\neg\varphi \vee \neg\psi) \geq q_3\})| \text{ where}$$

$$q_1 + q_2 + q_3 = 1.6, q_1 \in [0, 0.3], q_2 \in [0, 0.6], q_3 \in [0, 1]\}$$

The restriction $q_1+q_2+q_3 \leq 1.6$ prevents an element from the remainder set from implying α, while $q_1 + q_2 + q_3 = 1.6$ selects such maximal elements. The other restrictions, $q_1 \in [0, 0.3]$, $q_2 \in [0, 0.6]$ and $q_3 \in [0, 1]$, guarantee each element of the remainder set be a subset of $Cn_{ew}(\Gamma)$.

Applying the selection functions γ_1, γ_2 to the remainder set, we have:

$$\gamma_1(Cn_{ew}(\Gamma)\bot\alpha) = Cn_{ew}(\Gamma)\bot\alpha$$

$$\gamma_2(Cn_{ew}(\Gamma)\bot\alpha) = \{Cn_{ew}(\{P(\varphi) \geq 0.2, P(\psi) \geq 0.5, P(\neg\varphi \vee \neg\psi) \geq 0.9\}$$

The ineffectiveness of γ_1 is explained by the fact that $q_1+q_2+q_3 = 1.6$ implies $d_1(\langle q_1, q_2, q_3\rangle, \langle 0.3, 0.6, 1\rangle) = 0.3$ for any lower bounds $q_1 \in [0, 0.3], q_2 \in [0, 0.6], q_3 \in [0, 1]$. To minimise $d_2(\langle q_1, q_2, q_3\rangle, \langle 0.3, 0.6, 1\rangle)$, each lower bound has to be decreased by the same amount, 0.1. One can also note that γ_∞ would yield the same result as γ_2.

A partial meet contraction $Cn_{ew}(\Gamma) \overset{\star}{-}_p \alpha$, for $p = 1$ and $p = 2$, can be obtained taking the intersection of the selected elements:

$$Cn_{ew}(\Gamma) \overset{\star}{-}_1 \alpha = \bigcap \gamma_1(Cn_{ew}(\Gamma)\bot\alpha) = Cn_{ew}(\{P(\varphi) \geq 0, P(\psi) \geq 0.3, P(\neg\varphi \vee \neg\psi) \geq 0.7\})$$

$$Cn_{ew}(\Gamma) \overset{\star}{-}_2 \alpha = \bigcap \gamma_2(Cn_{ew}(\Gamma)\bot\alpha) = Cn_{ew}(\{P(\varphi) \geq 0.2, P(\psi) \geq 0.5, P(\neg\varphi \vee \neg\psi) \geq 0.9\})$$

We point out that $Cn_{ew}(\Gamma) \overset{\star}{-}_1 \alpha$ is a full meet contraction and $Cn_{ew}(\Gamma) \overset{\star}{-}_2 \alpha$ is a maxichoice contraction. To see how the full meet contraction was obtained, note three particular elements in the remainder set, whose intersection is $Cn_{ew}(\Gamma) \overset{\star}{-}_1 \alpha$:

$$Cn_{ew}(\{P(\varphi) \geq 0, P(\psi) \geq 0.6, P(\neg\varphi \vee \neg\psi) \geq 1\})$$

$$Cn_{ew}(\{P(\varphi) \geq 0.3, P(\psi) \geq 0.3, P(\neg\varphi \vee \neg\psi) \geq 1\})$$

$$Cn_{ew}(\{P(\varphi) \geq 0.3, P(\psi) \geq 0.6, P(\neg\varphi \vee \neg\psi) \geq 0.7\})$$

To construct liftable partial meet contractions on Γ, we just need a $\Gamma - \alpha = \Psi$ such that $Cn_{ew}(\Psi) = Cn_{ew}(\Gamma) \overset{\star}{-}_p \alpha$, for $p = 1$ or $p = 2$. But we want the contracted base to be canonical as

well, thus the canonical form is taken:

$$\Gamma -_1 \alpha = \{P(\varphi) \geq 0, P(\psi) \geq 0.3, P(\neg\varphi \vee \neg\psi) \geq 0.7\}$$

$$\Gamma -_2 \alpha = \{P(\varphi) \geq 0.2, P(\psi) \geq 0.5, P(\neg\varphi \vee \neg\psi) \geq 0.9\}$$

These two operators are liftable partial meet contractions satisfying (f_*-Canonicity). In particular, they are extreme cases: a liftable full meet contraction and a liftable maxichoice contraction, respectively. □

The fact that only contractions by negated conditionals can be well characterised by the postulates and the partial meet construction within the probabilistic logic may be somewhat strange at first, but thinking about belief revision, such oddness disappears. To revise a probabilistic knowledge base by a conditional $\alpha = P(\varphi|\psi) \geq q$, one can apply these two operations, in some order: contract by $\neg\alpha$; expand by α. And as contractions by $\neg\alpha$ are indeed well-founded on the AGM paradigm, one can define revision as well. For instance, using Levi's identity, to revise a canonical base Γ by a conditional $\alpha = P(\varphi|\psi) \geq q$, one could perform a liftable contraction by $\neg\alpha$, which is a well-defined operation, followed by consistently adding α to the base.

5.5 AGM Paradigm and ε-Consolidations

In this section, we investigate how to adapt the obtained results to consolidation methods via violations, introduced in Sections 4.4. As the consolidation operator $\mathcal{C}_p^\varepsilon$ may violate (Pareto-Optimality), we consider now only the extended language $\mathcal{L}_{X_n}^\varepsilon$, due to the fact that the consolidation operator $\mathcal{V}_p$, which returns ε-bases, is Pareto-ε-optimal. The standard probabilistic consequence operator Cn_{Pr} is generalised to $\mathcal{L}_{X_n}^\varepsilon$ in the obvious way: $Cn_{Pr}(\Gamma)$ is the set of conditionals $\alpha \in \mathcal{L}_{X_n}^\varepsilon$ such that, if a probabilistic interpretation π satisfies Γ, it satisfies α.

In the same way that maximal consolidation operators can be characterised through liftable contractions using the consequence operation Cn_{ew}, we can characterise maximal ε-consolidation operators. To achieve that, we introduce a consequence operation for the language $\mathcal{L}_{X_n}^\varepsilon$, over which ε-bases are built, and its corresponding canonical form:

Definition 5.5.1. The function $Cn_\varepsilon : \mathbb{K}^\varepsilon \to \mathbb{K}^\varepsilon$ is a consequence operation such that:

- For all $\Gamma \in \mathbb{K}^\varepsilon$, $Cn_\varepsilon(\Gamma) = \bigcup\{Cn_\varepsilon(\{\alpha\})|\alpha \in \Gamma\}$;
- For all $\alpha = P(\varphi|\psi) \geq_\varepsilon q \in \mathcal{L}_{X_n}^\varepsilon$, $Cn_\varepsilon(\{\alpha\}) = \bigcup\{P(\varphi|\psi) \geq_{\varepsilon'} q|\varepsilon' \in [0,1], \varepsilon' \geq \varepsilon\}$.

Definition 5.5.2. For any probabilistic ε-base Γ such that $Cn_\varepsilon(\Gamma) = Cn_\varepsilon(\Psi)$ for a canonical $\Psi \in \mathbb{K}_c^\varepsilon$[4]:

$$Cn_\varepsilon^{-1}(\Gamma) = \{P(\varphi|\psi) \geq_\varepsilon q \in \Gamma| \text{ there is no } P(\varphi|\psi) \geq_{\varepsilon'} q \in \Gamma \text{ with } \varepsilon' < \varepsilon\}$$

While $Cn_{ew}(\alpha)$ returns the set of all conditionals on the same propositions as α with smaller lower bounds, $Cn_\varepsilon(\beta)$ is the set of ε-conditionals with greater violations.

Proposition 5.5.3. *Cn_ε is Tarskian and satisfies Cn_{Pr}-dominance.*

[4]This avoid ill-defined cases as $\Gamma = \{P(\varphi) \geq_\varepsilon q|\varepsilon > 0\}$.

For partial meet constructions on ε-bases closed under Cn_ε to be well-defined, we need the upper bound property:

Lemma 5.5.4 (Upper bound property for probabilistic ε-consolidation)**.** *Let $\Gamma \in \mathbb{K}_c^\varepsilon$ be a canonical ε-base. For every consistent $\Delta \subseteq Cn_\varepsilon(\Gamma)$, there is a Δ' such that $\Delta \subseteq \Delta' \subseteq Cn_{ew}(\Gamma)$ and $\Delta' \in Cn_\varepsilon(\Gamma)\perp^\perp$.*

The remainder set $Cn_\varepsilon(\Gamma)\perp^\perp$ can also be characterised in terms of maximal ε-consolidations:

Lemma 5.5.5. *Let $\Gamma \in \mathbb{K}_c^\varepsilon$ be a canonical ε-base. $\Psi \in \mathbb{K}_c^\varepsilon$ is a maximal ε-consolidation of Γ iff $Cn_\varepsilon(\Psi) \in Cn_\varepsilon(\Gamma)\perp^\perp$.*

Note that the function Cn_ε^{-1} also satisfies the properties of a canonical form $f_\star$, now considering the consequence operation Cn_ε, it follows that maximal ε-consolidations satisfy the corresponding ($f_\star$-Canonicity) — indeed, any $\Gamma \in \mathbb{K}_c^\varepsilon$ is such that $Cn_\varepsilon^{-1}(\Gamma) = \Gamma$. Using Lemma 5.5.4, we can prove the analogous representation results for maximal ε-consolidation operators:

Lemma 5.5.6. *Consider $Cn^\star = Cn_\varepsilon$ and $f_\star = Cn_\varepsilon^{-1}$. The operation $\Gamma!$ satisfies (*Lifted Success), (*Lifted Inclusion), (*Lifted Fullness) and ($f_\star$-Canonicity) for all $\Gamma \in \mathbb{K}_c$ iff $\mathcal{V} : \mathbb{K}_c \to \mathbb{K}_c^\varepsilon$, with $\mathcal{V}(\Gamma) = \Gamma!$ for all $\Gamma \in \mathbb{K}_c$, is a maximal ε-consolidation operator.*

If we require only (*Lifted Relevance), instead of (*Lifted Fullness), the resulting construction selects the greatest violation for each conditional $P(\varphi|\psi) \geq q \in \Gamma$ out of a set of maximal ε-consolidations. That is, one could prove for ε-consolidations a result analogous to Lemma 5.3.16.

Proceeding to adapt all the results to ε-bases, we could define selection functions for maximal ε-consolidation using d_p^ε and even define general contractions in this extended language (for negated ε-conditionals), but this is not in the scope of this work. The main concern in this chapter was to characterise these probabilistic consolidation operators with the terms and concepts of the well-established AGM framework of belief revision.

Chapter 6

Localising Inconsistency and Fixing the Postulates

Back to the main thread of this thesis — and to the content of (De Bona and Finger, 2015) —, we proceed to argue for the rejection of the postulate of (Independence) while seeking a way to reconcile the desirable properties for measuring inconsistency in probabilistic knowledge bases. In Chapter 4, we saw how minimal inconsistent sets derive from the consolidation method via discarding formulas, and different consolidation procedures were shown to underlie continuous inconsistency measures tailored to probabilistic logic. To reject (Independence), it remains to check if minimal inconsistent sets are a suitable way to characterise the primitive inconsistencies in a base to be consolidated with these intrinsically probabilistic consolidation operators — and if free conditionals capture the harmless formulas indeed. On a negative answer, instead of completely abandoning (Independence), we look for refinements of the minimal inconsistent set and free conditional concepts in order to reconcile the postulates. This is done by investigating how one should localise inconsistencies while applying consolidation operators. By "localise inconsistency", we mean "find the sets of formulas causing the inconsistency", which can be split into two tasks: to characterise such sets and to compute them. We focus on the former, for this characterisation is sufficient to rework the postulates for inconsistency measures.

Apart from analysing the postulate of (Independence), characterising the sets of formulas that are forming the inconsistency is useful on its own. In practice, an inconsistent set of pieces of information has consistent subsets, and the conflicting elements forms only a rather small portion. When a knowledge base is inconsistent, one may question which elements of the base are in fact causing the inconsistency. The answer to this question can serve different purposes, for instance:

- to assess the severity of the inconsistency by counting the number of different sets causing it or quantifying their sizes;
- to focus on a subset while repairing the inconsistency, ignoring the rest of the base;
- to find the sources of the inconsistent information (experts, databases, sensors, etc.), for they to be fixed or to reach an agreement.

The question "which subsets are causing the inconsistency?" is ill-posed, requiring some assumptions. For instance, in any monotonic logic, a base containing an inconsistent subset of formulas (a *conflict*) is also inconsistent, so any conflict could naively be called "guilty" for the inconsistency.

Nonetheless, it is clearly not the case that any formula in a conflict can be essentially used in a proof to derive the contradiction, for instance. While defining when a set is properly causing the inconsistency, or is a *primitive conflict*, one wants to avoid that any "innocent" formula, as a tautology, be blamed for the inconsistency just for being within an inconsistent set. Hence, primitive conflicts must possess further desiderata, addressing, for instance, the ends just listed above. Within classical propositional logic, for the reasons presented at the end of Section 4.2, minimal inconsistent subsets are *the* characterisation of primitive conflicts, and their use was inherited by other logics, as the probabilistic one. As we have seen, it suffices to focus on MISes in an abrupt consolidation, and free conditionals can be ignored. Nevertheless, in probabilistic logic, its proper consolidation operators lead to new characterisations of where to focus on and what to ignore while consolidating.

In this chapter, we develop definitions to capture when a conditional is essentially contributing to the inconsistency of a probabilistic knowledge base. Every definition of primitive conflict also characterises when a formula is not causing the inconsistency, so we investigate these dual concepts. We start in Section 6.1 by quickly reviewing minimal inconsistent sets and their counterpart, free conditionals, showing that these concepts do not fit quantitative consolidation (via probability changing). As a conclusion, (Independence) and (MIS-Separability) must be abandoned, then we search for weaker versions. Pointing out the close relation among consolidation operators, minimal inconsistent sets and free conditionals, we adapt the latter, in Section 6.2, to suit (quantitative) maximal consolidation operators. The dual characterisation of primitive conflict is in Section 6.3. Afterwards, this approach is employed in Section 6.4 to derive the primitive conflict — and the corresponding free conditional — characterisation related to maximal ε-consolidation operators. Finally, in Section 6.5, we exploit these new concepts to weaken the postulates of (Independence) and (MIS-Separability) in a suitable way.

6.1 Minimal Inconsistent Sets and Free Conditionals

Minimal inconsistent sets are the basis for consolidation methods via discarding formulas, as we have seen in Section 4.2. For this reason, they are the primitive conflict definition used to state (Independence) and (MIS-Separability) and form the basis of some inconsistency measures in classical logic (Hunter and Konieczny, 2005, 2006, 2008, 2010).

The focus on MISes dates at least back to the seminal work of Reiter (1987), on the diagnosis problem. Reiter investigated how to restore the consistency of a knowledge base by discarding elements from each MIS, computing thus a hitting set of their collection. Moreover, MISes play a central role in AGM belief contraction, as we have seen in Section 4.2. For all these reasons, minimal inconsistent sets can be considered the purest form of inconsistency (Hunter and Konieczny, 2008). Consequently, formulas that do not participate in the MISes of a base can be considered as "innocent" regarding its inconsistency — or *free*. Methods to compute MISes and maximal consistent subsets — thus, to find free conditionals — are not in the scope of the present study, and the main techniques are referred to in Section 4.2.

The notion of free conditional is strongly related to the idea that minimal inconsistent sets are the causes of inconsistencies, as suggested by Hunter and Konieczny (2006). Thimm (2013) says that free conditionals are "harmless", in some sense, to the consistency of a knowledge base. As explained in Section 4.2, what is behind this view is the classical way of handling inconsistency

through ruling out formulas, as Reiter (1987) proposed in his diagnosis problem and as the standard AGM paradigm of belief revision defines base contraction (see (Hansson, 1999) for a general view of the AGM paradigm). Reiter's hitting sets technique views a repair of some inconsistent set of formulas as giving up at least one element from each minimal inconsistent set. For such repair to be minimal, no free formula should be discarded. In the AGM paradigm, the consolidation of a belief base is defined as the contraction by the contradiction (Hansson, 1997), as detailed in Sections 5.1. The postulate of (Inclusion) claims that the result of a contraction is a subset of the belief base in question, and the postulate of (Core-retainment) — like (Relevance) — forces the contraction of $\perp$ to contain all free formulas of the base. Consequently, as pointed out in Section 4.2, any abrupt consolidation via partial meet or kernel can ignore free conditionals, for they are consistent with any consistent subset of the base. Nonetheless, this is not the case when a base is consolidated via probability bounds adjustment, which is a more natural way to consolidate probabilistic bases, as argued in Section 4.1:

Example 6.1.1. Recall Example 1.0.2, and suppose we are focusing only on patients of disease D; that is, assuming $P(D) = 1$. Associating the atomic propositions x_1 and x_2 to the presence of the symptoms S_1 and S_2, respectively, we construct the canonical base:

$$\Gamma = \{P(x_1 \wedge x_2) \geq 0.6, P(x_1 \wedge \neg x_2) \geq 0.5, P(x_1) \leq 0.8\}\,.$$

Firstly, we prove that $\Delta = \{P(x_1 \wedge x_2) \geq 0.6, P(x_1 \wedge \neg x_2) \geq 0.5\}$ is the only MIS in Γ. Note that $\{P(x_1 \wedge x_2) \geq 0.6, P(x_1) \leq 0.8\}$ is consistent, for such set is satisfied by the probability measure induced by the following probability mass: $\pi_1(x_1 \wedge x_2) = 0.6, \pi_1(x_1 \wedge \neg x_2) = 0.2, \pi_1(\neg x_1 \wedge x_2) = \pi_1(\neg x_1 \wedge \neg x_2) = 0.1$. To prove that $\{P(x_1 \wedge \neg x_2) \geq 0.5, P(x_1) \leq 0.8\}$ is consistent, consider the following probability mass: $\pi_2(x_1 \wedge x_2) = 0.3, \pi_2(x_1 \wedge \neg x_2) = 0.5, \pi_2(\neg x_1 \wedge x_2) = \pi_2(\neg x_1 \wedge \neg x_2) = 0.1$. Hence, all MISes of Γ must contain $\Delta = \{P(x_1 \wedge x_2) \geq 0.6, P(x_1 \wedge \neg x_2) \geq 0.5\}$, for other subsets are all consistent. Furthermore, note that Δ is inconsistent and minimal, so it is a MIS. We can conclude that Γ is the only MIS in Δ.

Now suppose that one wants to consolidate Γ by relaxing its probability bounds. To consolidate the single MIS of Γ, which is $\Delta = \{P(x_1 \wedge x_2) \geq 0.6, P(x_1 \wedge \neg x_2) \geq 0.5\}$, both lower bounds must sum at most 1; for instance $\Delta' = \{P(x_1 \wedge x_2) \geq 0.5, P(x_1 \wedge \neg x_2) \geq 0.5\}$ is a (maximal) consolidation of Δ. However, when looking at the whole base Γ, such consolidation of the single MIS yields $\Gamma' = \{P(x_1 \wedge x_2) \geq 0.5, P(x_1 \wedge \neg x_2) \geq 0.5, P(x_1) \leq 0.8\}$, which is not a consolidation of Γ at all, for Γ' is inconsistent. Note that any π satisfying Δ' must be such that $P_\pi(x_1) = P_\pi(x_1 \wedge x_2) + P_\pi(x_1 \wedge \neg x_2) = 1$, hence not satifying $P(x_1) \leq 0.8 \in \Gamma'$. □

From Example 6.1.1, we can conclude that free conditionals cannot be ignored while consolidating bases by adjusting probability bounds, and the concept of minimal inconsistent set does not capture all causes of inconsistency to be fixed when considering a quantitative consolidation. For instance, recall the situation in Example 1.0.2, formalised in Example 6.1.1, where the conditionals in the base Γ came from three different experts. By attributing the whole inconsistency to the single MIS, a knowledge engineer could ask for the experts responsible for $P(x_1 \wedge x_2) \geq 0.6$ and $P(x_1 \wedge \neg x_2) \geq 0.5$ to together reassess their opinions in a consistent way, ignoring the responsible for $P(x_1) \leq 0.8$. An agreement could result in $P(x_1 \wedge x_2) \geq 0.5, P(x_1 \wedge \neg x_2) \geq 0.5$, but, as we have seen, this do not consolidate the entire knowledge base.

As a conclusion, we can definitely reject (Independence) — and logically (MIS-Separability) —, as a legitimate rationality constraint on inconsistency measures for probabilistic logic. Instead of completely forgetting these postulates, we can look for suitable refinements of them, which depend on the definition of primitive conflict.

6.2 Refining the Free Conditional Concept

To weaken (Independence) and (MIS-Separability), we want to find a characterisation of primitive conflicts that fits quantitative consolidation operators — *i.e.*, whose consolidation imply restoring the consistency of the whole base, allowing one to ignore the rest of the base while adjusting the probability bounds to reach consistency. However, for the sake of the presentation, we start by adapting the concept of free conditional, to then independently look for a suitable primitive conflict characterisation. In the end, we prove the dual relation between the derived concepts.

A weaker form of free conditional has already been suggested in the literature. Thimm (2013) defines a *safe* conditional as one whose atomic propositions are disjoint from those in the rest of the base. We also demand that the conditional be satisfiable in order to be safe[1]. Based on safe conditionals, Thimm (2013) adapted to the probabilistic logic the following postulate:

Postulate 6.2.1 (**Weak Independence**). *If α is a safe probabilistic conditional in Γ, then $\mathcal{I}(\Gamma) = \mathcal{I}(\Gamma \setminus \{\alpha\})$.*

Hunter and Konieczny (2010) have previously suggested the same weakening for (Independence), in the classical setting, when they acknowledge that the original postulate may be too strong a property to require. It is easy to see that safe formulas may be ignored while repairing an inconsistent base via probabilities adjustment and imply no extra effort to consolidate. Hence, the postulate of (Weak Independence) is compatible with (Consistency) and (Continuity), since they are satisfied by minimal violation measures (Potyka, 2014), the whole family $\mathcal{I}_p$ (Thimm, 2013) and also $\mathcal{I}_{CRV}$:

Proposition 6.2.2. *$\mathcal{I}_{CRV}$ satisfies (Weak Independence).*

Although safe conditionals are easily recognisable, we expect them to be rare in practice, due to the natural logical dependencies among propositions within a knowledge base. Furthermore, safe conditionals are a primitive concept, not derived from a characterisation of atomic inconsistencies, and do not yield a suitable refinement of (MIS-Separability). We are looking for a stronger, more useful notion of "innocent" formulas, between safe and free, which is dual to some primitive conflict definition.

We have seen in Chapter 4 that free conditionals (and MISes) are linked to abrupt consolidations, although no neat formal connection has been provided. By formally redefining a free conditional in terms of abrupt consolidations, we shall get a definition parametrised by the consolidation procedure. By inserting then the consolidation procedures that underlie the typical probabilistic inconsistency measures, more suitable forms of free conditionals might arise, possibly yielding a compatible version of (Independence). To follow this path, some concepts are introduced:

Definition 6.2.3. Let Γ be a knowledge base in $\mathbb{K}$. An *abrupt repair* of Γ is any set $\Delta \subseteq \Gamma$ such that $\Gamma' = \Gamma \setminus \Delta$ is consistent — Γ' is an abrupt consolidation. If an abrupt repair Δ is such that,

[1] Thimm (2013) only considers conditionals $P(\varphi|\psi) = q$ such that $\varphi \wedge \psi$ and $\neg\varphi \wedge \psi$ are (classically) satisfiable, so any such conditional is also satisfiable.

for every $\Psi \subsetneq \Delta$, $\Gamma \setminus \Psi$ is inconsistent, Δ is a *minimal abrupt repair* — and $\Gamma' = \Gamma \setminus \Delta$ is a *maximal abrupt consolidation.*

We can now prove[2] a result that states different ways to define a free conditional: as being part of no minimal abrupt repairs (of all maximal abrupt consolidations) or being consistent with any abrupt consolidation. We say a conditional α (or the base $\{\alpha\}$) is *consistent with* a knowledge base Γ if the set $\Gamma \cup \{\alpha\}$ is consistent.

Theorem 6.2.4. *Consider a knowledge base $\Gamma \in \mathbb{K}$ and a probabilistic conditional $\alpha \in \Gamma$. The following statements are equivalent:*

1. *There is no minimal abrupt repair Δ of Γ such that $\alpha \in \Delta$.*
2. *For all maximal abrupt consolidations Γ' of Γ, $\alpha \in \Gamma'$.*
3. *If $\Gamma' = \Gamma \setminus \Delta$ is an abrupt consolidation of Γ (equivalently, Δ is an abrupt repair of Γ), then α is consistent with Γ'.*
4. *There is no minimal inconsistent set $\Delta \subseteq \Gamma$ such that $\alpha \in \Delta$.*

Note that the fourth statement above is the definition of free conditional given in Section 3.1. The first and the second statements are clearly dual to each other, so we have presented two new ways of equivalently defining a free conditional without mentioning minimal inconsistent sets, but using abrupt repair and abrupt consolidation. As suggested in the Section 4.2, ruling a conditional out is equivalent to relaxing the corresponding lower bound to 0 — that is why we call it an *abrupt* repair. However, a probabilistic logic allows for a more general notion of consolidation, via probability bounds adjustment, as shown in Section 4.3.

We are interested in consolidation operators that satisfy (Non-Strengthening), only decreasing the lower bounds (increasing the upper bounds). Recall that a weakening of a base is the result of relaxing its probability bounds and a maximal consolidation is a natural consolidation (a consistent weakening) where the bounds could not be less relaxed. Using these concepts, two new definitions for free conditional arise: a conditional is free if it is in any maximal consolidation; or a conditional is free if it is consistent with any natural consolidation. We can prove these definitions are actually equivalent:

Lemma 6.2.5. *Consider a knowledge base $\Gamma \in \mathbb{K}$ and a conditional $\alpha \in \Gamma$. The following statements are equivalent:*

1. *For all maximal consolidations Γ' of Γ, $\alpha \in \Gamma'$.*
2. *If Γ' is a natural consolidation of Γ, then α is consistent with Γ'.*

A modification of the free conditional concept is suggested by the comparison of Lemma 6.2.5 with Theorem 6.2.4. In order not to overload the concept of free conditional, we say these probabilistic conditionals are *innocuous*, for they are consistent with any natural consolidation of the knowledge base.

Definition 6.2.6. An *innocuous probabilistic conditional* of Γ is a probabilistic conditional $\alpha \in \Gamma$ such that, for every maximal consolidation Γ' of Γ, $\alpha \in \Gamma'$.

[2] All technical results in this section and in Section 6.3 were originally presented in (De Bona and Finger, 2015).

The difference between free and innocuous conditionals can be seen in the knowledge base from Example 6.1.1, as the following example shows.

Example 6.2.7. Consider the following knowledge base:

$$\Gamma = \{P(x_1 \wedge x_2) \geq 0.6, P(x_1 \wedge \neg x_2) \geq 0.5, P(x_1) \leq 0.8\} \ .$$

As it was claimed in the Example 6.1.1, $\{P(x_1 \wedge x_2) \geq 0.6, P(x_1 \wedge \neg x_2) \geq 0.5\}$ is the only minimal inconsistent set of Δ; so $\alpha = P(x_1) \leq 0.8$ is a free probabilistic conditional. Nonetheless, Γ has no innocuous probabilistic conditional. This can be noted through the following maximal consolidation of Γ:

$$\Gamma' = \{P(x_1 \wedge x_2) \geq 0.55, P(x_1 \wedge \neg x_2) \geq 0.45, P(x_1) \leq 1\} \ .$$

Γ' is consistent and any natural consolidation $\Psi \neq \Gamma'$ has at least one smaller lower bound; so Γ' is a maximal consolidation. But no original conditional of Γ is in Γ', so none is innocuous. Equivalently, any $\beta \in \Gamma$ is inconsistent with Γ'. An example of innocuous conditional can be given in the knowledge base $\Psi = \Gamma \cup \{P(x_2) \geq 0.3\}$, since $P(x_2) \geq 0.3$ would be consistent with any natural consolidation of Ψ. □

An innocuous probabilistic conditional of Γ is consistent with any abrupt consolidation of Γ, since such consolidation is semantically equivalent to a natural consolidation in which any relaxed probability lower bound is null; furthermore, a safe conditional of Γ is clearly consistent with any natural consolidation of Γ:

Proposition 6.2.8. *Consider a probabilistic conditional $\alpha \in \Gamma$. If α is safe, it is innocuous; if α is innocuous, it is free.*

6.3 Refining the Primitive Conflict Concept

Before investigating the sort of (Independence) postulate that can be derived from innocuous conditionals, we analyse the related form of localising conflicts, which shall enable us to restate (MIS-Separability) as well. Furthermore, localising conflicts in probabilistic logic has its own benefits, as we argued in the beginning of this chapter. To localise each conflict that has to be repaired in a base under consolidation via probability adjustment, not only the conditionals that can be ignored, we need to adapt the primitive conflict characterisation to quantitative consolidations, as just done for the definition of free conditional.

Note that the union of minimal inconsistent sets is equal to the union of minimal abrupt repairs of a knowledge base, so that it forms the complement of the set of free probabilistic conditionals. To be consistent, we should provide a definition of conflicting sets such that their union is complementary to the set of innocuous conditionals. A set with all probabilistic conditionals that are not innocuous would be inconsistent when not empty, but would not have the minimality we are looking for. Such a set would be analogous to the union of all minimal inconsistent sets, but we search for a more fundamental, atomic notion of conflict, that can be derived by analysing the consolidation properties of minimal inconsistent sets.

A minimal inconsistent set is minimal regarding set inclusion, and this is formally related to the abrupt consolidation:

Proposition 6.3.1. *A knowledge base Γ is a minimal inconsistent set iff Γ is inconsistent and there are no $\Delta_1, \ldots, \Delta_k \subsetneq \Gamma$, with $k \geq 1$, such that:*

1. $\bigcup_{i=1}^{k} \Delta_i = \Gamma$;
2. *For every $\Gamma' \subseteq \Gamma$, if $\Gamma' \cap \Delta_i$ is an abrupt consolidation of Δ_i for all $1 \leq i \leq k$, then Γ' is an abrupt consolidation of Γ.*

Intuitively, a minimal inconsistent set Γ is a conflict that cannot be analysed in smaller subsets such that abruptly consolidating them implies abruptly consolidating Γ. Starting with a single inconsistent base Γ, we can find smaller subsets Δ_i satisfying both items of 6.3.1. We can do this recursively on the inconsistent sets Δ_i until we reach unanalysable conflicts, which happens to be minimal inconsistent sets. So, abruptly consolidating these sets is abruptly consolidating Γ. Substituting natural consolidation for abrupt consolidation, we have an analogous definition of conflict:

Definition 6.3.2. A knowledge base Γ is an *inescapable conflict* if Γ is inconsistent and there are no $\Delta_1, \ldots, \Delta_k \subsetneq \Gamma$, with $k \geq 1$, such that:

1. $\bigcup_{i=1}^{k} \Delta_i = \Gamma$;
2. If Δ'_i is a natural consolidation of Δ_i for all $1 \leq i \leq k$ and $\bigcup_{i=1}^{k} \Delta'_i$ is a weakening of Γ, then $\bigcup_{i=1}^{k} \Delta'_i$ is a natural consolidation of Γ.

The extra condition in the second item of Definition 6.3.2 forces weakenings of different knowledge bases $\Delta_i, \Delta_j \subsetneq \Gamma$ with some probabilistic conditional in common to agree in that probability bound; otherwise, $\bigcup_{i=1}^{k} \Delta'_i$ would not be a knowledge base in canonical form. In other words, the second item says that if we relax the probability bounds of Γ making each Δ_i consistent, then Γ becomes consistent. Inescapable conflicts could equivalently be defined in an alternative way:

Lemma 6.3.3. *A knowledge base Γ is an inescapable conflict iff there is a weakening Γ' of Γ such that Γ' is a minimal inconsistent set.*

Lemma 6.3.3 is illustrated by Example 6.1.1, where there is a weakening that consolidates any proper subset of the knowledge base while turning the whole base into a MIS — so it was an inescapable conflict. As it happens with abrupt consolidation and MISes, to consolidate Γ, one only needs to weaken (relax) its probability bounds in such a way that each inescapable conflict is solved.

Corollary 6.3.4. *Consider two knowledge bases $\Gamma, \Gamma' \in \mathbb{K}$ such that Γ' is a weakening of Γ. If for every inescapable conflict $\Delta \subseteq \Gamma$ its weakening $\{\beta \in \Gamma' \mid \alpha \in \Delta$ and $\alpha \subseteq \beta\}$ is consistent, then Γ' is a natural consolidation of Γ.*

In other words, to consolidate a probabilistic base via relaxing its probability bounds, one can focus on inescapable conflicts, ignoring the rest of the base. Considering again that each conditional in the base Γ from Example 6.1.1 is due to a different expert, if a knowledge engineer were to request the assessors responsible for the inconsistency to make their opinions compatible, all three experts should be involved, since Γ is an inescapable conflict.

Given that all abrupt consolidations can be viewed as natural consolidations and each knowledge base is a weakening of itself, an inescapable conflict is something weaker than a minimal inconsistent set:

Corollary 6.3.5. *If Δ is a minimal inconsistent set, then Δ is an inescapable conflict.*

Example 6.3.6. Consider again the knowledge base from Example 6.2.7:

$$\Gamma = \{P(x_1 \wedge x_2) \geq 0.6, P(x_1 \wedge \neg x_2) \geq 0.5, P(x_1) \leq 0.8\} \ .$$

As it was already shown, $\{P(x_1 \wedge x_2) \geq 0.6, P(x_1 \wedge \neg x_2) \geq 0.5\}$ is the only minimal inconsistent set of Γ — and, by Corollary 6.3.5, it is an inescapable conflict. Nevertheless, it can be proved that the whole Γ is an inescapable conflict as well.

Suppose, by contradiction, there are $\Delta_1, \ldots, \Delta_k \subsetneq \Gamma$ such that $\bigcup_{i=1}^k \Delta_i = \Gamma$ and, if Δ'_i is a natural consolidation of Δ_i for all $1 \leq i \leq k$ and $\bigcup_{i=1}^k \Delta'_i$ is a weakening of Δ, then $\bigcup_{i=1}^k \Delta'_i$ is a natural consolidation of Δ. To build $\bigcup_{i=1}^k \Delta'_i$, we pick a natural consolidation Δ'_i for each $\Delta_i \subsetneq \Delta$. There are two cases: (a) $P(x_1 \wedge x_2) \geq 0.6 \in \Delta_i$; and (b) $P(x_1 \wedge x_2) \geq 0.6 \notin \Delta_i$. In case (a), we construct Δ'_i by relaxing the probability bound of the conditional $P(x_1 \wedge x_2) \geq 0.6$ to $P(x_1 \wedge x_2) \geq 0.5$; formally, $\Delta'_i = (\Delta_i \setminus \{P(x_1 \wedge x_2) \geq 0.6\}) \cup \{P(x_1 \wedge x_2) \geq 0.5\}$. In case (b), we choose the trivial weakening $\Delta'_i = \Delta_i$. Consider then the following knowledge base:

$$\Gamma' = \bigcup_{i=1}^{k} \Delta'_i = \{P(x_1 \wedge x_2) \geq 0.5, P(x_1 \wedge \neg x_2) \geq 0.5, P(x_1) \leq 0.8\} \ .$$

Note that each $\Delta'_i \subsetneq \Gamma'$ is consistent and thus a natural consolidation of the corresponding Δ_i. By the premises, Γ' should be consistent, but it is inconsistent, since $\Gamma' \setminus \{P(x_1) \leq 0.8\}$ implies $P(x_1) \geq 1$ (as shown in the Example 6.1.1). Finally, there cannot exist such $\Delta_1, \ldots, \Delta_k \subsetneq \Gamma$, and Γ is an inescapable conflict. □

Recall that a free conditional is defined in the standard way as not belonging to any minimal inconsistent set. We prove the analogous result for innocuous conditionals and inescapable conflicts, linking all concepts introduced in this section.

Theorem 6.3.7. *The following statements are equivalent:*

1. *For all maximal consolidations Γ' of Γ, $\alpha \in \Gamma'$.*
2. *If Γ' is a natural consolidation of Γ, then α is consistent with Γ'.*
3. *There is no inescapable conflict Δ in Γ such that $\alpha \in \Delta$.*
4. *α is an innocuous conditional in Γ.*

As one can focus on inescapable conflicts while consolidating a probabilistic base by relaxing the probability bounds, innocuous conditionals can be ignored in this process, which validates their role of characterising when a conditional is not causing the inconsistency.

6.4 Primitive Conflicts and ε-Consolidations

If a base is to be consolidated via a quantitative operator satisfying (Non-Strengthening) — via a weakening —, we saw that it suffices to watch for the inescapable conflicts while relaxing the bounds, and innocuous conditionals can be ignored. This applies to the consolidation operators $\mathcal{C}_p$,

$\mathcal{C}_{CRV}$ and $\mathcal{C}_p^\varepsilon$ presented in Section 4.3. Nevertheless, we argued in Section 4.4 that the consolidation via violations aims at finding probabilistic interpretations π which are not well-characterised by the probability bounds returned by $\mathcal{C}_p^\varepsilon$. We introduced the ε-consolidations to properly encode the intended probabilistic interpretations π resulting from minimising some p-norm of the violations vector. With this extended logic in mind, we can investigate the primitive conflict characterisation related to the ε-consolidation operators $\mathcal{V}_p$.

Consider for instance a base $\Gamma = \Delta \cup \{\alpha\}$ in $\mathbb{K}$, where $\alpha \in \mathcal{L}_{X_n}^P$ is an innocuous conditional in Γ and Δ is its single inescapable conflict. In principle, it is not necessarily the case that an ε-consolidation of Δ is consistent with α. However, we can directly adapt the concepts of inescapable conflict and innocuous conditional to attain the same results for ε-consolidations:

Definition 6.4.1. An *ε-innocuous probabilistic conditional* of Γ is a probabilistic conditional $\alpha \in \Gamma$ such that, for every maximal ε-consolidation Γ' of Γ, $\alpha \in \Gamma'$.

Analogously to the free and the innocuous, the ε-innocuous conditionals has an alternative definition:

Lemma 6.4.2. *Consider a knowledge base $\Gamma \in \mathbb{K}$ and a conditional $\alpha \in \Gamma$. The following statements are equivalent:*

1. *For all maximal ε-consolidations Γ' of Γ, $\alpha \in \Gamma'$.*
2. *If Γ' is an ε-consolidation of Γ, then α is consistent with Γ'.*

From which follows the result analogous to Proposition 6.2.8:

Proposition 6.4.3. *Consider a probabilistic conditional $\alpha \in \Gamma$. If α is safe, it is ε-innocuous; if α is ε-innocuous, it is free.*

Regarding the inescapable conflicts, their translation to ε-consolidations is also straightforward:

Definition 6.4.4. A knowledge base Γ is an *ε-inescapable conflict* if Γ is inconsistent and there are no $\Delta_1, \ldots, \Delta_k \subsetneq \Gamma$, with $k \geq 1$, such that:

1. $\bigcup_{i=1}^k \Delta_i = \Gamma$;
2. If Δ_i' is an ε-consolidation of Δ_i for all $1 \leq i \leq k$ and $\bigcup_{i=1}^k \Delta_i'$ is a ε-weakening of Γ, then $\bigcup_{i=1}^k \Delta_i'$ is a ε-consolidation of Γ.

The alternative definition also holds:

Lemma 6.4.5. *A knowledge base Γ is an ε-inescapable conflict iff there is a ε-weakening Γ' of Γ such that Γ' is a minimal inconsistent set.*

And the following corollary states that ε-inescapable conflicts capture the all the effort to ε-consolidate a base:

Corollary 6.4.6. *Consider two knowledge bases $\Gamma, \Gamma' \in \mathbb{K}$ such that Γ' is a ε-weakening of Γ. If, for every ε-inescapable conflict $\Delta \subseteq \Gamma$, the corresponding ε-weakening $\{\beta \in \Gamma' \mid \alpha \in \Delta$ and $\alpha \leq_\varepsilon \beta\}$ is consistent, then Γ' is an ε-consolidation of Γ.*

Note that all abrupt consolidations can also be viewed as ε-consolidations. We can insert violations $\varepsilon_i = 1$ instead of discarding conditionals, making then trivially satisfied by any probability mass. Hence, as each knowledge base is a ε-weakening of itself, an ε-inescapable conflict is something weaker than a minimal inconsistent set:

Corollary 6.4.7. *If Δ is a minimal inconsistent set, then Δ is an ε-inescapable conflict.*

Finally, we can connect the concepts for ε-consolidations:

Theorem 6.4.8. *The following statements are equivalent:*

1. *For all maximal ε-consolidations Γ' of Γ, $\alpha \in \Gamma'$.*
2. *If Γ' is an ε-consolidation of Γ, then α is consistent with Γ'.*
3. *There is no ε-inescapable conflict Δ in Γ such that $\alpha \in \Delta$.*
4. *α is an ε-innocuous conditional in Γ.*

Regarding the relation between inescapable conflicts and ε-inescapable conflicts, we are not in a position to prove a general link nor to tell them apart by a counter-example. Nonetheless, for unconditional probabilistic bases, these concepts collapse:

Lemma 6.4.9. *For any unconditional probabilistic knowledge base $\Gamma \in \mathbb{K}_c$, Γ is an inescapable conflict iff Γ is an ε-inescapable conflict.*

Corollary 6.4.10. *For any unconditional probabilistic knowledge base $\Gamma \in \mathbb{K}_c$ and conditional $\alpha \in \Gamma$, α is innocuous in Γ iff α is ε-innocuous in Γ.*

6.5 Reconciling the Postulates

Finally, we show how the postulates for measuring inconsistency in probabilistic knowledge bases can be reconciled. Supposing an underlying consolidation procedure via probabilities adjustment, we showed in Section 6.2 how innocuous conditionals can be ignored when one is restoring the consistency of a base. Afterwards, in Section 6.3, we posited that inescapable conflicts — and not MISes — capture the causes of inconsistency in probabilistic bases. Hence, innocuous conditionals, which are outside any inescapable conflict, are a natural candidate to substitute for free conditionals within the formulation of (Independence) in order to obtain a version compatible with (Consistency) and (Continuity). To avoid confusion, we call it (i-Independence) (De Bona and Finger, 2015):

Postulate 6.5.1 (**i-Independence**). *If α is an innocuous conditional of Γ, then $\mathcal{I}(\Gamma) = \mathcal{I}(\Gamma \setminus \{\alpha\})$.*

From Proposition 6.2.8, follows the next relation among (Weak Independence), (i-Independence) and (Independence) (De Bona and Finger, 2015):

Corollary 6.5.2. *If $\mathcal{I}$ satisfies (Independence), then $\mathcal{I}$ satisfies (i-Independence). If $\mathcal{I}$ satisfies (i-Independence), then $\mathcal{I}$ satisfies (Weak Independence).*

In a similar way, we can employ the characterisation of primitive conflict related to innocuous conditional in order to modify (MIS-separability). If Γ is a base in $\mathbb{K}$, Let $\mathrm{IC}(\Gamma)$ denote the set of all inescapable conflicts $\Psi \subseteq \Gamma$.

Property 6.5.3 (**IC-Separability**). *For any* $\Gamma \in \mathbb{K}_c$*, if* $\Gamma = \Delta \cup \Psi$*,* $\Delta \cap \Psi = \varnothing$ *and* $\mathrm{IC}(\Gamma) = \mathrm{IC}(\Delta) \cup \mathrm{IC}(\Psi)$*, then* $\mathcal{I}(\Gamma) = \mathcal{I}(\Delta) + \mathcal{I}(\Psi)$*.*

As inescapable conflict is a weaker concept than MIS, (MIS-separability) is stronger than (IC-separability) (De Bona and Finger, 2015).

Corollary 6.5.4. *If* $\mathcal{I}$ *satisfies (MIS-separability), then* $\mathcal{I}$ *satisfies (IC-separability).*

Similarly to Proposition 3.1.7, which relates (MIS-separabiliby) and (Independence), we can prove:

Corollary 6.5.5. *If* $\mathcal{I}$ *satisfies (IC-separability) and (Consistency), then* $\mathcal{I}$ *satisfies (i-Independence).*

These new desirable properties are indeed compatible with (Consistency) and (Continuity), as we can exhibit inconsistency measures simultaneously satisfying them all:

Lemma 6.5.6. *For any* $p \in \overline{\mathbb{N}}_{>0}$*,* $\mathcal{I}_p$ *satisfies (i-Independence). Furthermore,* $\mathcal{I}_p$ *satisfies (IC-Separability) iff* $p = 1$*.*

Corollary 6.5.7. *There is an inconsistency measure* $\mathcal{I} : \mathbb{K}_c \to [0, \infty)$ *that satisfies (Consistency), (Continuity), (i-Independence) and (IC-Separability).*

Regarding (i-Independence), we were not able to prove it for $\mathcal{I}_p^\varepsilon$ nor to find a counter-example. Despite the fact that an innocuous formula α can be ignored while consolidating a base $\Delta \cup \{\alpha\}$ via $\mathcal{C}_p^\varepsilon$, in the sense that α will be consistent with $\mathcal{C}_p^\varepsilon(\Delta)$, this consolidation operator does not capture the essence of minimal violation measures. It may be the case that probabilistic interpretations not minimising the discrepancy $d_p^\varepsilon(\Gamma)$ satisfies the consolidation $\mathcal{C}_p^\varepsilon(\Gamma)$. To avoid this, we introduced ε-consolidations and the consolidation operator $\mathcal{V}_p$ in Section 4.4, and it was showed in Section 6.4 that ε-inescapable conflicts is the suitable primitive conflict characterisation regarding such operator. By Theorem 6.4.8, a conditional is ε-innocuous if it is not in any ε-inescapable conflicts of the base, what motivates the following postulate:

Postulate 6.5.8 (**ε-Independence**). *If* α *is an* ε*-innocuous conditional of* Γ*, then* $\mathcal{I}(\Gamma) = \mathcal{I}(\Gamma \setminus \{\alpha\})$*.*

From Proposition 6.4.3, it follows the relation among (Weak Independence), (ε-Independence) and (Independence):

Corollary 6.5.9. *If* $\mathcal{I}$ *satisfies (Independence), then* $\mathcal{I}$ *satisfies (*ε*-Independence). If* $\mathcal{I}$ *satisfies (*ε*-Independence), then* $\mathcal{I}$ *satisfies (Weak Independence).*

Regarding (IC-separability), we can modify it using ε-inescapable conflicts. If Γ is a base in $\mathbb{K}$, Let $\mathrm{IC}_\varepsilon(\Gamma)$ denote the set of all ε-inescapable conflicts $\Psi \subseteq \Gamma$.

Property 6.5.10 (**ε-Separability**). *For any* $\Gamma \in \mathbb{K}_c$*, if* $\Gamma = \Delta \cup \Psi$*,* $\Delta \cap \Psi = \varnothing$ *and* $\mathrm{IC}_\varepsilon(\Gamma) = \mathrm{IC}_\varepsilon(\Delta) \cup \mathrm{IC}_\varepsilon(\Psi)$*, then* $\mathcal{I}(\Gamma) = \mathcal{I}(\Delta) + \mathcal{I}(\Psi)$*.*

As ε-inescapable conflict is a weaker concept than MIS, (MIS-separability) is stronger than (ε-Separability).

Corollary 6.5.11. *If* $\mathcal{I}$ *satisfies (MIS-separability), then* $\mathcal{I}$ *satisfies (IC-separability).*

Similarly to Proposition 3.1.7, which relates (MIS-separabiliby) and (Idependence), we can prove:

Corollary 6.5.12. *If $\mathcal{I}$ satisfies (ε-Separability) and (Consistency), then $\mathcal{I}$ satisfies (ε-Independence).*

Using these new postulates, we complete the list of properties held by minimal violation measures:

Lemma 6.5.13. *For any $p \in \mathbb{N}_{>0} \cup \{\infty\}$, $\mathcal{I}_p^\varepsilon$ satisfies (ε-Independence). $\mathcal{I}_p^\varepsilon$ satisfies (ε-Separability) iff $p = 1$.*

The result above implies the compatibility of these new desirable properties with (Consistency) and (Continuity):

Corollary 6.5.14. *There is an inconsistency measure $\mathcal{I} : \mathbb{K}_c \to [0, \infty)$ that satisfies (Consistency), (Continuity), (ε-Independence) and (ε-Separability).*

In the end, we have two relaxed versions of the duo (Independence) and (MIS-Separability), and each new pair is compatible with (Consistency) and (Continuity). While any $\mathcal{I}_p$ satisfies (i-Independence), each $\mathcal{I}_p^\varepsilon$ satisfies (ε-Independence). These postulates are based on characterisations of primitive conflicts derived from different consolidation procedures. That is, imposing one or another version of (Independence) is suggesting the corresponding method to restore consistency. However, one can be interested only in measuring the inconsistency, without consolidating the base. In this case, further criteria are needed in order that a specific inconsistency measure (or a family) may be single out. Computational aspects have already been discussed in Chapter 3, where it is shown that the computation of $\mathcal{I}_1^\varepsilon$ and $\mathcal{I}_\infty^\varepsilon$ can be done through linear programs, being no harder than solving PSAT in practice. In fact, these two measures have further reasons to be endorsed, as we give them meaningful interpretations in the next Chapter.

Chapter 7

Inconsistency Measures via Dutch Books

The question of how to gauge the inconsistency degree of a set of probability assessments is not exclusive to probabilistic logic. In Bayesian statistics and in formal epistemology, where probabilities are understood as subjective degrees of belief, this issue has been tackled with an entirely different approach. By associating an agent's degrees of belief to her betting behaviour, statisticians and philosophers have been measuring the inconsistency of her subjective probabilities via the money loss she would be exposed to. Intuitively, the more inconsistent the probabilities, the greater the amount of money an agent can lose for sure while gambling. We show there is an equivalence between inconsistency measures devised in this approach and measures based on distance minimisation, providing an interpretation for the latter.

In Section 7.1, we introduce the betting concept of Dutch book, which imposes a sure loss to an incoherent agent. In Section 7.2, we present two inconsistency measures based on Dutch books and prove their equivalence to two minimal violation measures. Other measures via sure loss in Dutch books from the literature are reviewed in Section 7.3. Most of the content in these three sections was originally presented in (De Bona and Finger, 2015). Finally, we enhance the syntax with confidence factors in Section 7.4, following Nau (1981), to show a more general equivalence between the two approaches to measuring inconsistency in probabilistic bases.

7.1 Introducing Dutch Books

In formal epistemology and Bayesian statistics, there is an interest in measuring the incoherence[1] of an agent whose beliefs are given as probabilities or lower previsions over propositions or random variables — a *Bayesian* agent. If we have propositions from classical logic, the formalised problem at hand is exactly the one the present study investigates. When the agent's degrees of belief are represented by a probabilistic knowledge base, to measure the agent's incoherence is to measure the inconsistency of such knowledge base. Schervish, Seidenfeld, and Kadane (1998, 2002b) have proposed ways to measure incoherence of an agent based on Dutch books within the Statistics community, and Staffel (2015) analyses them in the context of formal epistemology.

Dutch book arguments are based on the agent's betting behaviour induced by her degrees of belief and are typically used to show their irrationality. To introduce the concept of Dutch book, we start with the unconditional case. Dutch book arguments rely on an operational interpretation

[1] This term is not related to the coherence of probabilities defined in Section 2.2.3, from de Finetti's theory; we are still within the probabilistic semantics we have been employing.

of (imprecise) degrees of belief, in which their lower/upper bounds are defined through bet buying/selling prices. Suppose an agent believes that the probability of proposition φ_i is at least (or at most) q_i, for $1 \leq i \leq m$. Consider a bet ticket on the proposition φ_i that returns a prize from the ticket seller of $\lambda_i \geq 0$ if φ_i is the case; otherwise it is worthless. Dutch book arguments generally use a *willingness-to-bet assumption* that this agent is willing to buy (or to sell, if q_i is an upper probability) such a bet ticket on φ_i for $q_i\lambda_i \geq 0$, for any $\lambda_i \geq 0$. Then, if a bettor can buy and/or sell a set of bet tickets from/to the agent, according to her probabilities, that will cause her a sure loss no matter which possible world is the case, we say she is exposed to a Dutch book. This set of bet tickets that causes a guaranteed loss to the agent is called a Dutch book.

Example 7.1.1. Alice (the agent) and Bob (the bettor) are flying to the beach. To spend the time on the plane, they discuss and gamble on the destination weather, to be checked upon arrival — will it be sunny and/or hot (say, at least $20°C$)? Alice assigns probability bounds for three propositions, formed by the atoms x_1 =“the weather is sunny” and x_2 =“the weather is hot”:

- she believes the probability of the weather being hot is at least 70%; which we represent by the conditional $P(x_1) \geq 0.7$;
- she thinks the probability of the weather being sunny is at least 50%; which is represented by $P(x_2) \geq 0.5$;
- she also says that the probability of the weather being both hot and sunny is at most 10%; formalised into $P(x_1 \wedge x_2) \leq 0.1$.

Now Bob can choose which bet tickets he wants to buy from Alice, and which ones he wants to sell, under the willingness-to-bet assumption. He can also set the prize λ_i each ticket will return if the corresponding proposition turns out to be the case. So Bob decides to trade the following bet tickets with Alice:

- He sells to Alice a bet ticket that pays back \$10 if the weather is hot (x_1 is true); and \$0 otherwise. Alice pays $\$10 \times 0.7 = \$7$ for it, which she considers fair;
- Bob also makes Alice buy for $\$10 \times 0.5 = \$5$ a bet ticket that will return \$10 to her only if the weather is sunny (x_2 is true);
- Finally, Bob buys from Alice a bet ticket whose prize \$10 is paid back only in case the weather is hot and sunny ($x_1 \wedge x_2$ is true), and Alice sells it for $\$10 \times 0.1 = \$1$.

Combining the three bet tickets, Bob received $\$7 + \$5 = \$12$ from Alice and returned \$1 to her, so that before they arrive and prizes are paid Alice is losing \$11. Table 7.1 show the prizes Alice can earn from each bet ticket traded for each possible weather; a negative quantity means Alice has to pay it to Bob.

Note that, no matter how the weather is when they arrive and pay the prizes, the total quantity Alice can receive from Bob is at most \$10. Since she is losing \$11 before landing and checking the weather, she will still be losing at least \$1 in the end. Given this sure loss scenario, this set of three bets is said to be a Dutch book against her. □

The agent buying a bet ticket on φ for $q\lambda$ is equivalent to her selling a bet ticket on $\neg\varphi$ for $(1-q)\lambda$, for in both transactions her net profit would be:

Table 7.1: *Prizes earned by Alice (in \$) for each bet and each possible world.*

possible world	$x_1 \wedge x_2$	$x_1 \wedge \neg x_2$	$\neg x_1 \wedge x_2$	$\neg x_1 \wedge \neg x_2$
bet on Hot (x_1)	10	10	0	0
bet on Sunny (x_2)	10	0	10	0
bet on Hot and Sunny ($x_1 \wedge x_2$)	-10	0	0	0
Total	10	10	10	0

- $(1-q)\lambda$ if φ happens to be the case;
- $-q\lambda$ if $\neg\varphi$ happens to be the case.

Since the belief in $P(\varphi) \leq q$ can be expressed via $P(\neg\varphi) \geq 1-q$, every bet ticket transaction has an equivalent one with the agent buying, corresponding to a lower probability bound. Hence, any Dutch book can be formally expressed with the agent only buying bet tickets. For instance, we can equivalently rewrite the example above:

Example 7.1.2. Consider Example 7.1.1, with the same beliefs for Alice, just replacing the transaction

- Bob buys from Alice a bet ticket whose prize \$10 is paid back only in case the weather is hot and sunny ($x_1 \wedge x_2$ is true), and Alice sells it for $\$10 \times 0.1 = \$1$.

By the equivalent one:

- Alice buys from Bob a bet ticket whose prize \$10 is paid back only in case the weather is not hot or is not sunny ($\neg x_1 \vee \neg x_2$ is true, $x_1 \wedge x_2$ is false), and Bob sells it for $\$10 \times 0.9 = \$9$.

In that case, Alice would to buying bet tickets on x_1, x_2 and $\neg x_1 \vee \neg x_2$ for $\$10\times0.7$, $\$10\times0.5$ and $\$10\times0.9$, respectively. Before they arrive and prizes are paid Alice would be losing $\$7+\$5+\$9 = \$21$. Table 7.2 show the prizes Alice could earn from each bet ticket traded for each possible weather.

Table 7.2: *Prizes earned by Alice (in \$) for each bet and each possible world.*

possible world	$x_1 \wedge x_2$	$x_1 \wedge \neg x_2$	$\neg x_1 \wedge x_2$	$\neg x_1 \wedge \neg x_2$
bet on Hot (x_1)	10	10	0	0
bet on Sunny (x_2)	10	0	10	0
bet on not Hot or not Sunny ($\neg x_1 \vee \neg x_2$)	0	10	10	10
Total	20	20	20	10

That is, Alice could earn at most 20\$ from Bob after checking the actual weather. Equivalently to Example 7.1.1, buying this triple of bet tickets would yield a net guaranteed loss of 1\$ to Alice. □

From here on we consider that the agent only buys bet tickets from the bettor. Instead of the agent paying for a bet ticket and eventually getting its prize back from the bettor, we can view this whole operation as a single contract between these two players.

Definition 7.1.3. A *gamble* on $\varphi \in \mathcal{L}_{X_n}$ is an agreement between the agent and the bettor with two parameters, the *stake* $\lambda \geq 0$ and the *relative price* $q \in [0,1]$, stating that:

- the agent pays $\lambda \times q$ to the bettor if φ is false;

- the bettor pays $\lambda \times (1 - q)$ to the agent if φ is true.

A gamble on φ with stake $\lambda \geq 0$ and relative price q is equivalent to the agent buying from the bettor a bet ticket for $\lambda \times q$ that returns λ only if φ is the case. The willingness-to-bet assumption translates to gambles in the following way: if an agent believes that the probability of a proposition φ being true is at least q, she finds acceptable gambles on φ with any stake $\lambda \geq 0$ and relative price q. In Example 7.1.2, the tickets trading is equivalent to a set of three gambles: a gamble on x_1 with stake \$10 and relative price 0.7; a gamble on x_2 with stake \$10 and relative price 0.5; and a gamble on $\neg x_1 \vee \neg x_2$ with stake \$10 and relative price 0.9.

A gamble on φ can be generalised to consider a conditioning event ψ. Consider a bet ticket that, when ψ is true, pays a prize of λ if φ is the case and returns 0 if φ is false. In other words, this bet ticket works as a regular gamble on φ when ψ is the case. However, suppose this bet ticket pays back to the agent the same amount that was spent in its buying if ψ is false — that is, the gamble is canceled. The following generalisation of gambles capture these "conditional bets":

Definition 7.1.4. A *(conditional) gamble* on $\varphi \in \mathcal{L}_{X_n}$ given $\psi \in \mathcal{L}_{X_n}$ (or a gamble on $\varphi|\psi$) is an agreement between the agent and the bettor with two parameters, the *stake* $\lambda \geq 0$ and the *relative price* $q \in [0, 1]$, stating that:

- the agent pays $\lambda \times q$ to the bettor if ψ is true and φ is false;
- the bettor pays $\lambda \times (1 - q)$ to the agent if ψ is true and φ is true;
- the gamble is called off, causing neither profit nor loss to the involved parts, if ψ is false.

Accordingly, we generalise the willingness-to-bet assumption: if an agent believes that the probability of a proposition φ being true given that ψ is true is at least q, she finds acceptable gambles on $\varphi|\psi$ with stake $\lambda \geq 0$ and relative price q. A *Dutch book* is a set of (conditional) gambles that the agent sees as fair, under the willingness-to-bet assumption, that causes her a guaranteed loss no matter which possible world is the case. We assume Dutch books contain exactly one gamble on $(\varphi_i|\psi_i)$ with stake $\lambda_i \geq 0$ and relative price q_i per each conditional $P(\varphi_i|\psi_i) \geq q_i \in \Gamma$, the base formalising the agent's beliefs. This is not restrictive, since gambles on the same $(\varphi_i|\psi_i)$ with the same relative price[2] can be merged by summing the stakes, and the absence of a gamble is equivalent to a stake equal to zero. We can thus denote a Dutch book simply by the value of its stakes $\lambda_1, \ldots, \lambda_m \geq 0$, where $m = |\Gamma|$. Actually, any set of gambles involving an agent whose epistemic state is represented by Γ can be represented by these m values of stakes, since the relative prices are set in Γ.

If the set of probabilistic conditionals that represents an agent's epistemic state turns out to be inconsistent, she is said to be *incoherent*. An incoherent agent is always exposed to a Dutch book, and vice-versa (Nau, 1981). In other words, an agent sees as fair a set of gambles that causes her a guaranteed loss if, and only if, the knowledge base codifying her (conditional) degrees of belief is inconsistent. We can check this connection in Example 7.1.1: $P(x_1) \geq 0.7$ and $P(x_2) \geq 0.5$ imply a probability of at least 0.2 for $x_1 \wedge x_2$, which Alice violates. Consequently, she is exposed to a Dutch book, for her three probability bound assessments are not jointly satisfiable. In this way,

[2]Since we focus on canonical bases, there cannot be two conditionals one the same $\varphi_i|\psi_i$ and different lower bounds, so that different gambles on $\varphi_i|\psi_i$ shall have the same relative price.

Dutch book arguments were introduced to show that a set of rational degrees of belief must obey the axioms of probability and are a standard proof of incoherence (introductions to Dutch books and their relation to incoherence can be found in (Shimony, 1955) and (de Finetti, 1974)).

7.2 Minimal Violation Measures as Guaranteed Sure Losses

Since Dutch books are the footprint of inconsistent degrees of belief, a natural approach to measuring an agent's incoherence is through the magnitude of the sure loss she is vulnerable to. The intuition says that the more incoherent an agent is, the greater the guaranteed loss that can be imposed on her through a Dutch book. Nonetheless, with no bounds on the stakes, such loss would also be unlimited for incoherent agents. For instance, in Example 7.1.2, if stakes were \$100, \$100 and \$100, Alice would have a net loss of at least \$10, instead of \$1, regardless of the weather on arrival. To better understand this relation between stakes and Dutch book losses, we formalise the gambling setting in the following.

Consider the knowledge base $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq m\}$ representing an agent's epistemic state. Let $\lambda_i \geq 0$ denote a gamble on $(\varphi_i|\psi_i)$ with relative price q_i, for $1 \leq i \leq m$. A set of gambles can then be represented by the vector $\langle \lambda_1, \ldots, \lambda_m \rangle$. If a possible world w_j is the case, the net profit for the agent regarding a bet on $\varphi|\psi$ with stake λ and relative price q can be computed via

$$\lambda(I_{w_j}(\varphi \wedge \psi) - qI_{w_j}(\psi)),$$

in which $I_{w_j} : \mathcal{L}_{X_n} \to \{0, 1\}$ is the indicator function of the set $\{\varphi \in \mathcal{L}_{X_n} | w_j \models \varphi\}$ — a valuation. For a gamble on $(\varphi_i|\psi_i)$ with stake λ_i, the agent's net profit in a possible word w_j is $\lambda_i(I_{w_j}(\varphi_i \wedge \psi_i) - q_i I_{w_j}(\psi_i))$. Recall (from (2.4)–(2.5)) that $a_{ij} = I_{w_j}(\varphi_i \wedge \psi_i) - q_i I_{w_j}(\psi_i)$. If a given possible world w_j is the case, the set of gambles $\langle \lambda_1, \ldots, \lambda_m \rangle$ gives the agent a profit of $\sum_{i=1}^{m} a_{ij}\lambda_i$. Let ℓ be a sure loss ($-\ell$ is profit) a set of gambles yields to the agent; *i.e.*, no matter which possible world is the case, the agent loses at least ℓ. Thus, ℓ is such that $\sum_{i=1}^{m} a_{ij}\lambda_i \leq -\ell$ for all possible worlds w_j. To measure the agent's incoherence as loss exposure, one needs to find the maximum amount of money she can lose for sure through a Dutch book (a set of bets). When there is no restriction on the stakes, to find a set of gambles $\langle \lambda_1, \ldots, \lambda_m \rangle$ that maximises the sure loss is to solve the following linear program:

$$\max \ell \text{ subject to:} \tag{7.1}$$

$$\begin{bmatrix} 1 & a_{11} & \ldots & a_{m1} \\ 1 & a_{12} & \ldots & a_{m2} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & a_{12^n} & \ldots & a_{m2^n} \end{bmatrix} \begin{bmatrix} \ell \\ \lambda_1 \\ \vdots \\ \lambda_m \end{bmatrix} \leq \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} \tag{7.2}$$

$$\lambda_1, \ldots, \lambda_m, \geq 0 \ . \tag{7.3}$$

The linear program above can be viewed as the dual of that in lines (2.4)–(2.6), which checks the consistency of Γ, if we consider that 0 is the function being minimised in the latter, since we are interested only in its feasibility (for duality theory in linear programing, see, for instance, (Vanderbei , 1996)). Note that, in (2.4)–(2.6), $\sum \pi = 1$ can be inserted into A as a row of 1's. By duality theory, as the program above is feasible, it is unbounded iff (2.4)–(2.6) is infeasible. Consequently, if Γ is

inconsistent, sure loss via Dutch book is unlimited.

Different strategies to circumvent this in order to measure incoherence as a finite loss are found in the literature on the foundations of Bayesian statistics. Schervish *et al.* (2002b) propose a flexible formal approach to limiting these stakes generating a family of incoherence measures for upper and lower previsions on bounded random variables. In this section, we are interested in two of them, which we simplify to our case.

Their whole family of incoherence measures is based on the maximum guaranteed loss an agent is exposed to via a Dutch book, varying only on how stakes are limited. The first incoherence measure Schervish *et al.* introduce that concerns us is when the sum of the values of the stakes is lesser than or equal to one, $\sum_i \lambda_i \leq 1$. The second incoherence measure we investigate is defined as the maximum guaranteed loss when each stake has value no greater than one, or $\lambda_i \leq 1$ [3]. We define the inconsistency measures $\mathcal{I}^{sum}_{SSK} : \mathbb{K}_c \to [0, \infty)$ and $\mathcal{I}^{max}_{SSK} : \mathbb{K}_c \to [0, \infty)$ on canonical knowledge bases as these two incoherence measures on the corresponding agents represented by these knowledge bases. That is, we equate $\mathcal{I}^{sum}_{SSK}(\Gamma)$ (and $\mathcal{I}^{max}_{SSK}(\Gamma)$), for any $\Gamma \in \mathbb{K}_c$, to the maximum sure loss an agent whose epistemic state is represented by Γ is exposed to through a Dutch book when the sum (maximum) of the stakes' values is at most one.

Example 7.2.1. Recall Example 7.1.2, in which there are three gambles, with stakes \$10, \$10 and \$10. These gambles guarantee a loss of at least \$1 to Alice. But now suppose that Bob, while choosing the gambles, must do it so that the absolute values of the stakes sum up to one. He could so arrange the same gambles but changing the stakes to 1/3, 1/3 and 1/3. In this new scenario, Alice would have a sure loss of 1/30. Similarly, if the absolute value of each stake is limited to the interval $[0, 1]$, stakes could be 1, 1 and 1, yielding a guaranteed loss of 1/10 to Alice. In fact, it can be checked (by solving the linear programs) that 1/30 and 1/10 are the greatest amounts one can take for sure from Alice via Dutch book if stakes have absolute values summing up to one or are all in $[0, 1]$, respectively. Formalizing, with $\Gamma = \{P(x_1) \geq 0.7, P(x_2) \geq 0.5, P(x_1 \wedge x_2) \leq 0.1\}$ codifying Alice's epistemic state, we have $\mathcal{I}^{sum}_{SSK}(\Gamma) = 1/30$ and $\mathcal{I}^{max}_{SSK}(\Gamma) = 1/10$. □

Even though incoherence measures based on Dutch books, from the Bayesian statistics and the formal epistemology communities, and inconsistency measures based on distance minimisation, from Artificial Intelligence researchers, may seem unrelated at first, they are actually two sides of the same coin. The programs that compute the maximum guaranteed loss an agent is exposed to are technically dual to those that minimise distances to measure inconsistency. Nau (1981) has already investigated this matter, mentioning results (discussed in Section 7.4) similar to the following:

Theorem 7.2.2. *For any* $\Gamma \in \mathbb{K}_c$, $\mathcal{I}^{sum}_{SSK}(\Gamma) = \mathcal{I}^{\varepsilon}_{\infty}(\Gamma)$.

Proof. Just add the constraint $\lambda_1 + \cdots + \lambda_m \leq 1$ to the linear program (7.1)–(7.3). The dual of this new program would become the program (3.19)–(3.21), which computes $\mathcal{I}^{\varepsilon}_{\infty}(\Gamma)$. So, by the strong duality theorem, $\mathcal{I}^{sum}_{SSK}(\Gamma) = \mathcal{I}^{\varepsilon}_{\infty}(\Gamma)$, for both programs are always feasible. □

Recall that $\mathcal{I}^{\varepsilon}_{\infty}$ is exactly one of the two feasible measures proposed by Potyka (2014). Far from meaningless, such measure quantifies the maximum sure loss an agent is exposed to when the sum of the stakes is no greater than one — or, equivalently, fixed at one.

[3] Schervish *et al.* (2002b) actually measure the incoherence as maximum rates between the guaranteed loss and the sum (the maximum) of the stakes' values. Clearly, this is equivalent to maximizing the guaranteed loss when the sum of the stakes' values is no greater than 1 (or these values are in $[0, 1]$)

As to Potyka's other feasible proposal, $\mathcal{I}_1^\varepsilon$, duality in linear programming provides a correspondence with the second incoherence measure we presented from Schervish *et al.*:

Theorem 7.2.3. *For any* $\Gamma \in \mathbb{K}_c$, $\mathcal{I}_{SSK}^{max}(\Gamma) = \mathcal{I}_1^\varepsilon(\Gamma)$.

Proof. Similarly to the proof of Theorem 7.2.2, insert the constraints $\lambda_i \leq 1$, for $1 \leq i \leq m$, for $m = |\Gamma|$, into the linear program (7.1)–(7.3). The dual of this new program would become the program (3.14)–(3.17), with $p = 1$, which computes $\mathcal{I}_1^\varepsilon(\Gamma)$. Again, by the strong duality theorem, $\mathcal{I}_{SSK}^{max}(\Gamma) = \mathcal{I}_1^\varepsilon(\Gamma)$, since both programs are always feasible. □

Theorem 7.2.3 states the extensional identity between $\mathcal{I}_1^\varepsilon$ and $\mathcal{I}_{SSK}^{max}$. Within the unconditional probabilities scenario, this means that the Manhattan distance between the agent's probabilities and the closest consistent probabilities is equal to the maximum sure loss she is exposed to when stakes' values are not higher than one.

Theorem 7.2.2 and Theorem 7.2.3 give an operational interpretation for the inconsistency measures $\mathcal{I}_\infty^\varepsilon$ and $\mathcal{I}_1^\varepsilon$ based on betting behaviour. It was remarked in Section 3.2.2 that $\mathcal{I}_p^\varepsilon$ and $\mathcal{I}_p$ give the same inconsistency degrees to unconditional knowledge bases. Thus, Dutch books with limited stakes ($\lambda_i \leq 1$ or $\sum_i \lambda_i \leq 1$) can be used to rationalise also $\mathcal{I}_1$ and $\mathcal{I}_\infty$ in the unconditional setting. However, when we take into account conditional probabilities, only $\mathcal{I}_1^\varepsilon$ and $\mathcal{I}_\infty^\varepsilon$ measure the maximum guaranteed loss an agent would be exposed to, when stakes are limited via $\lambda_i, \leq 1$ or $\sum_i \lambda_i \leq 1$, respectively.

Different ways of bounding stakes can lead to different inconsistency measures, but our motivation in this section was not to use Dutch books to determine which measures should be adopted — that is the reason of the postulates. The point here is that these two measures ($\mathcal{I}_1^\varepsilon$ and $\mathcal{I}_\infty^\varepsilon$), besides satisfying some postulates and being computable through linear programs, have a meaningful interpretation. In the next section, we show that other measures based on Dutch books have these qualities as well.

7.3 Other Feasible Principled Measures

In order to measure incoherence as the greatest guaranteed loss in a Dutch book, Schervish *et al.* (1998) have firstly proposed two different ways of normalizing such loss: by limiting either the agent's or the bettor's resources. The authors introduce the concept of *escrow* as the amount committed into a gamble by the agent (or the bettor). For instance, consider a gamble on $\varphi_i|\psi_i$ with stake $\lambda_i \geq 0$ and relative price q_i. The agent might lose $q_i\lambda_i$ with this gamble, while the bettor is exposed to a loss of $(1 - q_i)\lambda_i$. Schervish *et al.* call these quantities the agent's and the bettor's *escrows*. Equivalently, the agent's (or bettor's) escrow for a gamble is how much she (he) has to commit from her (his) resources to cover an eventual loss.

Instead of bounding the sum of the stakes, an agent's degree of incoherence can be measured, as the maximum guaranteed loss in a Dutch book, by limiting the agent's (or the bettor's) total escrow to one[4]. In other words, we are limiting how much the agent (or the bettor) could lose in case that every gamble resolves unfavourably, inflicting a loss to her (him). Schervish *et al.* (1998) give market

[4]Again, this is equivalent to measuring incoherence as the maximum ratio between the sure loss and the agent's (the bettor's) total escrow.

situations that justify these choices. We denote by $\mathcal{I}_{SSK}^{a,sum} : \mathbb{K} \to [0, \infty)$ and $\mathcal{I}_{SSK}^{b,sum}\mathbb{K} \to [0, \infty) \cup \{\infty\}$[5] the inconsistency measures corresponding to these two incoherence measures, when the agent's or the bettor's total escrow is at most one, respectively.

Formally, starting with the linear program of lines (7.1)–(7.3), $\mathcal{I}_{SSK}^{a,sum}(\Gamma)$ and $\mathcal{I}_{SSK}^{b,sum}(\Gamma)$ are obtained via the maximisation of ℓ by adding further constraints. Let $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq m\}$ be a knowledge base. To compute $\mathcal{I}_{SSK}^{a,sum}(\Gamma)$, one needs to insert the restriction $\sum_{i=1}^{m} q_i\lambda_i \leq 1$ into (7.1)–(7.3). Similarly, $\mathcal{I}_{SSK}^{b,sum}(\Gamma)$ is the solution (on ℓ) of the program (7.1)–(7.3) incremented with the constraint $\sum_{i=1}^{m}(1 - q_i)\lambda_i \leq 1$.

The fact that $\mathcal{I}_{SSK}^{b,sum}$ may be unbounded is acknowledged by Schervish *et al.* (1998). For instance, consider an agent whose belief state is given by $\Gamma = \{P(\varphi) \geq 1, P(\neg\varphi) \geq 1\}$. The agent finds acceptable pairs of gambles on (φ and $\neg\varphi$) in which the bettor has escrows equal to zero ($\lambda_i(1-q_i) = 0$, for $q_i = 1$), and sure loss can be scaled arbitrarily up. In such cases, we define $\mathcal{I}_{SSK}^{b,sum}(\Gamma) = \infty$.

Example 7.3.1. Recall Example 7.1.2, its three gambles, with stakes \$10, \$10 and \$10, and the implied loss of at least \$1 to Alice. But now suppose that Bob has to choose gambles in such a way that his (or Alice's) total escrow sum up to 1. Note that, with stakes \$10, \$10 and \$10, his total escrow is $\$10\times(1-0.7)+\$10\times(1-0.5)+\$10\times0.1 = \$9$ (Alice's is $\$10\times0.7+\$10\times0.5+\$10\times(1-0.1) = \$21$). He could then arrange the same gambles but changing the stakes to 10/9, 10/9 and 10/9 (or 10/21, 10/21 and 10/21) in order to his (Alice's) total escrow be equal to one. In this new scenario, Alice would have a sure loss of 1/9 (or 1/21). Once again, one could verify, by solving the linear programs, that 1/9 and 1/21 are the greatest amount one can take for sure from Alice via Dutch book if Bob's or Alice's total escrow is no greater than 1, respectively. Formalizing, with $\Gamma = \{P(x_1) \geq 0.7, P(x_2) \geq 0.5, P(x_1 \wedge x_2) \geq 0.1\}$ codifying Alice's epistemic state, we have $\mathcal{I}_{SSK}^{b,sum}(\Gamma) = 1/9$ and $\mathcal{I}_{SSK}^{a,sum}(\Gamma) = 1/21$. □

Schervish *et al.* (2002b) contemplate in detail a whole spectrum of ways to bound the agent's escrows, the bettor's, or their sum in order to measure the maximum sure loss. For each of these three quantities, the author note that the two extreme functions in their framework used to normalize the guaranteed loss are the maximum and the sum, from which six different inconsistency measures arise Schervish *et al.* (2003). Note that the sum of the agent's and the bettor's escrows for a single gamble is equal to the value of its stake, so $\mathcal{I}_{SSK}^{sum}$ and $\mathcal{I}_{SSK}^{max}$ are two inconsistency measures from this same framework. To build the remaining two measures, escrows could be bounded via their maximum, instead of their total. Intuitively, this corresponds to limiting the quantity the agent (or the bettor) accepts to eventually lose in each individual gamble.

We define the inconsistency measure $\mathcal{I}_{SSK}^{a,max} : \mathbb{K} \to [0, \infty)$ (and $\mathcal{I}_{SSK}^{b,max}\mathbb{K} \to [0, \infty) \cup \{\infty\}$) on knowledge bases as the degree of incoherence of the corresponding agents measured via the maximum sure loss she is exposed through a Dutch book if the agent's (the bettor's) escrow for each gamble in no greater than one. To compute $\mathcal{I}_{SSK}^{a,max}(\Gamma)$, for $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq m\}$, we may again use the linear program (7.1)–(7.3) and compute the maximum value of ℓ with extra constraints $q_i\lambda_i \leq 1$, for $1 \leq i \leq m$. Similarly, $\mathcal{I}_{SSK}^{b,max}(\Gamma)$ is the solution (on ℓ) to the program formed by inserting the restrictions $(1 - q_i)\lambda_i \leq 1$, for $1 \leq i \leq m$, into (7.1)–(7.3). As with $\mathcal{I}_{SSK}^{b,sum}$, we define $\mathcal{I}_{SSK}^{b,max}(\Gamma) = \infty$ when such program is unbounded.

[5] For reasons that will be clear soon, we relax in this section the definition of inconsistency measures, allowing their range to include ∞.

Example 7.3.2. Remember the scenario from Example 7.1.2, in which three gambles are considered, with stakes \$10, \$10 and \$10, and Alice has a guaranteed loss of at least \$1. Now suppose that Bob can only choose gambles in which his (Alice's) eventual loss — the escrow —is lesser than or equal to one. In other words, his (her) maximum escrow is no greater than 1. With these constraints, Bob can choose the same three gambles, but with stakes 2, 2 and 2: his escrows are $2 \times (1 - 0.7) = 0.6$, $2 \times (1 - 0.5) = 1$ and $2 \times 0.1 = 0.2$ (with stakes 10/9, 10/9 and 10/9, Alice's escrow are $(10/9) \times 0.7 = 7/9$, $10/9 \times 0.5 = 5/9$ and $10/9 \times (1 - 0.1) = 1$). Note that Bob (Alice) can eventually lose at most 1 in a single gamble. In this new setting, Alice would have a guaranteed loss of 1/5 (or 1/9). By solving the corresponding linear programs, we would find that 1/5 and 1/9 are the greatest amount one can take for sure from Alice via Dutch book if Bob's or Alice's maximum escrow is no greater than 1, respectively. Formalizing, with $\Gamma = \{P(x_1) \geq 0.7, P(x_2) \geq 0.5, P(x_1 \wedge x_2) \leq 0.1\}$ codifying Alice's epistemic state, we have $\mathcal{I}^{b,max}_{SSK}(\Gamma) = 1/5$ and $\mathcal{I}^{a,max}_{SSK}(\Gamma) = 1/9$. □

These four inconsistency measures ($\mathcal{I}^{a,sum}_{SSK}$, $\mathcal{I}^{b,sum}_{SSK}$, $\mathcal{I}^{a,max}_{SSK}$ and $\mathcal{I}^{b,max}_{SSK}$) based on limiting the escrows have most of the desirable properties we presented.

Lemma 7.3.3. *$\mathcal{I}^{a,sum}_{SSK}$, $\mathcal{I}^{b,sum}_{SSK}$, $\mathcal{I}^{a,max}_{SSK}$ and $\mathcal{I}^{b,max}_{SSK}$ are well-defined and satisfy (Consistency), (ε-Independence) and (Monotonicity). $\mathcal{I}^{a,max}_{SSK}$ and $\mathcal{I}^{b,max}_{SSK}$ also satisfy (Super-additivity) and (ε-separability).*

Lemma 7.3.4. *$\mathcal{I}^{a,sum}_{SSK}$, $\mathcal{I}^{a,max}_{SSK}$, $\mathcal{I}^{b,sum}_{SSK}$ and $\mathcal{I}^{b,max}_{SSK}$ are continuous for probabilities within $(0, 1)$.*

Lemma 7.3.5. *$\mathcal{I}^{a,sum}_{SSK}$ satisfies (Normalisation).*

Altogether, $\mathcal{I}^{a,sum}_{SSK}$, $\mathcal{I}^{b,sum}_{SSK}$, $\mathcal{I}^{a,max}_{SSK}$ and $\mathcal{I}^{b,max}_{SSK}$ are all computable through linear programs, have the core desirable properties, including a version of (Independence), and can be given an operational interpretation. $\mathcal{I}^{a,max}_{SSK}$ and $\mathcal{I}^{b,max}_{SSK}$ also satisfy (Super-additivity) and (ε-separability), while $\mathcal{I}^{a,sum}_{SSK}$ is normalised. These measures can be good alternatives to measure inconsistency in some contexts, as the market scenarios described by Schervish *et al.* (1998).

It is worth noting that $1 - \mathcal{I}^{a,sum}_{SSK}$, $1 - \mathcal{I}^{sum}_{SSK}$ and $1 - \mathcal{I}_{\infty}$ can be seen as possible generalisations of the measure of consistency proposed by Knight (2002) for classical propositional logic, seen in Section 3.1. For a given a set $\Gamma \subsetneq \mathcal{L}_{X_n}$, recall that it is maximally $\eta^{\star}$-consistent if $\eta^{\star}$ is the maximum value of η for which the knowledge base $\Gamma_{\eta} = \{P(\varphi) \geq \eta \,|\, \varphi \in \Gamma\}$ is satisfiable. If we assign probability one to each element of Γ, building the base $\Gamma' = \{P(\varphi) \geq 1 \,|\, \varphi \in \Gamma\}$, then Γ is maximally η-consistent iff $\mathcal{I}^{a,sum}_{SSK}(\Gamma') = \mathcal{I}^{sum}_{SSK}(\Gamma') = \mathcal{I}^{\varepsilon}_{\infty}(\Gamma') = \mathcal{I}_{\infty}(\Gamma') = 1 - \eta$. Note that $\mathcal{I}^{a,sum}_{SSK}(\Gamma)$ and $\mathcal{I}^{sum}_{SSK}(\Gamma)$ are equal for all probabilities in Γ' are 1. Hence, Theorems 7.2.2 and 7.4.7 can rationally support the use of Knight's measure in the classical setting as well. Suppose an agent's belief base Γ contains φ iff she sees as fair a gamble on φ with $q = 1$, then Γ is maximally η-consistent iff the agent is exposed to a maximum sure loss of $1 - \eta$ when her resources are fixed at one (equivalently, stakes sum up to one).

7.4 Measuring Inconsistency with Confidence Factors

The measures $\mathcal{I}^{a,sum}_{SSK}$, $\mathcal{I}^{b,sum}_{SSK}$, $\mathcal{I}^{a,max}_{SSK}$ and $\mathcal{I}^{b,max}_{SSK}$ can also be related to Potyka's minimal violation measures, through a generalisation of the language to include confidence weights associated to the probability bounds. Such weights are proposed by Nau (1981), and find application in merging

probabilities from different experts, for instance. Let a *weighted conditional* be formed by attaching a *confidence factor* δ_i, with $\delta_i \in \mathbb{R}_{>0} \cup \{\infty\}$, to a conditional $P(\varphi_i|\psi_i) \geq q_i$, forming $P(\varphi_i|\psi_i) \geq q_i\langle\delta_i\rangle$, which reads "the probability of φ_i being true given that ψ_i is true is at least q_i with confidence δ_i". A common interpretation for these confidence factors is through how much the agent is willing to bet (the stake or the escrow) on that relative price. We define a *weighted probabilistic knowledge base* as a finite set of weighted conditionals $P(\varphi_i|\psi_i) \geq q_i\langle\delta_i\rangle$. We denote by $\mathbb{K}^w$ the set of all weighted knowledge bases. We extend the notion of canonical to weighted bases with no repeated pair $\varphi_i|\psi_i$, with $\mathbb{K}^w_c$ denoting the set of all such bases. If confidence factors are omitted, we define them as $\delta_i = 1$, so that, in some sense, $\mathbb{K} \subseteq \mathbb{K}^w$.

The consistency of weighted bases is given by ignoring the confidence factors: a probabilistic interpretation satisfies a weighted conditional $P(\varphi_i|\psi_i) \geq q_i\langle\delta_i\rangle$ iff it satisfies the (non-weighted) conditional $P(\varphi_i|\psi_i) \geq q_i$. The concepts of free, safe, MIS, (ε-)innocuous and (ε-)inescapable conflicts apply to weighted conditionals and bases, just ignoring the confidence factors.

As a weighted base is consistent iff the corresponding non-weighted base is consistent, inconsistency measures for the latter can be employed to measure the inconsistency of the former. Nonetheless, weighted bases convey more information, which can be employed to a more fine-grained measurement. For instance, consider the bases:

$$\begin{aligned}\Gamma &= \{P(x_1) \geq 0.7\langle 0.1\rangle, P(\neg x_1) \geq 0.4\langle 0.1\rangle, P(x_2) \geq 0.3\langle 0.9\rangle\}\,;\\ \Psi &= \{P(x_1) \geq 0.7\langle 0.5\rangle, P(\neg x_1) \geq 0.4\langle 0.5\rangle, P(x_2) \geq 0.3\langle 0.1\rangle\}\,.\end{aligned}$$

Both weighted bases are inconsistent, containing the same MIS and (ε-)inescapable conflict, $\{P(x_1) \geq 0.7, P(\neg x_1) \geq 0.4\}$, and the safe conditional $P(x_2) \geq 0.3$. However, in Ψ, the probability bounds in the MIS have higher confidence factors. It means that an agent represented by Ψ is more confident about these inconsistent lower bounds than it is the agent whose epistemic state is encoded in Γ. Therefore, one expects that Ψ be more inconsistent than Γ.

Potyka's minimal violation measures $\mathcal{I}^{\varepsilon}_p : \mathbb{K} \to [0, \infty)$ can be generalised to measure inconsistency in weighted knowledge bases. The intuition says the more confident an agent is about a bound, the higher the penalty for changing it while consolidating. Consider a weighted base $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i\langle\delta_i\rangle | 1 \leq i \leq m\}$ and a probabilistic interpretation π. We adapt the discrepancy d^{ε}_p introduced in Section 4.3.2 to allow for these factors:

$$\begin{aligned}d^{\delta}_p(\Gamma, \pi) &= \|\langle\delta_1\varepsilon_1, \ldots, \delta_m\varepsilon_m\rangle\|_p, \quad \text{where}\\ \varepsilon_i &= \max\{0, q_i P_\pi(\psi_i) - P_\pi(\varphi_i \wedge \psi_i)\} \text{ for all } 1 \leq i \leq m\,.\end{aligned}$$

In such definition, when $\delta_i = \infty$, we assume: if $\varepsilon_i = 0$, then $\delta_i\varepsilon_i = 0$; else, $\delta_i\varepsilon_i = \infty$ and $\|\langle\delta_1\varepsilon_1, \ldots, \delta_m\varepsilon_m\rangle\|_p = \infty$.

For any $p \in \overline{\mathbb{N}}_{>0}$, the measure $\mathcal{I}^{\delta}_p : \mathbb{K}^w_c \to [0, \infty) \cup \{\infty\}$[6] is defined as following:

$$\mathcal{I}^{\delta}_p(\Gamma) = \min\{d^{\delta}_p(\Gamma, \pi)\}\,.$$

The measures $\mathcal{I}^{\delta}_1(\Gamma)$ and $\mathcal{I}^{\delta}_\infty(\Gamma)$ had already been implicitly introduced by Nau (1981), as will

[6] We relax again in this section the definition of an inconsistency measure, allowing them to return infinite.

be later detailed.

The family of inconsistency measures $\mathcal{I}_p^\delta(\Gamma)$ also generalises minimal violation measures with integrity constraints, introduced by Potyka and Thimm (2015). The authors extend minimal violation measures $\mathcal{I}_p^\varepsilon$ to cope with conditionals that need to be satisfied, and cannot be violated. Their formalism can be thought of as representing knowledge with pairs $\langle \Gamma, \Psi \rangle \in \mathbb{K}_c \times \mathbb{K}_c$, where Ψ represents the set of integrity constraints. The idea is that, while measuring the inconsistency of $\Gamma \cup \Psi$ [7] via minimising the discrepancy $d_p^\varepsilon(\Gamma \cup \Psi, \pi)$, only probabilistic interpretations satisfying Ψ — which is assumed consistent — are considered. This is equivalent to minimising $d_p^\varepsilon(\Gamma, \pi)$ with π satisfying Ψ, since any conditional in Ψ would have null violation. This situation can be formalised as an inconsistency measurement for Γ parametrised by a consistent set $\Psi \in \mathbb{K}_c$ of integrity constraints. Hence, supposing a consistent base $\Psi \in \mathbb{K}_c$, the authors introduce the inconsistency measure $\mathcal{I}_p^\Psi : \mathbb{K} \to [0, \infty)$, defined, for all $\Gamma \in \mathbb{K}_c$, as:

$$\mathcal{I}_p^\Psi(\Gamma) = \min\{d_p^\varepsilon(\Gamma, \pi) \mid \pi \in \Pi_n \text{ satisfies } \Psi\}\,.$$

Potyka and Thimm (2015) are interested in generalised models for inconsistent bases. These models are the probabilistic interpretations that minimise the discrepancy d_p^ε. Allowing integrity constraints enables one to assure the satisfaction of a "core" of inviolable conditionals, which can represent information taken for granted, or modelling choices one does not want to change. For instance, the set of integrity constraints can represent physical laws or hierarchies among concepts, similarly to a TBox in description logics.

Dropping the premise that Ψ is consistent, we can assign infinite confidence factors to all its members, while assigning 1 to the conditionals in Γ, forming the weighted base $\Gamma_\Psi \in \mathbb{K}_c^w$:

$$\Gamma_\Psi = \{\alpha\langle 1\rangle \mid \alpha \in \Gamma\} \cup \{\beta\langle \infty\rangle \mid \beta \in \Psi\}$$

The reader can note that, if Ψ is consistent $\mathcal{I}_p^\delta(\Gamma_\Psi) = \mathcal{I}_p^\Psi(\Gamma)$. Since the confidence factors associated to the conditionals in Ψ are infinite, any π minimising $d_p^\delta(\Gamma_\Psi, \pi)$ will satisfy Ψ while minimising $d_p^\varepsilon(\Gamma, \pi)$. Hence, integrity constraints can be simulated in a weighted base by employing infinite confidence factors. Note that, if Ψ is inconsistent, $d_p^\delta(\Gamma_\Psi, \pi)$ is infinite for any π, for at least one conditional in Ψ is always violated.

Potyka and Thimm (2015) illustrate integrity constraints with a version of the lottery paradox (Kyburg Jr, 1961):

Example 7.4.1. Suppose there is a lottery with n tickets, numbered from 1 to n, and it is known, due to the drawing mechanism, that there will be a winner. Nevertheless, for each particular ticket, we do not believe it will win. Let x_i represent the fact that ticket i will win the lottery. Potyka and Thimm (2015) model this situation in two probabilistic bases:

$$\begin{aligned} \Gamma &= \{P(x_i) \leq 0 | 1 \leq i \leq n\} \\ \Psi &= \left\{ P(\bigvee_{i=1}^{n} x_i) \geq 1 \right\} \end{aligned}$$

The base Ψ encodes the knowledge that at least one ticket will win the lottery, and is an integrity

[7] We assume, without loss of generality, that $\Gamma \cap \Psi = \varnothing$ and $\Gamma \cup \Psi$ forms a canonical base.

constraint. To measure the inconsistency of Γ with $\mathcal{I}_{\varepsilon}^{\Psi}$, one looks for the minimal p-norm of the violations vector considering a probabilistic interpretation satisfying Ψ, obtaining, for any $p \geq 2$:

$$\mathcal{I}_p^{\Psi}(\Gamma) = \left\| \langle \frac{1}{n}, \ldots, \frac{1}{n} \rangle \right\|_p = \frac{\sqrt[p]{n}}{n}$$

When $p = 1$, we have $\sum_{i=1}^{n} \varepsilon_i = 1$ and $\mathcal{I}_p^{\Psi}(\Gamma) = 1$, so the general form above also applies.

Now consider the following weighted base in $\mathbb{K}_c^w$:

$$\Gamma_{\Psi} = \{P(x_i) \leq 0\langle 1 \rangle | 1 \leq i \leq n\} \cup \{P(\bigvee_{i=1}^{n} x_i)\langle \infty \rangle \geq 1\}$$

To compute $\mathcal{I}_p^{\delta}$, one has to find the violations minimising $\|\langle 1\varepsilon_1, \ldots, 1\varepsilon_n, \infty\varepsilon_{n+1} \rangle\|$. Again, for any $p \geq 2$, this yields a unique violations vector :

$$\mathcal{I}_p^{\delta}(\Gamma_{\Psi}) = \left\| \langle \frac{1}{n}, \ldots, \frac{1}{n}, 0 \rangle \right\|_p = \frac{\sqrt[p]{n}}{n}$$

If $p = 1$, we have once more $\sum_{i=1}^{n} \varepsilon_i = 1$ and $\mathcal{I}_p^{\delta}(\Gamma_{\Psi}) = 1$.

To see the difference of considering the integrity constraint, or the infinite confidence factor, consider the non-weighted base $\Gamma \cup \Psi$. Applying the minimal violation measure $\mathcal{I}_p^{\varepsilon}$, with $p \geq 2$, yields:

$$\mathcal{I}_p^{\varepsilon}(\Gamma \cup \Psi) = \left\| \langle \frac{1}{n+1}, \ldots, \frac{1}{n+1}, \frac{1}{n+1} \rangle \right\|_p = \frac{\sqrt[p]{n+1}}{n+1}$$

Now the conditional in Ψ is being violated by the probabilistic interpretation π that minimises d_p^{ε}. That is, we are implicitly assuming that in the closest consistent scenario there is a probability of no one winning the lottery, which contradicts the premise. For example, in a situation with two tickets, $\mathcal{I}_p^{\varepsilon}$ considers a closest consistent setting, for $p \geq 2$, with each ticket having 1/3 of winning probability and with the probability of the lottery having having a winner being 2/3. For $p = 2$, this yields $\mathcal{I}_2^{\varepsilon}(\Gamma \cup \Psi) = \sqrt{3}/3$. When we consider the weighted base and compute $\mathcal{I}_p^{\delta}(\Gamma_{\Psi})$ for a $p \geq 2$, the probabilistic interpretation minimising the weighted violations gives each ticket a probability 1/2 for winning the lottery — which still must have a winner. For $p = 2$, this yields $\mathcal{I}_2^{\delta}(\Gamma_{\Psi}) = \sqrt{2}/2$. □

To compute $\mathcal{I}_p^{\delta}$ for a base $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i]\langle \delta_i \rangle | 1 \leq i \leq m\} \in \mathbb{K}_c^w$, we could use δ_i as coefficients to ε_i within the objective function of the program in (3.14)–(3.17), but we choose the equivalent option of dividing the violations within the constraints by the confidence factors. We rewrite the program from (3.14)–(3.17) to include these factors, where ε_{δ} is a $(m \times 1)$-vector whose elements are ε_i/δ_i, for $1 \leq i \leq m$, where ε_i/δ_i is defined as 0 whenever $\delta_i = \infty$. $\mathcal{I}_p^{\delta}(\Gamma)$ is the minimum value of the objective function in the following program:

$$\min \sqrt[p]{\sum_{i=1}^{m} \varepsilon_i^p} \quad \text{subject to:} \tag{7.4}$$

$$A\pi \geq -\varepsilon_\delta \tag{7.5}$$

$$\sum \pi = 1 \tag{7.6}$$

$$\pi, \varepsilon \geq 0\ . \tag{7.7}$$

$\mathcal{I}_p^\delta$ works as a regular inconsistency measure when it is well-defined, but some extreme δ may turn the program infeasible — for instance, when all δ_i are infinite and Γ is inconsistent. Intuitively, the penalty for changing any probability bound in Γ would be infinite, and at least one bound should be relaxed, for Γ is inconsistent. We define $\mathcal{I}_p^\delta(\Gamma) = \infty$ in such cases.

When $p = 1$, (7.4)–(7.7) is a linear program. Again, any finite $p > 1$ leads to non-linear terms within the objective function, but this is not the case for $p = \infty$. Considering all ε_i equal to a single scalar $\varepsilon_{max} \geq 0$, we can modify the linear program (3.19)–(3.21) to allow for the confidence factors. The measure $\mathcal{I}_\infty^\delta$ is the solution of the following program, in which ε_δ is a $(m \times 1)$-vector whose elements are $\varepsilon_{max}/\delta_i$, for $1 \leq i \leq m$:

$$\min \varepsilon_{max} \quad \text{subject to:} \tag{7.8}$$

$$A\pi \geq -\varepsilon_\delta \tag{7.9}$$

$$\sum \pi = 1 \tag{7.10}$$

$$\pi, \varepsilon_{max} \geq 0\ . \tag{7.11}$$

From the $\mathcal{I}_p^\delta$ family, Potyka's measures are recovered when each confidence factor δ_i is equal to 1. Recall that non-weighted bases denote bases where the confidence factors are all equal to 1:

Proposition 7.4.2. *For any $p \in \overline{\mathbb{N}}_{>0}$ and $\Gamma \in \mathbb{K}_c$, $\mathcal{I}_p^\delta(\Gamma) = \mathcal{I}_p^\varepsilon(\Gamma)$.*

Properties from Lemma 3.2.8 also hold for the generalised measures, if their meaning are extended to inconsistency measures on $\mathbb{K}_c^w$ in the obvious way:

Theorem 7.4.3. *For any $p \in \overline{\mathbb{N}}_{>0}$, $\mathcal{I}_p^\delta : \mathbb{K} \to [0, \infty) \cup \{\infty\}$ is well-defined and satisfies (Consistency), (Continuity), (ε-Independence) and (Monotonicity). $\mathcal{I}_1^\delta$ also satisfies (Super-additivity) and (ε-Separability).*

Sufficient, but not necessary conditions for $\mathcal{I}_p^\delta$ to be finite are stated below:

Lemma 7.4.4. *If for any weighted conditional $P(\varphi_i|\psi_i) \geq q_i\langle\delta_i\rangle$ in a base $\Gamma \in \mathbb{K}_c^w$, $\delta_i = \infty$ implies $q_i = 0$, then $\mathcal{I}_p^\delta(\Gamma)$ is finite.*

To explore the conditions on which $\mathcal{I}_p^\delta$ is continuous, we need to define the characteristic function $\Lambda_\Gamma : [0,1]^{|\Gamma|} \to \mathbb{K}_c^w$ of a weighted base $\Gamma \in \mathbb{K}_c^w$. We suppose such function only changes the probabilities, not the confidence factors: if $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i'\langle\delta_i\rangle | 1 \leq i \leq m\}$ and $q = \langle q_1, \ldots, q_m\rangle \in [0,1]^m$, $\Lambda_\Gamma(q) = \{P(\varphi_i|\psi_i) \geq q_i\langle\delta_i\rangle | 1 \leq i \leq m\}$.

Lemma 7.4.5. *For any weighted knowledge base* $\Gamma = \{P(\varphi_i|\psi_i) \geq q'_i\langle\delta_i\rangle | 1 \leq i \leq m\}$ *in* $\mathbb{K}^w_c$*, if* δ_i *is finite for any* $1 \leq i \leq m$*, then* $\mathcal{I}^\delta_p \circ \Lambda_\Gamma(q) : [0,1]^m \to [0,\infty)$ *is continuous.*

Another constraint on the confidence factors ensure normalisation:

Lemma 7.4.6. *For all* $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i]\langle\delta_i\rangle | 1 \leq i \leq m\} \in \mathbb{K}^w_c$*, if* $\delta_i \leq 1/q_i$ *for every* $1 \leq i \leq m$*, then* $\mathcal{I}^\delta_\infty(\Gamma) \in [0,1]$.

After generalizing Potyka's measures based on distance minimisation, their correspondence to the measures based on Dutch books follows. While investigating how to reconcile inconsistent probabilities assigned to events, Nau (1981) developed a framework similar to ours. He started considering a program that minimises the weighted distance between the original and the consolidated probability bounds, focusing on the consolidation problem. Due to non-convexity, he linearises the restrictions in that program — using violations, like Potyka's ε's — to obtain constraints like (7.5)–(7.7), from the linear program that compute $\mathcal{I}^\delta_p$. Nau left the confidence factors within the objective function, not in the linear restrictions, but both approaches are clearly equivalent. Then, he considered the minimisation of the p-norm of these weighted violations, for $p = 1$ and $p = \infty$, proposing linear programs that compute $\mathcal{I}^\delta_1$ and $\mathcal{I}^\delta_\infty$, which he employed to consolidate bases. Nau was fully aware of the duality between the programs that minimise violations and the programs that maximise sure loss in Dutch books. He focused on the program that computes $\mathcal{I}^\delta_\infty$, together with its dual. It is mentioned that such dual computes the maximum sure loss when the δ-weighted sum of bet units (stakes' values) is limited. Nau notes a single interesting particular case, which corresponds to the result below, ignoring the other similar results we prove — although he was probably aware of them.

Lemma 7.4.7. *Consider the knowlegde base* $\Gamma = \{P(\varphi_i|\psi_i)q_i | 1 \leq i \leq m\} \in \mathbb{K}_c$ *and the weighted knowledge base* $\Gamma' = \{P(\varphi_i|\psi_i)q_i\langle\delta_i\rangle | 1 \leq i \leq m\} \in \mathbb{K}^w_c$*. If* $\delta_i = 1/q_i$ *for every* $1 \leq i \leq m$*, then* $\mathcal{I}^{a,sum}_{SSK}(\Gamma) = \mathcal{I}^\delta_\infty(\Gamma')$ *and both are finite.*

When probabilities are assigned only through unconditional lower bounds, $\Gamma = \{P(\varphi_i) \geq q_i \,|1 \leq i \leq m\}, \mathcal{I}^{a,sum}_{SSK}$ is the minimum $\varepsilon \in [0,1]$ such that $\{P(\varphi_i) \geq q_i(1-\varepsilon) | 1 \leq i \leq m\}$ is consistent, which is to compute the least relative decrement in all lower bounds to reach consistency. Interpreting probability lower bounds as the corresponding upper bounds, there is the dual result:

Lemma 7.4.8. *Consider the knowlegde base* $\Gamma = \{P(\varphi_i|\psi_i)q_i | 1 \leq i \leq m\} \in \mathbb{K}_c$ *and the weighted knowledge base* $\Gamma' = \{P(\varphi_i|\psi_i)q_i\langle\delta_i\rangle | 1 \leq i \leq m\} \in \mathbb{K}^w_c$*. If* $\delta_i = 1/(1-q_i)$ *for every* $1 \leq i \leq m$*, then* $\mathcal{I}^{b,sum}_{SSK}(\Gamma) = \mathcal{I}^\delta_\infty(\Gamma')$.

Suppose each conditional $P(\varphi|\psi) \geq q$ is written as $P(\neg\varphi|\psi) \leq 1 - q$. For unconditional probabilities, the result above means that maximising sure loss when the bettor's resources are limited to one is equivalent to minimizing the relative increment in all upper bounds one has to perform in order to consolidate the base. For instance, if Γ contains only unconditional assessments written as probability upper bounds, $\mathcal{I}^b_{SSK}(\Gamma) = 0.05$ means that an increase of 5% in each upper bound is needed to restore consistency.

In general, $\mathcal{I}^\delta_\infty$ measures the maximum guaranteed loss an agent is exposed to through a Dutch book when we limit to one the sum of the stakes' values (λ_i) weighted by $\gamma_i = 1/\delta_i$ — which we call *volatility factors*. This is achieved by adding a constraint $\sum \lambda_i\gamma_i \leq 1$ to the program (7.1)–(7.3),

which maximises sure loss. In other words, the higher the volatility factor γ_i of a probability bound, the more the corresponding stake is constrained.

A more meaningful interpretation to the volatility factors $\gamma_i = 1/\delta_i$ when computing $\mathcal{I}^{\delta}_{\infty}$ can be given by considering the *coverage ratios* $c_i = \gamma_i/q_i$. The number c_i is the ratio between how much the agent has to commit from her resources to cover a gamble on $\varphi_i|\psi_i$ and how much she can lose in that bet. For instance, the weighted conditional $P(\varphi) \geq 0.75\langle 4\rangle$, with a confidence factor 4, corresponds to a volatility factor 0.25. The coverage ratio is $0.25/0.75 = 1/3$, thus the agent has to commit \$1/3 for each \$1 she might lose in a gamble in that conditional. In this setting, $\mathcal{I}^{\delta}_{\infty}$ maximises the sure loss when her resources — the total she can commit — are fixed at 1 and her coverage ratios are $c_i = \gamma_i/q_i$. Analogously, if we define $c'_i = (1-\gamma_i)/q_i$, $\mathcal{I}^{\delta}_{\infty}$ maximises sure profit for the bettor when his resources are limited to 1 and his coverage ratios are c'_i. When the agent's coverage ratios c_i (or the bettor's c'_i) are all equal to one, we have that $\mathcal{I}^{a,sum}_{SSK}(\Gamma) = \mathcal{I}^{\delta}_{\infty}(\Gamma)$ ($\mathcal{I}^{b,sum}_{SSK}(\Gamma) = \mathcal{I}^{\delta}_{\infty}(\Gamma)$).

When $p = 1$, the generalisation of Potyka's measure can be brought to the measure $\mathcal{I}^{max}_{SSK}$, having $\mathcal{I}^{a,max}_{SSK}$ and $\mathcal{I}^{b,max}_{SSK}$ as particular cases. Consider a knowledge base $\Gamma = \{P(\varphi_i|\psi_i)q_i\langle\delta_i\rangle | 1 \leq i \leq m\}$ representing an agent's belief state. We define $\mathcal{I}^{\delta,max}_{SSK}(\Gamma)$ as the maximum sure loss such an agent would be exposed to when gambles have stakes limited via $\lambda_i \leq \delta_i$ for all $1 \leq i \leq m$. Note that, the greater a confidence factor is, the higher the amount the agent accepts to risk at a gamble on the corresponding probability bound.

Theorem 7.4.9. *For any weighted knowledge base* $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i\langle\delta_i\rangle | 1 \leq i \leq m\} \in \mathbb{K}^w_c$, $\mathcal{I}^{\delta,max}_{SSK}(\Gamma) = \mathcal{I}^{\delta}_{1}(\Gamma)$.

Corollary 7.4.10. *Consider a weighted knowledge base* $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i\langle\delta_i\rangle | 1 \leq i \leq m\} \in \mathbb{K}^w_c$ *and a knowledge base* $\Gamma' = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq m\} \in \mathbb{K}_c$. *If* $\delta_i = 1/q_i$ *for all* $1 \leq i \leq m$, *then* $\mathcal{I}^{\delta,max}_{SSK}(\Gamma) = \mathcal{I}^{a,max}_{SSK}(\Gamma') = \mathcal{I}^{\delta}_{1}(\Gamma)$. *If* $\delta_i = 1/(1-q_i)$ *for all* $1 \leq i \leq m$, *then* $\mathcal{I}^{\delta,max}_{SSK}(\Gamma) = \mathcal{I}^{b,max}_{SSK}(\Gamma') = \mathcal{I}^{\delta}_{1}(\Gamma)$.

We even allow confidence factors δ_i to be ∞, defining $\varepsilon_i/\infty = 0$. In such a case, the sure loss may be again unbounded, and the program that compute $\mathcal{I}^{\gamma}_{1}(\Gamma)$, infeasible — $\mathcal{I}^{\delta}_{1}(\Gamma) = \mathcal{I}^{\delta,max}_{SSK}(\Gamma)$ are then defined as ∞.

Properties from Theorem 7.4.3 are satisfied by $\mathcal{I}^{\delta,max}_{SSK}$, including (Super-additivity) and (ε-Separability). Furthermore, it is computable by means of a linear program and has a meaningful interpretation in the gambling scenario. Therefore, the measure $\mathcal{I}^{\delta,max}_{SSK}$ might be appropriate to handle cases where (Super-additivity) and (ε-Separability) are desirable — and confidence factors, available. Additionally, if (Normalisation) is required, we could follow Muiño's (2011) ideas to normalise $\mathcal{I}^{\delta,max}_{SSK}$, although (Super-additivity) and (ε-separability) would no longer hold.

Chapter 8

Conclusion

8.1 Summary

Handling inconsistency has been receiving increased attention from the AI community since most inference methods rely on the consistency of the premises; and such requirement is commonly violated in large bases of probabilistic knowledge. A reasonable starting point for dealing with the inconsistency in probabilistic bases is to know where it is, how severe it is, and how this severity changes with the probabilities. To address these issues, in this thesis we studied different ways of measuring inconsistency in probabilistic knowledge bases, and concepts to localise the inconsistency were developed along the way. Three aspects of inconsistency measures were discussed: postulates they should satisfy, the efficiency of the methods used to compute them, and possible meaningful interpretations for them.

We reviewed the main proposals for measuring inconsistency found in the literature, briefly discussing their computational aspects. At the same time, we explored the rationality postulates these measures should satisfy. We proved that (Consistency), (Independence) and (Continuity) are incompatible postulates. To understand such inconsistency and construct an argument to drop (Independence), we analysed the consolidation methods for probabilistic knowledge bases. Both (Independence) and (MIS-Separability) are expressed in terms of minimal inconsistent sets. As it was argued for, the central role of MISes is due to the classical method to restore consistency, which is performed by ruling formulas out. We presented different methods for consolidating probabilistic knowledge bases, via probability bounds adjustment or by adding violations, linked to the inconsistency measures tailored to probabilistic logic.

Stepping aside the central topic of this thesis, we investigated these probabilistic consolidation procedures under the well-established AGM framework of belief revision. To characterise consolidation via probability changing, we defined the concept of liftable contraction, which is parametrised by a consequence operation. Using this consequence operation, we lifted the postulates for base contraction operations, in order that probabilistic consolidations may be fully characterised. In the end, we devised a probabilistic contraction operation for a specific language, since in the general case such operation is ill-defined.

Back to the main subject of our work, we analysed the relation between consolidation methods and characterisation of primitive conflicts. We found that minimal inconsistent sets do not capture all the inconsistency causes in a probabilistic base that is to be consolidated by relaxing the probability bounds. The connection between MISes and abrupt consolidations was formalized, and by

supposing different underlying consolidation procedures we obtained the concepts of (ε-)innocuous conditional and (ε-)inescapable conflicts. Such new definitions enabled us to formulate new postulates and properties: (i-Independence), (ε-Independence), (IC-separability) and (ε-Separability). We proved that both forms of weakening (Independence) and (MIS-Separability) are compatible with (Consistency) and (Continuity).

Finally, it was proved that incoherence measures via Dutch books are extensionally equivalent to inconsistency measures via minimal violation. In other words, measures that minimise the violations can be interpreted as the sure loss an agent is exposed to in a gambling scenario. We showed how this betting context provides other incoherence measures that can be applied to inconsistency measurement in knowledge bases as well. By considering confidence factors associated to conditionals in a base, we generalised the connection between the two approaches to measuring inconsistency.

8.2 Concluding Remarks

Localising, measuring and repairing probabilistic inconsistencies are intrinsically entangled. The severity of a base's inconsistency is usually understood as the effort needed for consolidating it, and a consolidation method can be defined through the minimisation of an inconsistency measure. Depending on how one wants to restore consistency — for instance, by discarding formulas or changing probabilities —, different it will be the characterisation of the sets of formulas causing the inconsistency. The dual of such definition yields the formulas that are not collaborating to the inconsistency, and it is from these formulas that we want our inconsistency measure to be independent. Hence, by first assuming a consolidation procedure, one can define an independence postulate — as (i-Independence) or (ε-Independence) — to then tell apart rational from irrational inconsistency measures. Similarly, to the question "which probabilities are causing the inconsistency?", one should reply: "how do you intend to consolidate them?", as Example 6.1.1 makes clear. Nevertheless, there is no need to lift the consolidation to the more foundational status, being inconsistency measurement and localisation only secondary operations. For instance, Potyka's minimal violation measures came earlier than the corresponding consolidation procedure, from which we derived the ε-innocuous notion. In principle, it is also possible to depart from a characterisation of primitive conflict, pass through the inconsistency measures that satisfy the corresponding independence postulate and arrive at consolidation methods that minimise these measures. The conclusion is that this triple must be chosen together, or in a consistent way, to avoid situations in which repairing the primitive conflicts does not repair the whole base, or measuring the inconsistency does not reflect the effort to consolidate.

In AI, inconsistency measures for probabilistic knowledge bases have been based on distance minimisation, while in Statistics and formal epistemology incoherence measures for Bayesian agents have focused on Dutch books vulnerability. The connections here established can help both communities to investigate their corresponding problems under a different angle. Statisticians and philosophers can adopt (and probably improve) the rationality postulates approach from computer scientists. Furthermore, they can at least take a look at the computational issues involved, such as the difficulties of non-convex optimisation and the relative benefits of linear programming approaches. Computer scientists, on their turn, can make use of the arguments for the probabilism thesis (that degrees of belief should be consistent) in formal epistemology to devise meaningful inconsistency measures

for probabilistic knowledge bases. For instance, besides Dutch book arguments, probabilism can be endorsed via accuracy considerations.

8.3 Future Work

From a theoretical point of view, this thesis leaves an open question whose answer could bring it a neater closure: are inescapable conflicts and ε-inescapable conflicts the same thing? An affirmative answer would endorse (i-Independence) as a legitimate rationality postulate, for it would be satisfied by the two major families of inconsistency measures, which are linked to two rather different consolidation procedures. We know that this is the case in the unconditional context, but most of the probabilistic knowledge bases practical use comes from conditional probability assignments. Otherwise, if that answer was negative, it would point out the already mentioned intertwining among localising, measuring and repairing probabilistic inconsistency.

The introduced concepts of (ε-)innocuous conditional and (ε-)inescapable conflict will have more practical use in measuring inconsistency only if their instances are recognisable with admissible computational cost. Nothing was said here about the complexity of the computational task of finding innocuous conditionals and inescapable conflicts within a knowledge base, but they are clearly very hard problems. Thus, future work includes investigating these problems aiming at devising algorithms to solve them.

Another possible subject of future work is to investigate the link between proofs of contradiction constructible from an inconsistent probabilistic base and its (ε-)inescapable conflicts. Intuitively, it seems that if a conditional is not essential to any proof of the contradiction, then it is (ε-)innocuous. Equivalently, an (ε-)inescapable conflict is apparently the set of conditionals essentially used in such a proof. Hence, an isomorphism like Curry-Howard's (which links proofs to lambda terms) tailored to the probabilistic logic might be needed to formalise a normal form for these proofs.

A more tangible future work is related to the application of our methods to classical propositional logic. Even though there is no (Continuity) postulate for inconsistency measures in classical logic, and consequently no incompatibility result like ours, there is still the perception that the (Independence) may be too strong in some cases (Hunter and Konieczny, 2010). Based on the same relation among localising, measuring and repairing inconsistency, we could derive new primitive conflict characterisations from different consolidation methods, yielding suitable versions of (Independence). Besnard (2014) has recently parametrised Hunter and Konieczny's postulates by the primitive conflict characterisation, and taking the consolidation procedure as primitive, the resulting postulates of both approaches could be compared.

Under the AGM theory of belief change, future work includes the investigation of probabilistic revision from the (liftable) contraction we defined. Revising probabilities under new evidence is a central problem in epistemology and philosophy of science, and Bayesian conditionalisation is the default approach employed. Nevertheless, it yields some difficulties regarding null probabilities and unknown hypotheses. Furthermore, Dubois and Prade (1997) argued that focusing and revising are fundamentally distinct operations, and conditionalisation fits better the former, as a reference class change. Our proposal for probabilistic contraction can be used to define a revision through expansion, using Levi's identity. Maybe this framework could address some cases where conditionalisation seems irrational, as when the epistemic input brings generic knowledge, instead of factual evidence.

On the practical side, it remains to employ these inconsistency measures to solve real-world problems. Inconsistency measures can be used to define generalised models for (possibly inconsistent) probabilistic knowledge bases, as done by Potyka and Thimm (2015). de Morais, De Bona, and Finger (2015) have recently proposed an approach to the problem of part-of-speech tagging using these tools. Their idea is to capture dependencies between arbitrarily distant words in a sentence while guessing the part-of-speech tags, relying on the absence of independence assumptions in probabilistic logic, in contrast to the standard solution via Markov Hidden Models. An implementation of such approach is currently under construction by the first author of that work.

Appendix A

Proofs of Technical Results

A.1 Technical Results from Chapter 3

Proposition 3.1.7 . *If $\mathcal{I}$ satisfies (MIS-separability) and (Consistency), then $\mathcal{I}$ satisfies (Independence).*

Proof. Let Γ be a knowledge base and $\alpha \in \Gamma$ a free conditional. By MIS-separability, as α is free and all MISes of Γ are in $\Gamma \setminus \{\alpha\}$, we have $\mathcal{I}(\Gamma) = \mathcal{I}(\Gamma \setminus \{\alpha\}) + \mathcal{I}(\alpha)$. As $\{\alpha\}$ is not a MIS, it is consistent; and, by (Consistency), $\mathcal{I}(\{\alpha\}) = 0$. Finally, $\mathcal{I}(\Gamma) = \mathcal{I}(\Gamma \setminus \{\alpha\})$. □

Proposition 3.1.8 . *$\mathcal{I}_{MIS}$ and $\mathcal{I}_{MIS^C}$ satisfy (Consistency), (Monotonicity), (Independence) and (MIS-separability).*

Proof. **(Consistency):** A given base $\Gamma \in \mathbb{K}_c$ is consistent iff it contains no MIS. Note that $\mathcal{I}_{MIS}(\Gamma) = 0$ (and $\mathcal{I}_{MIS^C}(\Gamma) = 0$) if Γ contains no MIS; and any MIS in Γ implies $\mathcal{I}_{MIS}(\Gamma) > 0$ (and $\mathcal{I}_{MIS^C}(\Gamma) > 0$).

(MIS-Separability): Consider $\Psi, \Delta\Gamma \in \mathbb{K}_c$ such that $\Gamma = \Delta \cup \Psi$, $\mathrm{MIS}(\Gamma) = \mathrm{MIS}(\Delta) \cup \mathrm{MIS}(\Psi)$ and $\mathrm{MIS}(\Delta) \cap \mathrm{MIS}(\Psi) = \varnothing$. If $\mathcal{I}$ is either $\mathcal{I}_{MIS}$ or $\mathcal{I}_{MIS^C}$, then it holds that:

$$\begin{aligned}
\mathcal{I}(\Gamma) &= \sum_{\Theta \in \mathrm{MIS}(\Gamma)} \mathcal{I}(\Theta) \\
\mathcal{I}(\Gamma) &= \sum_{\Theta \in \mathrm{MIS}(\Psi)} \mathcal{I}(\Theta) + \sum_{\Theta \in \mathrm{MIS}(\Delta)} \mathcal{I}(\Theta) \\
\mathcal{I}(\Gamma) &= \mathcal{I}(\Psi) + \mathcal{I}(\Delta)\,.
\end{aligned}$$

(Independence): As both measures satisfy (Consistency) and (MIS-Separability), the result follows from Proposition 3.1.7.

(Monotonicity): Consider the probabilistic bases $\Gamma, \Gamma \cup \{\alpha\}\mathbb{K}_c$ Again, if $\mathcal{I}$ is either $\mathcal{I}_{MIS}$ or $\mathcal{I}_{MIS^C}$, then it holds that:

$$\mathcal{I}(\Gamma) = \sum_{\Theta \in \mathrm{MIS}(\Gamma)} \mathcal{I}(\Theta)$$

For (Monotonicity) to hold, just note that $\mathrm{MIS}(\Gamma) \subseteq \mathrm{MIS}(\Gamma \cup \{\alpha\})$. □

Proposition 3.1.11 . *If $\mathcal{I}$ satisfies (Super-additivity), then $\mathcal{I}$ satisfies (Monotonicity).*

Proof. Consider an arbitrary base $\Gamma \in \mathbb{K}_c$ and an arbitrary conditional $\alpha \notin \Gamma$. If $\mathcal{I}$ satisfies (Super-additivity), then $\mathcal{I}(\Gamma \cup \{\alpha\}) \geq \mathcal{I}(\Gamma) + \mathcal{I}(\{\alpha\})$. As $\mathcal{I}(\{\alpha\}) \geq 0$, $\mathcal{I}(\Gamma \cup \{\alpha\}) \geq \mathcal{I}(\Gamma)$ and (Monotonicity) holds. □

Proposition 3.1.12 . *$\mathcal{I}_\eta : \mathbb{K}_c \to [0, \infty)$ satisfies (Consistency), (Monotonicity), (Independence), and (Normalisation) and violates (MIS-Separability).*

Proof. **(Consistency):** Consider a base $\Gamma \in \mathbb{K}_c$. If it is consistent, there is a probability mass $\pi_1 : W_n \to [0, 1]$ satisfying all its conditionals. Hence, there is a probability mass $\pi_2 : \Pi_n \to [0, 1]$ with $\pi_2(\pi_1) = 1$ that $P_{\pi_2}(\alpha) \geq 1$ for all $\alpha \in \Gamma$; and $\mathcal{I}_\eta(\Gamma) = 0$. Now suppose Γ is inconsistent. For any probability mass $\pi_2 : \Pi_n \to [0, 1]$, there must be some probabilistic interpretation $\pi_1 : W_n \to [0, 1]$ such that $\pi_2(\pi_1) > 0$. Since π_1 cannot satisfy all conditionals in Γ, for it is inconsistent, there is some $\alpha \in \Gamma$ not satisfied by it. Hence, for π_2 is a probability mass, $(\sum_{\pi \in \Pi_n} \pi_2(\pi)) - \pi_2(\pi_1) < 1$ and $P_{\pi_2}(\alpha) < 1$. Consequently:

$$\max \{\eta | \{P(\alpha_i) \geq \eta | \alpha_i \in \Gamma\} \text{ is consistent}\} < 1 \,,$$

and $\mathcal{I}_\eta(\Gamma) > 0$.

(Monotonicity): Consider an arbitrary base $\Gamma \in \mathbb{K}_c$ and an arbitrary conditional $\beta \notin \Gamma$. Let $\eta^\star$ be defined as:

$$\eta^\star = \max \{\eta | \{P(\alpha) \geq \eta | \alpha \in \Gamma \cup \{\beta\}\} \text{ is consistent}\} \,,$$

so that $\mathcal{I}_\eta(\Gamma \cup \{\beta\}) = 1 - \eta^\star$. If $\{P(\alpha) \geq \eta^\star | \alpha \in \Gamma \cup \{\beta\}\}$ is consistent, then $\{P(\alpha) \geq \eta^\star | \alpha \in \Gamma\}$ is consistent. Hence, $\mathcal{I}_\eta(\Gamma) \leq 1 - \eta^\star = \mathcal{I}_\eta(\Gamma \cup \{\beta\})$, and (Monotonicity) holds.

(Independence): Consider an arbitrary base $\Gamma \in \mathbb{K}_c$ and a free conditional $\beta \notin \Gamma$. Let Ψ denote the base $\Gamma \setminus \{\beta\}$. Suppose that $\mathcal{I}_\eta(\Psi) = 1 - \eta^\star$, so that $\{P(\alpha) \geq \eta^\star | \alpha \in \Psi\}$ is consistent. Let $\pi_2 : \Pi_n \to [0, 1]$ be such that $P_{\pi_2}(\alpha) \geq \eta^\star$ for all $\alpha \in \Psi$. We are going to construct a probability mass $\pi_2' : \Pi_n \to [0, 1]$ such that $P_{\pi_2'}(\alpha) \geq \eta^\star$ for all $\alpha \in \Psi$. Then, for each $\pi \in \Pi_n$, consider the set $\Delta \subseteq \Psi$ of conditionals satisfied by it. As Δ is consistent and β is free in Γ, $\Delta \cup \{\beta\}$ is consistent, being satisfied by some π'. Let $f : \Pi_n \to \Pi_n$ be a function such that for any $\pi \in \Pi_n$, $f(\pi)$ returns the corresponding π'. Now define $\pi_2' : \Pi_n \to [0, 1]$ for any $\pi_1 \in \Pi_n$ as:

$$\pi_2'(\pi_1) = \sum \{\pi_2(\pi) \,|\, f(\pi) = \pi_1\}$$

By construction, $P_{\pi_2'}(\alpha) = P_{\pi_2}(\alpha)$ for all $\alpha \in \Psi$ and $P_{\pi_2'}(\beta) = 1$. Hence, the set $\{P(\alpha) \geq \eta^\star | \alpha \in \Gamma\}$ is consistent and $\mathcal{I}_\eta(\Gamma) \leq 1 - \eta^\star$. By (Monotonicity), $\mathcal{I}_\eta(\Gamma) = 1 - \eta^\star$.

(Normalisation): For any base $\Gamma \in \mathbb{K}_c$, $\{P(\alpha) \geq 0 | \alpha \in \Gamma\}$ is consistent, so $\mathcal{I}_\eta(\Gamma) \geq 0$. Furthermore, it is clear that the set $\{P(\alpha) \geq \eta | \alpha \in \Gamma\}$ is inconsistent when $\eta > 1$, yielding $\mathcal{I}_\eta(\Gamma) \leq 1$.

(MIS-Separability): To show a counter-example, consider the bases $\Delta = \{P(\top) \geq 1\}$, $\Psi = \{P(\top \wedge \top) \geq 1\}$ and $\Gamma = \Delta \cup \Psi$ in $\mathbb{K}_c$. Note that the only MISes in Γ are Δ and Ψ. As no probabilistic interpretation $\pi : W_n \to [0, 1]$ can satisfy $\alpha = P(\top) \geq 1$, the set $\{P(\alpha) \geq \eta\}$ is inconsistent for any $\eta > 0$. Hence, $\mathcal{I}_\eta(\Delta) = 1$ and, analogously, $\mathcal{I}_\eta(\Psi) = 1$. By (Normalisation), $\mathcal{I}_\eta(\Gamma) \leq 1 < 2 = \mathcal{I}_\eta(\Delta) + \mathcal{I}_\eta(\Psi)$, and (MIS-Separability) fails.

□

Theorem 3.2.3 . *For any $p \in \overline{\mathbb{N}}_{>0}$, $\mathcal{I}_p$ is well-defined and satisfies (Consistency), (Continuity) and (Monotonicity), but not (Independence).*

Proof. To show that $\mathcal{I}_p$ is well-defined, we use results from the proof of Theorem 1 in (Thimm, 2013). For any $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq m\}$, Thimm shows that the set $Q_\Gamma = \{\langle q_1, \ldots, q_m\rangle \in \mathbb{R}^m | \Lambda_\Gamma^P(\langle q_1, \ldots, q_m\rangle)$ is consistent$\}$ is compact and closed, where $\Lambda_\Gamma^P : [0,1]^m \to \mathbb{K}$ is the characteristic function of Γ that assigns precise probabilities. Let $h : \mathbb{R}^2 \to \mathbb{R}$ be a function such that $h(a,b) = \max(a-b, 0)$ for any $a, b \in \mathbb{R}$. The measure $\mathcal{I}_p$ is the minimum of $\|f_q(r)\|_p$ with $r \in Q_\Gamma$, where $f_q : \mathbb{R}^m \to \mathbb{R}^{2m}$ is a function such that $f_q(\langle r_1, \ldots, r_m\rangle) = \langle h(q_1, r_1), \ldots, h(q_m, r_m)\rangle$. Intuitively, $f_q(r)$ measures, for a given point r, how much the lower bounds q_i have to be relaxed for we to have $q_i \leq r_i$. Finally, $\mathcal{I}_p$ is well defined, for Q_Γ is closed and compact (Thimm, 2013).

(Consistency): By definition, a p-norm is never negative, thus $\mathcal{I}_p(\Gamma) \geq 0$. Suppose $\Gamma = \Lambda_\Gamma(q)$ is consistent. A vector $q' = q$ is such that $\|q' - q\|_p = 0$ for any $p \in \overline{\mathbb{N}}_{>0}$, thus $\mathcal{I}_p(\Gamma) = 0$. Now suppose $\Gamma = \Lambda_\Gamma(q)$ is inconsistent. For every $q' \in Q_\Gamma$, $q' \neq q$, then $\|q' - q\|_p > 0$ and $\mathcal{I}_p(\Gamma) > 0$ for any $p \in \overline{\mathbb{N}}_{>0}$.

(Continuity): Given a base $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq m\}$, its characteristic function $\Lambda_\Gamma : [0,1]^m \to \mathbb{K}$ and a fixed $r \in Q_\Gamma$, define the function $g_r : \mathbb{R}^m \to \mathbb{R}$ such that $g_r(\langle q_1, \ldots, q_m\rangle) = \|f_q(r)\|_p$. Note that $\mathcal{I}_p \circ \Lambda_\Gamma(\langle q_1, \ldots, q_m\rangle)$ is computed as the minimum of $\{g_r(\langle q_1, \ldots, q_m\rangle) | r \in Q_\Gamma\}$. Each g_r is continuous, and the minimum of continuous functions is continuous, hence $\mathcal{I}_p \circ \Lambda_\Gamma$ is continuous.

(Monotonicity): Let $\Lambda_\Gamma(q')$ be a natural consolidation of $\Gamma = \Lambda_\Gamma(q)$ such that $\|q' - q\|_p$ is minimised, for a $p \in \overline{\mathbb{N}}_{>0}$, and $\mathcal{I}_p(\Gamma) = \|q' - q\|_p$. To prove by contradiction, suppose $\mathcal{I}_p(\Gamma \cup \{\alpha\}) < \mathcal{I}_p(\Gamma)$, for some $\Psi = \Gamma \cup \{\alpha\} \in \mathbb{K}_c$. Hence, there is a natural consolidation $\Psi' = \Lambda_\Psi(r')$ of $\Psi = \Lambda_\Psi(r)$ such that $\|r' - r\|_p < \|q' - q\|_p$. Consider the base $\Gamma' = \Psi' \setminus \{\beta\}$, such that $\alpha \geq \beta$. As Ψ' is consistent, $\Gamma' = \Lambda_\Gamma(q'')$ also is, and it is a natural consolidation of Γ. Since q and q'' are projections (subsets, in a sense) of r and r', $q'' - q$ is a projection of $r' - r$ and $\|q'' - q\|_p \leq \|r' - r\|_p < \|q' - q\|_p$. Finally, it would follow that $\mathcal{I}_p(\Gamma) \leq \|q'' - q\|_p < \|q' - q\|_p = \mathcal{I}_p(\Gamma)$, which is a contradiction.

That (Independence) does not hold is a consequence of Theorem 3.4.1. □

Lemma 3.2.4 . *$\mathcal{I}_p$ satisfies (Super-additivity) iff $p = 1$.*

Proof. ($\rightarrow$) To note that (Super-additivity) does not hold if $p > 1$, consider the bases $\Psi = \{P(\top) \leq 0.9\}, \Delta = \{P(\bot \wedge \bot) \geq 0.1\}, \Gamma = \Psi \cup \Delta$[1]. Note that Ψ and Δ are MISes. By the definition of maximal consolidation, if $\mathcal{I}_p(\Gamma) = d$, then there is maximal consolidation $\Lambda_\Gamma(q')$ of $\Gamma = \Lambda_\Gamma(q)$ such that $\|q' - q\|_p = d$. The only maximal consolidations of Ψ, Δ, Γ are $\Psi' = \{P(\top) \geq 1]\}, \Delta' = \{P(\bot \wedge \bot) \geq 0\}, \Gamma' = \Psi' \cup \Delta'$. For any finite p, $\mathcal{I}_p(\Psi) = \mathcal{I}_p(\Delta) = \sqrt[p]{0.1^p} = 0.1$, and $\mathcal{I}_p(\Gamma) = \sqrt[p]{0.1^p + 0.1^p} = 0.1\sqrt[p]{2}$. For $p = \infty$, $\mathcal{I}_p(\Psi) = \mathcal{I}_p(\Delta) = \max\langle 0.1\rangle = 0.1$ and $\mathcal{I}_p(\Gamma) = \max\langle 0.1, 0.1\rangle = 0.1$. Therefore, for any $p > 1 \in \mathbb{N} \cup \{\infty\}$, $\mathcal{I}_p(\Gamma) = 0.1\sqrt[p]{2} < 0.2 = \mathcal{I}_p(\Psi) + \mathcal{I}_p(\Delta)$, and (Super-additivity) fails.

($\leftarrow$) Now fix $p = 1$. To prove that (Super-additivity) holds, suppose there are bases $\Psi, \Delta, \Gamma = \Psi \cup \Delta$ in $\mathbb{K}_c$ such that $\Psi \cap \Delta = \varnothing$. Let $\Psi' = \Lambda_\Psi(q'), \Delta' = \Lambda_\Delta(r'), \Gamma' = \Lambda_\Gamma(s')$ be maximal consolidations of $\Psi = \Lambda_\Psi(q), \Delta = \Lambda_\Delta(r), \Gamma = \Lambda_\Gamma(s)$ that minimise $\|q' - q\|_1, \|r' - r\|_1, \|s' - s\|_1$, corresponding to $\mathcal{I}_1(\Psi), \mathcal{I}_1(\Delta), \mathcal{I}_1(\Gamma)$. Clearly, Γ' can be partitioned into $\Psi'' \cup \Delta''$ in such a way that $\Psi'' = \Lambda_\Psi(s'_\Psi), \Delta'' = \Lambda_\Delta(s'_\Delta)$ are natural consolidations of Ψ, Δ. By the construction of s'_Ψ and

[1]Note that $P(\top) \leq 0.9$ abbreviates $P(\bot) \geq 0.1$, so we use $\bot \wedge \bot$ in Δ to keep Γ canonical.

s'_Δ, $\|s' - s\|_1 = \|s'_\Psi - q\|_1 + \|s'_\Delta - r\|_1$. Hence, for $\mathcal{I}_1(\Psi) \leq \|s'_\Psi - q\|_1$ and $\mathcal{I}_1(\Delta) \leq \|s'_\Delta - r\|_1$, it follows that $\mathcal{I}(\Gamma) = \|s' - s\|_1 \geq \mathcal{I}_1(\Psi) + \mathcal{I}_1(\Delta)$. □

Lemma 3.2.5 . *$\mathcal{I}_p$ satisfies (Normalisation) iff $p = \infty$.*

Proof. ($\rightarrow$) To note that (Normalisation) does not hold if p is finite, consider the base $\Gamma = \{P(\top) \leq 0, P(\bot \wedge \bot) \geq 1\}$. The only maximal consolidation of Γ is $\Gamma' = \{P(\top) \geq 1, P(\bot \wedge \bot) \leq 0\}$. For any finite p, $\mathcal{I}_p(\Gamma) = \sqrt[p]{1^p + 1^p} = \sqrt[p]{2} > 1$, and (Normalisation) fails.

($\leftarrow$) By definition, $\mathcal{I}_\infty(\Gamma)$ is the minimum of $\|q' - q\|_\infty$ subject to $\Gamma = \Lambda_\Gamma(q)$ and $\Lambda_\Gamma(q')$ being consistent. As the vectors q, q' are in $[0,1]^{|\Gamma|}$, $\|q' - q\|_\infty \in [0,1]$, since $|q'_i - q_i| \in [0,1]$ for all elements q_i, q'_i of q, q'. □

Proposition 3.2.7 . *For any unconditional probabilistic knowledge base $\Gamma \in \mathbb{K}_c$ and any $p \in \overline{\mathbb{N}}_{>0}$, $\mathcal{I}^\varepsilon_p(\Gamma) = \mathcal{I}_p(\Gamma)$.*

Proof. Consider a base $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq m\}$. Just note that, for all $1 \leq i \leq m$, when $\psi_i = \top$ and consequently $P_\pi(\psi_i) = 1$, the program (3.9)-(3.13) becomes the linear program (3.14)-(3.17), reaching the same minimum for the objective function $\|\langle \varepsilon_1, \ldots, \varepsilon_m \rangle\|_p$ (Potyka, 2014). □

Theorem 3.2.8 . *For any $p \in \overline{\mathbb{N}}_{>0}$, $\mathcal{I}^\varepsilon_p : \mathbb{K} \rightarrow [0, \infty)$ is well-defined and satisfies (Consistency), (Continuity) and (Monotonicity), but not (Independence). $\mathcal{I}^\varepsilon_p$ satisfies (Super-additivity) iff $p = 1$; and $\mathcal{I}^\varepsilon_p$ satisfies (Normalisation) iff $p = \infty$.*

Proof. To see that, for any $p \in \overline{\mathbb{N}}_{>0}$, $\mathcal{I}^\varepsilon_p : \mathbb{K} \rightarrow [0, \infty)$ is well-defined and satisfies (Consistency) and (Monotonicity), and that $\mathcal{I}^\varepsilon_1$ satisfies (Super-additivity), see the proof of Theorem 7.4.3, taking $\delta_i = 1$ for any confidence factor δ_i, due to Proposition 7.4.2. For (Continuity), see Lemma 7.4.5, again taking $\delta_i = 1$ for any confidence factor δ_i, by Proposition 7.4.2.

To prove that $\mathcal{I}^\varepsilon_p$ violates (Super-additivity) for any $p > 1$, just look at the counterexample given in the proof of Lemma 3.2.4, since, by Proposition 3.2.7, $\mathcal{I}_p(\Gamma) = \mathcal{I}^\varepsilon_p(\Gamma)$ for any unconditional $\Gamma \in \mathbb{K}_c$.

For (Normalisation), we note that $\mathcal{I}^\varepsilon_\infty(\Gamma) = \mathcal{I}^{sum}_{SSK}(\Gamma)$ for any $\Gamma \in \mathbb{K}$, by Theorem 7.2.2. When we are computing $\mathcal{I}^{sum}_{SSK}$, we limit the sum of the absolute values of the stakes to one. As the agent cannot lose more than the value of the stake in each gamble in a Dutch book, and they sum up to one, both $\mathcal{I}^{sum}_{SSK}$ and $\mathcal{I}^\varepsilon_\infty$ satisfy (Normalisation). To prove that $\mathcal{I}^\varepsilon_p$ violates (Super-additivity) for any $p < \infty$, just look at the counterexample given in the proof of Lemma 3.2.5, since, by Proposition 3.2.7, $\mathcal{I}_p(\Gamma) = \mathcal{I}^\varepsilon_p(\Gamma)$ for any unconditional $\Gamma \in \mathbb{K}_c$.

That (Independence) does not hold is a consequence of Theorem 3.4.1. □

Theorem 3.3.2 . *$\mathcal{I}_{CRV}$ is well-defined and satisfies the postulates of (Consistency), (Continuity) and (Monotonicity), but not (Independence).*

Proof. Capotorti, Regoli, and Vattari's (2010) Theorem 3 implies that $\mathcal{I}^P_{CRV}$ (for precise bases) is well-defined, and his Theorem 1, that it satisfies (Consistency). To extend these results to $\mathcal{I}_{CRV}$, consider a base $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq m\}$. We firstly need to prove that $\mathcal{I}^P_{CRV}$ (for precise bases) is continuous; that is, that $\mathcal{I}^P_{CRV} \circ \Lambda^P_\Gamma : [0,1]^m \rightarrow [0, \infty)$ is continuous. Adopting the convention $0 \ln 0 = 0$, it is clear that function $f_\pi(q) = d_{CRV}(q, \pi)$ is continuous on q for any fixed probabilistic interpretation π. As $\mathcal{I}^P_{CRV}(\Gamma)$ is the minimum (on π) of continuous functions $f_\pi(q)$, it is also continuous.

Recall that $\Lambda^P_\Gamma(\langle r_1, \ldots, r_m\rangle)$ is the precise base $\Gamma = \{P(\varphi_i|\psi_i) = q_i | 1 \leq i \leq m\}$. $\mathcal{I}_{CRV}(\Gamma)$ is defined as the minimum $\mathcal{I}^P_{CRV}(\Lambda^P_\Gamma(r))$, subject to $r \geq q$. For a given $q \in [0,1]^m$, the set $S = \{r \in [0,1]^m | r_i \geq q_i,\ 1 \leq i \leq m\}$ is closed. As $\mathcal{I}^P_{CRV}$ is continuous and S is closed and bounded, the minimum of $\{\mathcal{I}^P_{CRV} \circ \Lambda^P_\Gamma(r) | r \in S\}$ is well-defined; therefore, $\mathcal{I}_{CRV}(\Gamma) = \min\{\mathcal{I}^P_{CRV} \circ \Lambda^P_\Gamma(r) | r \in S\}$ also is. Furthermore, since $\mathcal{I}^P_{CRV} \circ \Lambda^P_\Gamma$ is continuous, it can be proved that $\mathcal{I}_{CRV}$ satisfies (Continuity), even though the details are omitted.

To see that (Consistency) holds, note that $\mathcal{I}_{CRV}(\Gamma) = 0$ iff there is some $r \geq q$ with $\mathcal{I}^P_{CRV}(\Lambda^P_\Gamma(r)) = 0$, which is the case iff there is some $r \geq q$ such that $\Lambda^P_\Gamma(r)$ is consistent — by the consistency of $\mathcal{I}^P_{CRV}$. If $\Lambda^P_\Gamma(r)$ is consistent for some $r \geq q$, so is $\Lambda_\Gamma(r)$ and, consequently, Γ. Conversely, if Γ is consistent, there is a probabilistic interpretation π satisfying it. Taking $r_i = P_\pi(\varphi_i \wedge \psi_i)/P_\pi(\psi_i)$ if $P_\pi(\psi_i) > 0$ and $r_i = q_i$ otherwise, for each $1 \leq i \leq m$, we have a consistent $\Lambda^P_\Gamma(r)$ with $r \geq q$ — and (Consistency) is satisfied.

The monotonicity of $\mathcal{I}_{CRV}$ follows directly from the monotonicity of $\mathcal{I}^P_{CRV}$. To prove the latter, consider a precise base $\Psi = \Gamma \cup \{P(\varphi|\psi) = q\}$ in $\mathbb{K}_c$ and a probabilistic interpretation π. It holds that:

$$d_{CRV}(\Psi, \pi) = d_{CRV}(\Gamma, \pi) + d_{CRV}(\{P(\varphi|\psi) = q\}, \pi)\ .$$

Capotorti *et al.* (2010) proved (in Theorem 1) that the discrepancy d_{CRV} is non-negative, so $d_{CRV}(\Psi, \pi) \geq d_{CRV}(\Gamma, \pi)$ for any probabilistic interpretation π. Hence, $\mathcal{I}^P_{CRV}(\Psi) \geq \mathcal{I}^P_{CRV}(\Gamma)$, and $\mathcal{I}_{CRV}$ satisfies (Monotonicity).

That (Independence) does not hold is a consequence of Theorem 3.4.1. □

Corollary 3.4.2 . *There is no inconsistency measure $\mathcal{I} : \mathbb{K}_c \to [0, \infty)$ that satisfies (Consistency), (MIS-separability) and (Continuity).*

Proof. It follows directly from Theorem 3.4.1 and Proposition 3.1.7. □

A.2 Technical Results from Chapter 4

Proposition 4.2.8 . $\Gamma!$ *satisfies (Success), (Inclusion) and (Relevance) iff* $\Gamma!$ *is a partial meet consolidation.*

Proof. See the proof of Theorem 5.1.15. Note that a base $\Gamma \in \mathbb{K}_c$ is consistent iff $\bot \notin Cn_{Pr}(\Gamma)$; moreover, Cn_{Pr} is Tarskian and satisfies the upper bound property for consolidation (Lemma 5.2.8). □

Proposition 4.2.12 . $\Gamma!$ *satisfies (Success), (Inclusion), (Core-retainment) iff* $\Gamma!$ *is a kernel consolidation.*

Proof. See the proof of Theorem 5.2.10 in (Wassermann, 2000), noting that, despite the compactness of Cn being mentioned in the Theorem, only the upper bound property is used in the proof. Note also that the proof works for a general language $\mathcal{L}$ containing $\bot$. Note again that a base $\Gamma \in \mathbb{K}_c$ is consistent iff $\bot \notin Cn_{Pr}(\Gamma)$; moreover, Cn_{Pr} is Tarskian and satisfies the upper bound property for consolidation (Lemma 5.2.8). □

Proposition 4.3.9 . *For any canonical base $\Gamma \in \mathbb{K}_c$ and $p \in \overline{\mathbb{N}}_{>0}$, $\mathcal{C}_p(\Gamma)$ is a maximal consolidation of Γ.*

Proof. To prove by contradiction, suppose $\mathcal{C}_p(\Gamma) = \Lambda_\Gamma(q'')$ is not a maximal consolidation, for a $p \in \overline{\mathbb{N}}_{>0}$. Hence, there must be a maximal consolidation $\Gamma' = \Lambda_\Gamma(q')$ of $\Gamma = \Lambda_\Gamma(q)$ such that $q_i \geq q'_i \geq q''_i$ for all $1 \leq i \leq |\Gamma|$ and $q'_i > q''_i$ for some $1 \leq i \leq |\Gamma|$; otherwise $\mathcal{C}_p(\Gamma)$ would be a maximal consolidation. Now it follows that $\|q - q'\|_p < \|q - q''\|_p$, and $\mathcal{C}_p(\Gamma)$ is not a consistent $\Lambda_\Gamma(r)$ that minimizes $d_p(q, r)$, a contradiction. □

Lemma 4.3.10 . *For any $p \in \overline{\mathbb{N}}_{>0}$, $\mathcal{C}_p$ is well-defined and satisfies (Success), (Vacuity), (Structural Preservation) and (Pareto-Optimality).*

Proof. For $\mathcal{C}_p(\Gamma)$ to be well-defined, for some arbitrary $\Gamma = \Lambda_\Gamma(q) \in \mathbb{K}_c$, the set $D_p(\Gamma)$ must not be empty, so it suffices to prove that the minimum of $\{d_p(q, q') | \Lambda_\Gamma(q') \text{ is consistent}\}$ is well-defined. This can be found in the proof of Theorem 3.2.3.

By definition, $\mathcal{C}_p(\Gamma) \in D_p(\Gamma)$ is a consistent $\Lambda_\Gamma(r)$, so (Success) and (Structural Preservation) are trivially satisfied. As p-norms are non-negative, and $d_p(q, r) = 0$ iff $q = r$, (Vacuity) also holds. (Pareto-Optimality) follows from Proposition 4.3.9. □

Proposition 4.3.12 . *For any uncoditional probabilistic knowledge base $\Gamma \in \mathbb{K}_c$ and any integer $p > 1$, $\mathcal{C}^\varepsilon_p(\Gamma) = \mathcal{C}_p(\Gamma)$.*

Proof. Consider a base $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq m\}$. Just note that, for all $1 \leq i \leq m$, when $\psi_i = \top$ and thusly $P_\pi(\psi_i) = 1$, the program (3.9)-(3.13) becomes the linear program (3.14)-(3.17), reaching the same solution on $\varepsilon_1, \ldots, \varepsilon_m$ (Potyka, 2014). By definition, $\mathcal{C}^\varepsilon_p(\Gamma) = \{P(\varphi_i|\psi_i) \geq q_i - \varepsilon_i | 1 \leq i \leq m\}$. To compute $\mathcal{C}^\varepsilon_p$, consider a solution on π for the program (3.14)-(3.17). By definition, $\mathcal{C}^\varepsilon_p(\Gamma) = \{P(\varphi_i|\psi_i) \geq \min\{q_i, P_\pi(\varphi_i)\} | 1 \leq i \leq m\}$. As the values of $\varepsilon_1, \ldots, \varepsilon_m$ minimising $d^\varepsilon_p(\Gamma, \pi)$ are such that $-\varepsilon_i = \min\{0, P_\pi(\varphi) - q_i\}$ $\min\{q_i, P_\pi(\varphi_i)\} = q_i - \varepsilon_i$, finishing the proof. □

Proposition 4.3.13 . *For any $p \in \mathbb{N}_{>0} \cup \{\infty\}$, $\mathcal{C}^\varepsilon_p$ satisfies (Success), (Vacuity), (Structural Preservation) and (Non-Strengthening).*

Proof. To see that (Success) holds, consider a base $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq m\}$ and the solution on π to the corresponding program (3.14)-(3.17), which minimises $d^\varepsilon_p(\Gamma, \pi)$. Let $\mathcal{C}^\varepsilon_p(\Gamma)$ be denoted by $\Gamma' = \{P(\varphi_i|\psi_i) \geq q'_i | 1 \leq i \leq m\}$. For each $1 \leq i \leq m$, if $P_\pi(\psi_i)_i = 0$, then π trivially satisfies $P(\varphi_i|\psi_i) \geq q'_i$; else, $q'_i = \min\{\frac{P_\pi(\varphi_i \wedge \psi_i)}{P_\pi(\psi_i)}, q_i\}$, and $q'_i P_\pi(\psi_i) \leq P_\pi(\varphi_i \wedge \psi_i)$, so π satisfies $P(\varphi_i|\psi_i) \geq q'_i$ as well. Therefore, π satisfies $\mathcal{C}^\varepsilon_p(\Gamma)$, which is consistent. (Non-Strengthening) follows trivially from the definition of $\mathcal{C}^\varepsilon_p(\Gamma)$.

As p-norms are non-negative, and $d^\varepsilon_p(\Gamma, \pi) = 0$ iff π satisfies Γ, (Vacuity) also holds.

By definition, $\mathcal{C}^\varepsilon_p(\Gamma)$ has the form $\Lambda_\Gamma(r)$, so (Structural Preservation) is trivially satisfied. □

Proposition 4.3.14 . *For any $p \in \mathbb{N}_{>0} \cup \{\infty\}$, $\mathcal{C}^\varepsilon_p$ violates (Pareto-Optimality).*

Proof. We show two counterexamples, one for $p = 1$ and the other for the remaining cases. For $p \geq 2$, consider the base

$$\Gamma = \{P(x_1 \wedge x_2|x_3) \geq 0.55, P(x_1 \wedge \neg x_2|x_3) \geq 0.55, P(\neg x_1 \wedge x_2|x_3) \geq 0.55, P(\neg x_1 \wedge \neg x_2|x_3) \geq 0.55, P(x_3) \geq 1\} .$$

If $W_n = \{w_1, \ldots, w_{2^n}\}$ is the set of possible worlds over $\{x_1, \ldots, x_n\}$, Γ can be denoted by:

$$\Gamma = \{P(w_i|x_3) \geq 0.55 | w_i \in W_2\} \cup \{P(x_3) \geq 1\} .$$

As the probability of the possible words $w_1, \ldots, w_4$ given x_3 must sum to 1 if $P_\pi(x_3) > 0$, Γ is inconsistent. Note that, for symmetry reasons, a probability interpretation π such that $P_\pi(w_i \wedge x_3) = 0.25y, P_\pi(x_3) = y$, for any $w_i \in W_2$, is a solution to the minimisation of $d_p^\varepsilon(\Gamma, \pi)$ for some value of $y \in [0, 1]$ (as a function of p). Such π corresponds to the following values for the valuations $\varepsilon_1, \ldots, \varepsilon_5$:

$$\begin{aligned}\varepsilon_1 = -P_\pi(x_1 \wedge x_2 \wedge x_3) + 0.55P_\pi(x_3) = 0.3P_\pi(x_3) &= 0.3y \\ \varepsilon_2 = \varepsilon_3 = \varepsilon_4 = \varepsilon_1 &= 0.3y \\ \varepsilon_5 = -P_\pi(x_3) + 1 &= 1 - y\end{aligned}$$

Note that no ε_i is negative. The p-norm of $\varepsilon = \langle \varepsilon_1, \ldots, \varepsilon_5 \rangle$ is $\|\varepsilon\|_p = \sqrt[p]{4(0.3y)^p + (1-y)^p}$. We look for the y minimising $f_p(y) = \|\varepsilon\|_p$ for each $p \in \mathbb{N}_{>0} \cup \{\infty\}$, $p \geq 2$. One can note that $f_p(0) = 1$ and $f_p(1) = 0.3\sqrt[p]{4}$ for any $p \in \mathbb{N}_{>0} \cup \{\infty\}$. For $y = 1/1.3$, we have

$$f_p(\frac{0.3}{1.3}) = \sqrt[p]{4(\frac{0.3}{1.3})^p + (\frac{0.3}{1.3})^p} = \sqrt[p]{5(\frac{0.3}{1.3})^p} = \frac{0.3}{1.3}\sqrt[p]{5}\ .$$

It follows that $f_p(\frac{0.3}{1.3}) < 1 = f_p(0)$ for any $p \in \mathbb{N}_{>0} \cup \{\infty\}$ with $p \geq 2$, so no such π with $y = 0$ minimises $d_p^\varepsilon(\Gamma, \pi)$ when $p \geq 2$. Therefore, when $p \geq 2$, any π minimising $d_p^\varepsilon(\Gamma, \pi)$ is such that $P_\pi(x_3) > 0$. Since $\frac{P_\pi(w_i \wedge x_3)}{P_\pi(x_3)} = 0.25$, $P(w_i|x_3) \geq 0.25 \in \mathcal{C}_p^\varepsilon(\Gamma)$ for all $w_i \in W_2$. Additionally, for $p \geq 1$, $f_p(1) > f_p(0.3/1.3)$, so no such π with $y = 1$ minimises $d_p^\varepsilon(\Gamma, \pi)$, and $P(x_2) \geq y \in \mathcal{C}_p^\varepsilon(\Gamma)$ for some $y \in (0, 1)$:

$$\mathcal{C}_p^\varepsilon(\Gamma) = \{P(w_i|x_3) \geq 0.25 | w_i \in W_2\} \cup \{P(x_3) \geq y\}\ .$$

Finally, consider the base:

$$\Gamma' = \{P(w_i|x_3) \geq 0.25 | w_i \in W_2\} \cup \{P(x_3) \geq 1\}\ .$$

Clearly, Γ is consistent, and $\mathcal{C}_p^\varepsilon(\Gamma)$ is not a maximal consolidation, since $y < 1$.

Now, for $p = 1$, consider the base:

$$\Psi = \Gamma \cup \{P(x_3 \wedge x_3) \geq 0.5\}$$

Let π be a probabilistic interpretation minimising $d_1^\varepsilon(\Psi, \pi)$ and let z denote $P_\pi(x_3)$. Such π corresponds to the following values for $\varepsilon_1, \ldots, \varepsilon_5$:

$$\begin{aligned}\varepsilon_i &= \max\{-P_\pi(w_i \wedge x_3) + 0.55P_\pi(x_3), 0\}, \text{ for } w_i \in W_2 \\ \varepsilon_i &= \max\{-P_\pi(w_i \wedge x_3) + 0.55z, 0\}, \text{ for } w_i \in W_2 \\ \varepsilon_5 &= \max\{-P_\pi(x_3) + 1, 0\} = 1 - z \\ \varepsilon_6 &= \max\{-P_\pi(x_3) + 0.5, 0\} = \max\{0.5 - z, 0\}\end{aligned}$$

As $\sum_{w_i \in W_2} P_\pi(w_i \wedge x_3) = P_\pi(x_3) = z$, $\sum_{i=1}^{4} \varepsilon_i \geq 2.2z - z = 1.2z$. To compute the 1-norm of $\varepsilon =$

$\langle \varepsilon_1, \ldots, \varepsilon_6 \rangle$, we have two cases. If $z \leq 0.5$:

$$\|\varepsilon\|_1 = \sum_{i=1}^{6} \varepsilon_i \geq 1.2z + (1-z) + (0.5-z) = 1.5 - 0.8z \ ,$$

else:

$$\|\varepsilon\|_1 = \sum_{i=1}^{6} \varepsilon_i \geq 1.2z + (1-z) = 1 + 0.2z \ .$$

Note that $\sum_{i=1}^{4} \varepsilon_i = 1.2$ iff $-P_\pi(w_i \wedge x_3) + 0.55z \geq 0$ (equivalently, when $z > 0$, $\frac{P_\pi(w_i \wedge x_3)}{P_\pi(x_3)} \leq 0.55$) for all $w_i \in W_2$. Taking, for instance, $P_\pi(w_i \wedge x_3) = 0.25z$, these minimum values are achievable for any value of z. Since we are minimizing $\|\varepsilon\|_1$, we can write:

$$\|\varepsilon\|_1 = \begin{cases} 1.5 - 0.8z & , \text{ if } z \leq 0.5 \\ 1 + 0.2z & , \text{ if } z > 0.5 \end{cases}$$

The reader can note that the minimum of $\|\varepsilon\|_1$ is at $z = 0.5$, where $\|\varepsilon\|_1 = 1.1$. As $z = P_\pi(x_3) > 0$, $\frac{P_\pi(w_i \wedge x_3)}{P_\pi(x_3)} \leq 0.55$ for all $w_i \in W_2$, and we can write:

$$\mathcal{C}_1^\varepsilon(\Psi) = \left\{ P(w_i|x_3) \geq \frac{P_\pi(w_i \wedge x_3)}{P_\pi(x_3)} | w_i \in W_2 \right\} \cup \{P(x_3) \geq 0.5, P(x_3 \wedge x_3) \geq 0.5\} \ .$$

Finally, consider the base:

$$\Psi' = \left\{ P(w_i|x_3) \geq \frac{P_\pi(w_i \wedge x_3)}{P_\pi(x_3)} | w_i \in W_2 \right\} \cup \{P(x_3) \geq 1, P(x_3 \wedge x_3) \geq 1\} \ .$$

Clearly, Ψ' is consistent, and $\mathcal{C}_p^\varepsilon(\Psi)$ is not a maximal consolidation, since $0.5 < 1$. □

Proposition 4.4.6 . *For any $p \in \mathbb{N}_{>0} \cup \{\infty\}$, $\mathcal{C}_p^\varepsilon$ satisfies (Success), (Vacuity) and (Pareto-ε-Optimality).*

Proof. By definition, $\mathcal{V}_p(\Gamma)$ is a consistent, so (Success) is trivially satisfied. Let $\varepsilon = \langle \varepsilon_1, \ldots, \varepsilon_m \rangle$ be the vector of violations inserted in $\mathcal{V}_p(\Gamma)$. As p-norms are non-negative, and $\|\varepsilon\|_p = 0$ iff $\varepsilon_i = 0$ for all $1 \leq i \leq m$, a consistent Γ yields the solution $\varepsilon = \langle 0, \ldots, 0 \rangle$ and (Vacuity) also holds.

For (Pareto-Optimatility), consider an arbitrary $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq m\}$ in $\mathbb{K}_c$. To prove by contradiction, suppose $\mathcal{V}_p(\Gamma) = \{P(\varphi_i|\psi_i) \geq_{\varepsilon_i} q_i | 1 \leq i \leq m\}$ is not a maximal ε-consolidation, for a $p \in \overline{\mathbb{N}}_{>0}$. Hence, there must be a maximal ε-consolidation $\Gamma' = \{P(\varphi_i|\psi_i) \geq_{\varepsilon'_i} q_i | 1 \leq i \leq m\}$ of Γ such that $0 \leq \varepsilon'_i \leq \varepsilon_i$ for all $1 \leq i \leq m$ and $\varepsilon'_i < \varepsilon_i$ for some $1 \leq i \leq m$; otherwise $\mathcal{V}_p(\Gamma)$ would be a maximal ε-consolidation. Now it follows that $\|\varepsilon'\|_p < \|\varepsilon\|_p$, and $\mathcal{V}_p(\Gamma)$ is not the consistent $\{P(\varphi_i|\psi_i) \geq_{\delta_i} q_i | 1 \leq i \leq m\}$ with minimum $\|\delta\|_p$, a contradiction. □

Proposition 4.5.2 . *The consolidation operator $\mathcal{C}_{CRV} : \mathbb{K}_c \to \mathbb{K}_c$ satisfies (Success), (Vacuity), (Structural Preservation) and (Non-Strengthening) but not (Pareto-Optimality).*

Proof. $\Gamma' = \mathcal{C}_{CRV}(\Gamma)$ is constructed from a probability mass π in such a way that, for any $P(\varphi_i|\psi_i) \geq q_i \in \Gamma'$, either $P_\pi(\psi_i) = 0$ or $q_i \leq \frac{P_\pi(\varphi_i \wedge \psi_i)}{P_\pi(\psi_i)}$; thus (Success) is satisfied.

By definition, $\mathcal{C}_{CRV}(\Gamma)$ is a $\Lambda_\Gamma(r)$, so (Structural Preservation) is trivially satisfied.

If Γ is consistent, there is a probabilistic interpretation π such that, for each $P(\varphi_i|\psi_i) \geq q_I \in \Gamma$, either $P_\pi(\psi_i) = 0$, or $P_\pi(\varphi_i \wedge \psi_i)/P_\pi(\psi_i) = r_i \geq q_i$. Considering $r_i = q_i$ when $P_\pi(\psi_i) = 0$), r is such that $d_{CRV}(\Lambda^P_\Gamma(r), \pi) = 0$. As $d_{CRV}(.,.)$ is always non-negative, such π and r minimises $d_{CRV}(\Lambda^P_\Gamma(r), \pi) = 0$, and (Vacuity) holds.

(Non-Strengthening) follows trivially from the definition of $\mathcal{C}_{CRV}(\Gamma)$.

To see that (Pareto-Optimality), consider the base $\Gamma = \{P(x_1|x_2) \geq 0.6, P(\neg x_1|x_2) \geq 0.6, P(x_1) \geq 0.9\}$. The base $\Lambda^P_\Gamma(r)$ that minimises $d_{CRV}(\Lambda^P_\Gamma(r), \pi)$ for some π, subject to $r \geq \langle 0.6, 0.6, 0.9\rangle$, is $r = q$. Solving the numerical optimisation, the π minimising $d_{CRV}(\Lambda^P_\Gamma(r), \pi)$ is such that $P_\pi(x_1 \wedge x_2)/P_\pi(x_2) = 0.5$, $P_\pi(x_1 \wedge x_2)/P_\pi(x_2) = 0.5$, and $P_\pi(x_2) \approxeq 0.89627$. Hence, $\mathcal{C}_{CRV}(\Gamma) = \{P(x_1|x_2) \geq 0.5, P(\neg x_1|x_2) \geq 0.5, P(x_2) \geq 0.89627\}$. However, $\Psi = \{P(x_1|x_2) \geq 0.5, P(\neg x_1|x_2) \geq 0.5, P(x_1) \geq 0.9\}$ is consistent, relaxing strictly less the probability bounds, and (Pareto-Optimality) is violated. □

A.3 Technical Results from Chapter 5

Theorem 5.1.9 . *Consider the propositional language $\mathcal{L}_{X_n}$ and the classical consequence operation Cn_{Cl}. The operator $- : \mathcal{L}_{X_n} \to 2^{\mathcal{L}}_{X_n}$ for a base $B \subseteq 2^{\mathcal{L}}_{X_n}$ satisfies (Success), (Inclusion), (Relevance) and (Uniformity) iff $-$ is a partial meet contraction.*

Proof. See Theorem 2.2 in (Hansson, 1999). □

Theorem 5.1.11 . *Let Cn be a consequence operation satisfying monotonicity and the upper bound property. The operator $- : \mathcal{L} \to 2^{\mathcal{L}}$ for a base $B \subseteq 2^{\mathcal{L}}$ satisfies (Success), (Inclusion), (Relevance) and (Uniformity) iff $-$ is a partial meet contraction.*

Proof. See the proof of Theorem 5.2.8 in (Wassermann, 2000), noting that, despite the compactness of Cn being mentioned in the Theorem, only the upper bound property is used in the proof. Note also that the proof works for a general language $\mathcal{L}$ containing $\perp$. □

Theorem 5.1.13 . *Let Cn be a consequence operation satisfying monotonicity and the upper bound property. The operator $- : \mathcal{L} \to 2^{\mathcal{L}}$ for a base $B \subseteq 2^{\mathcal{L}}$ satisfies (Success), (Inclusion), (Fullness) and (Uniformity) iff $-$ is a maxichoice contraction.*

Proof. ($\rightarrow$) As (Fullness) implies (Relevance), $B - \alpha = \bigcap \gamma(B \perp \alpha)$, for some selection function γ. It remains to prove that such γ returns a singleton. If $\alpha \in Cn(\varnothing)$, then $B \perp \alpha = \varnothing$ and, by definition, $\gamma(B \perp \alpha) = \{B\}$. If $B - \alpha = B$, again it must hold that $\gamma(B \perp \alpha) = \{B\}$, for there are no two different subsets of B whose intersection is equal to B. Hence, we consider $B - \alpha \neq B$ and, by (Inclusion), $B - \alpha \subsetneq B$. Due to (Fullness), any $\beta \in B \setminus B - \alpha$ is such that $\alpha \in Cn(B \cup \{\beta\})$. Therefore, $B - \alpha$ is a maximal subset of B that does not imply α; *i.e.* $\cap\gamma(B \perp \alpha) \in B \perp \alpha$. Finally, as the intersections of two elements in $B \perp \alpha$ is not in $B \perp \alpha$ — for it is not maximal —, $\gamma(B \perp \alpha)$ has only one element.

($\leftarrow$) As maxichoice contraction is a particular case of partial meet contraction, $B - \alpha$ satisfies (Success), (Inclusion), (Relevance) and (Uniformity), by Theorem 5.1.11. Now suppose there is a $\beta \in B \setminus (B - \alpha)$, so that $B - \alpha \neq B$. In this case, $(B \perp \alpha) \neq \varnothing$ and $\alpha \notin Cn(\varnothing)$. Due to (Success), $\alpha \notin Cn(B - \alpha)$. As $B - \alpha \in B \perp \alpha$, $B - \alpha$ is a maximal subset of B not implying α. Finally, $\alpha \in Cn(B - \alpha \cup \{\beta\})$ and $B - \alpha$ satisfies (Fullness). □

Theorem 5.1.15 . *Let Cn be a consequence operation satisfying monotonicity, the upper bound property and such that $\bot \notin Cn(\varnothing)$. $B!$ satisfies (Success), (Inclusion), (Relevance) iff $B!$ is a partial meet consolidation.*

Proof. See the proof of Theorem 5.2.12 in (Wassermann, 2000), noting that, despite the compactness of Cn being mentioned in the Theorem, only the upper bound property is used in the proof. Note also that the proof works for a general language $\mathcal{L}$ containing $\bot$. □

Corollary 5.1.16 . *Let Cn be a consequence operation satisfying monotonicity, the upper bound property and such that $\bot \notin Cn(\varnothing)$. $B!$ satisfies (Success), (Inclusion), (Fullness) iff $B!$ is a maxichoice consolidation.*

Proof. ($\rightarrow$) As (Fullness) implies (Relevance), $B - \bot = \bigcap \gamma(B\bot^{\bot})$, for some selection function γ. It remains to prove that such γ returns a singleton. If $B - \bot \cap \gamma(B\bot^{\bot}) = B$, it must hold that $\gamma(B\bot^{\bot}) = \{B\}$, for there are no two different subsets of B whose intersection is equal to B. Hence, we consider $B - \bot \neq B$ and, by (Inclusion), $B - \bot \subsetneq B$. Due to (Fullness), any $\beta \in B \setminus B - \bot$ is such that $\bot \in Cn(B \cup \{\beta\})$. Therefore, $B - \bot$ is maximal subset of B such that $\bot \notin Cn(B)$; *i.e.* $\cap\gamma(B\bot^{\bot}) \in B\bot^{\bot}$. Finally, as the intersections of two elements in $B\bot^{\bot}$ is not in $B\bot^{\bot}$ — for it is not maximal —, $\gamma(B\bot^{\bot})$ has only one element.

($\leftarrow$) As maxichoice contraction is a particular case of partial meet contraction, $B - \bot$ satisfies (Success), (Inclusion), (Relevance), by Theorem 5.1.15. Now suppose there is a $\beta \in B \setminus (B - \bot)$, so that $B - \bot \neq B$. Due to (Success), $\bot \notin Cn(B - \bot)$, since $\bot \notin Cn(\varnothing)$. As $B - \bot \in B\bot^{\bot}$, $B - \bot$ is a maximal subset of B such that $\bot \notin Cn(B)$. Finally, $\bot \in Cn(B - \bot \cup \{\beta\})$ and $B - \bot$ satisfies (Fullness). □

Proposition 5.2.2 . *Cn_{ew} is Tarskian and satisfies Cn_{Pr}-dominance.*

Proof. **Inclusion**: Consider a base $\Gamma \in \mathbb{K}$ and a conditional $\alpha \in \Gamma$. Note that $\alpha \in Cn_{ew}(\{\alpha\}) \subseteq Cn_{ew}(\Gamma)$, thus inclusion is satisfied.

Monotonicity: Consider a $\Psi \subseteq \Gamma \in \mathbb{K}$. If $\beta \in Cn_{ew}(\Psi)$, there must be an $\alpha \in \Psi$ such that $\beta \in Cn_{ew}(\{\alpha\})$. As $\alpha \in \Psi$ implies $\alpha \in \Gamma$, $Cn_{ew}(\{\alpha\}) \subseteq Cn_{ew}(\Gamma)$ and monotonicity is satisfied.

Idempotence: Consider a conditional α in a base $\Gamma \in \mathbb{K}$. For every $\gamma \in Cn_{ew}(Cn_{ew}(\Gamma))$, there is a $\beta \in Cn_{ew}(\Gamma)$ such that $\gamma \in Cn_{ew}(\{\beta\})$. Similarly, $\beta \in Cn_{ew}(\Gamma)$ implies there is some $\alpha \in \Gamma$ such that $\beta \in Cn_{ew}(\{\alpha\})$. But note that, if $\beta \in Cn_{ew}(\{\alpha\})$, then $Cn_{ew}(\{\beta\}) \subseteq Cn_{ew}(\{\alpha\})$ and $\gamma \in Cn_{ew}(\{\alpha\}) \subseteq Cn_{ew}(\Gamma)$. Therefore, $Cn_{ew}(Cn_{ew}(\Gamma)) \subseteq Cn_{ew}(\Gamma)$ and, by monotonicity, $Cn_{ew}(Cn_{ew}(\Gamma)) = Cn_{ew}(\Gamma)$.

Cn_{Pr}-dominance: Consider a base $\Gamma \in \mathbb{K}$ and a conditional $\beta = P(\varphi|\psi) \geq q$ in $Cn_{ew}(\Gamma)$. There must be an $\alpha \in \Gamma$ such that $\beta \in Cn_{ew}(\{\alpha\})$, so that $\alpha = P(\varphi|\psi) \geq q'$, for some $q' \geq q$. Note that, if some probabilistic interpretation π satisfies α, then π also satisfies β and $\beta \in Cn_{Pr}(\{\alpha\})$. As Cn_{Pr} satisfies monotonicity, $Cn_{Pr}(\{\alpha\}) \subseteq Cn_{Pr}(\Gamma)$ and $\beta \in Cn_{Pr}(\Gamma)$, finishing the proof. □

Proposition 5.2.7 . *Let $\mathcal{C} : \mathbb{K}_c \rightarrow \mathbb{K}_c$ be a maximal consolidation operator and consider $Cn^{\star} = Cn_{ew}$. For any $\Gamma \in \mathbb{K}$, the consolidation operation $!$ for Γ defined as $\Gamma! = Cn_{ew}(\Gamma\mathcal{C})$ satisfies (Relevance*).*

Proof. Note that if $\alpha \in Cn_{ew}(\Gamma) \setminus \Gamma!$, then $\alpha \notin Cn_{ew}(\Gamma\mathcal{C})$. Therefore, $\alpha = P(\varphi_i|\psi_i) \geq q_i$ is such that there are a $\beta = P(\varphi_i|\psi_i) \geq q_i'$ in $\Gamma\mathcal{C}$ and a $\gamma = P(\varphi_i|\psi_i) \geq q_i''$ in Γ such that $q_i'' \geq q_i > q_i'$. Finally, $\Gamma! \cup \{\alpha\}$ is inconsistent, for $\Gamma\mathcal{C} \cup \{\alpha\}$ also is, since $\Gamma\mathcal{C}$ is a maximal consolidation. □

Lemma 5.2.8 (Upper Bound Property – for consolidation)**.** *Let $\Gamma \in \mathbb{K}_c$ be a canonical base. For every consistent $\Delta \subseteq Cn_{ew}(\Gamma)$, there is a Δ' such that $\Delta \subseteq \Delta' \subseteq Cn_{ew}(\Gamma)$ and $\Delta' \in Cn_{ew}(\Gamma)\perp^{\perp}$.*

Proof. See the proof of Theorem 5.4.3, taking $\alpha = P(\top) < 0.5$, for instance. □

Lemma 5.2.9 . *Let $\Gamma \in \mathbb{K}_c$ be a canonical base. $\Psi \in \mathbb{K}_c$ is a maximal consolidation of Γ iff $Cn_{ew}(\Psi) \in Cn_{ew}(\Gamma)\perp^{\perp}$.*

Proof. The proof of the Proposition is broken into two parts, corresponding to the directions of the bi-implication, where $r = \langle r_1, \ldots, r_m\rangle$ and $q = \langle q_1, \ldots, q_m\rangle$ are vectors in $[0,1]^m$ and $m = |\Gamma|$:

($\rightarrow$) If Γ and $Cn_{ew}(\Gamma)$ are consistent, Γ is the only maximal consolidation of Γ, for no lower bounds need to be decreased in order to consolidate it. Thus, $Cn_{ew}(\Gamma)$ is a maximal consistent subset of $Cn_{ew}(\Gamma)$ and it is in $Cn_{ew}(\Gamma)\perp^{\perp}$. So we consider an inconsistent Γ. Let $\Psi = \Lambda_\Gamma(r)$ be a maximal consolidation of $\Gamma = \Lambda_\Gamma(q)$. As Ψ is consistent, which must be different from Γ, which is consistent. $Cn_{ew}(\Psi)$ is also consistent, for each of its conditionals is implied by some $\alpha \in \Psi$. Any $\alpha \in Cn_{ew}(\Gamma) \setminus Cn_{ew}(\Psi)$ is such that $\alpha = P(\varphi_i|\psi_i) \geq r'_i$ and $r_i < r'_i \leq q_i$ for some $1 \leq i \leq m$. As $\Psi = \Lambda_\Gamma(r)$ is a maximal consolidation of Γ, $\Psi' = \Lambda_\Gamma(\langle r_1, \ldots, r'_i, \ldots, r_m\rangle)$ is inconsistent and so is $Cn_{ew}(\Psi) \cup \{\alpha\} \supset \Psi'$. Hence, $Cn_{ew}(\Psi)$ is a maximal consistent subset of $Cn_{ew}(\Gamma)$ and $Cn_{ew}(\Psi) \in Cn_{ew}(\Gamma)\perp^{\perp}$.

($\leftarrow$) Consider a $\Psi \in \mathbb{K}_c$ such that $Cn_{ew}(\Psi) \in Cn_{ew}(\Gamma)\perp^{\perp}$. Consider $\Gamma = \Lambda_\Gamma(q)$. Since $Cn_{ew}(\Psi) \subseteq Cn_{ew}(\Gamma)$, $\Psi = \Lambda_\Gamma(r)$ for some $r \leq q$. To prove by contradiction, suppose Ψ is not a maximal consolidation of Γ. As Ψ is consistent — for $Cn_{ew}(\Psi) \supset \Psi$ also is —, there must be a $r' \in [0,1]^m$ such that $r < r' \leq q$ and $\Psi' = \Lambda_\Gamma(r')$ is consistent. Furthermore, $Cn_{ew}(\Psi')$ is consistent, and $Cn_{ew}(\Psi) \subsetneq Cn_{ew}(\Psi') \subseteq Cn_{ew}(\Gamma)$. Hence, $Cn_{ew}(\Psi)$ is not a maximal consistent subset of $Cn_{ew}(\Gamma)$, a contradiction. □

Theorem 5.2.10 . *Consider a base $\Gamma \in \mathbb{K}_c$. An operation $\Gamma!$ satisfies (Success), (Inclusion*) and (Relevance*), for $Cn^\star = Cn_{ew}$, iff $\Gamma! = \bigcap \gamma(Cn_{ew}(\Gamma)\perp^{\perp})$, for some selection function γ.*

Proof. The proof is adapated from the proof of Theorem 8 in (Santos *et al.*, 2015), just taking $\alpha = \perp$. From the construction to the postulates, making $\alpha = \perp$ in the mentioned proof already yield our result, just ignoring (Uniformity). From the postulates to the construction, Santos *et al.* (2015) define a selection function γ and split the proof into 3 parts. As the first part proves that γ is indeed a function, it is not needed in our case, for $\perp$ is the only argument and the result is trivial. The remaining parts do not use (Uniformity) and hold for our case. The second part aims to prove that γ is a selection function, and it follows from the upper bound property, which we proved for probabilistic consolidation in Lemma 5.4.3. The third part proves that $\Gamma! = \Gamma - \perp = \cap\gamma(Cn^\star(\Gamma)\perp^{\perp})$, using again the upper bound result. □

Corollary 5.2.11 . *Consider the consequence operation $Cn_{ew} = Cn^\star$, a maximal consolidation operator $\mathcal{C}$ and a consolidation operation ! for each $\Gamma \in \mathbb{K}_c$ defined as $\Gamma! = Cn_{ew} \circ \mathcal{C}(\Gamma)$. For each canonical base $\Gamma \in \mathbb{K}_c$, $!(\Gamma)$ satisfies (Success), (Inclusion*) and (Relevance*).*

Proof. By Lemma 5.2.9, $\mathcal{C}(\Gamma)$ is such that $Cn_{ew}(\mathcal{C}(\Gamma)) \in Cn_{ew}(\Gamma)\perp^{\perp}$, and the result follows by Theorem 5.2.10. □

Corollary 5.2.13 . *Consider the consequence operation $Cn_{ew} = Cn^\star$. The consolidation operation $\Gamma!$, for any $\Gamma \in \mathbb{K}_c$, satisfies (Success), (Inclusion*) and (Fullness*) iff there is a maximal consolidation operator $\mathcal{C}$ such that $\Gamma! = Cn_{ew} \circ \mathcal{C}(\Gamma)$ for all $\Gamma \in \mathbb{K}_c$.*

Proof. ($\leftarrow$) From Corollary 5.2.11, it just remains to prove that $!\Gamma$ satisfies (Fullness*) for all $\Gamma \in \mathbb{K}_c$. Consider an arbitrary $\Gamma \in \mathbb{K}_c$. By Lemma 5.2.9, $\Gamma! \in Cn_{ew}(\Gamma)\perp^\perp$, for $\Gamma! = Cn_{ew}(\mathcal{C}(\Gamma))$ and $\mathcal{C}(\Gamma)$ is a maximal consolidation of Γ. If $\beta \in Cn_{ew}(\Gamma) \setminus \Gamma!$, then $\Gamma! \cup \{\beta\}$ is inconsistent, for $\Gamma!$ is maximally consistent. Moreover, as $\Gamma!$ is consistent — like $\mathcal{C}(\Gamma)$ —, (Fullness*) is satisfied.

($\rightarrow$) From Theorem 5.2.10, as (Fullness*) implies (Relevance*), we have that $\Gamma! = \cap\gamma(Cn_{ew}(\Gamma)\perp^\perp)$ for some selection function γ. We shall prove that such γ selects a single member of $Cn_{ew}(\Gamma)\perp^\perp$, and the result follows from Lemma 5.2.9 (Note that $Cn_{ew}(\Gamma)\perp^\perp$ is never empty, for $\Lambda_\Gamma(\langle 0, \ldots, 0\rangle)$ is always consistent). Suppose, to prove by contradiction, that $\gamma(Cn_{ew}(\Gamma)\perp^\perp)$ contains two distinct elements $Cn_{ew}(\Psi)$ and $Cn_{ew}(\Psi')$. It follows that $\cap\gamma(Cn_{ew}(\Gamma)\perp^\perp) \subsetneq Cn_{ew}(\Psi)$ and there is a conditional $\beta \in Cn_{ew}(\Psi) \setminus \Gamma!$ such that $\Gamma! \cup \{\beta\} \cup Cn_{ew}(\Psi)$ is consistent — for $Cn_{ew}(\Psi)$ is. Nevertheless, this contradicts (Fullness*). Consequently, $\gamma(Cn_{ew}(\Gamma)\perp^\perp)$ contains only one element, which, by Lemma 5.2.9, has the form $Cn_{ew}(\Delta)$, where Δ is a maximal consolidation of Γ. Finally, we can define a function $\mathcal{C} : \mathbb{K}_c \rightarrow \mathbb{K}_c$ such that $\Gamma! = Cn_{ew}(\mathcal{C}(\Gamma))$ for all $\Gamma \in \mathbb{K}_c$. □

Corollary 5.2.15 . *Consider a base $\Gamma = Cn_{ew}(\Gamma)$ in $\mathbb{K}$. The consolidation operation $\Gamma!$, for any $\Gamma = Cn_{ew}(\Gamma)$ in $\mathbb{K}$, satisfies (Success), (Inclusion) and (Fullness) iff there is a maximal consolidation operator $\mathcal{C}$ such that $\Gamma! = Cn_{ew} \circ \mathcal{C} \circ Cn_{ew}^{-1}(\Gamma)$ for all $\Gamma = Cn_{ew}(\Gamma)$ in $\mathbb{K}$.*

Proof. Note that there is a bijection between $\mathbb{K}_c$ and the set $S = \{\Gamma \in \mathbb{K} | \Gamma = Cn_{ew}(\Gamma)\}$, given by the functions Cn_{ew} and Cn_{ew}^{-1}. For each $\Gamma \in S$, let $\Psi \in \mathbb{K}_c$ be equal to $Cn_{ew}^{-1}(\Gamma)$, so that $Cn_{ew} \circ \mathcal{C} \circ Cn_{ew}^{-1}(\Gamma) = Cn_{ew} \circ \mathcal{C}(\Psi)$ for all $\Gamma \in S$ for some maximal consolidation operator $\mathcal{C}$. By Corollary 5.2.13, a consolidation operation $\Psi!' = \Delta$ satisfies (Success), (Inclusion*) and (Fullness*) for any $\Psi \in \mathbb{K}_c$ iff there is some maximal consolidation operator $\mathcal{C}$ such that $\Delta = Cn_{ew} \circ \mathcal{C}(\Psi)$ (equivalently, $\Delta = Cn_{ew} \circ \mathcal{C} \circ Cn_{ew}^{-1}(\Gamma)$). We need to prove that $\Psi!' = \Delta$ satisfies (Success), (Inclusion*) and (Fullness*) iff $\Gamma! = \Delta$ satisfies (Success), (Inclusion) and (Fullness) for any pair $\langle \Psi, \Gamma\rangle \in \mathbb{K}_c \times S$ such that $\Gamma = Cn_{ew}(\Psi)$. (Success) is trivial, for it depends only on Δ. As $\Gamma = Cn_{ew}(\Psi)$, $\Delta \subseteq \Gamma$ iff $\Delta \subseteq Cn_{ew}(\Psi)$; and $\Gamma! = \Delta$ satisfies (Inclusion) iff $\Psi!' = \Delta$ satisfies (Inclusion*). Finally, note that (Fullness) for $\Gamma! = \Delta$ and (Fullness*) for $\Psi!' = \Delta$ are captured by the same proposition: if $\beta \in Cn_{ew}(\Psi)$ and $\beta \notin \Delta$, then Δ is consistent, and $\Delta \cup \{\beta\}$ is not. We can conclude that $\Gamma! = \Delta$ satisfies (Fullness) iff $\Psi!' = \Delta$ satisfies (Fullness*), for each $\Psi \in \mathbb{K}_c$, finishing the proof. □

Proposition 5.3.2 . *If Cn satisfies monotonicity and idempotency, and $Cn^\star$, inclusion and Cn-dominance, (Success) is equivalent to (*Lifted Success).*

Proof. ($\rightarrow$) Suppose a contraction operator $-$ satisfies (Success); that is, for any $B \subseteq \mathcal{L}$ and $\alpha \in \mathcal{L} \setminus Cn(\varnothing)$, $\alpha \notin Cn(B - \alpha)$. For $Cn^\star$ satisfies Cn-dominance, $Cn^\star(B - \alpha) \subseteq Cn(B - \alpha)$, and the monotonicity of Cn implies $Cn(Cn^\star(B-\alpha)) \subseteq Cn(Cn(B-\alpha))$. Since Cn satisfies idempotence, $Cn(B - \alpha) = Cn(Cn(B - \alpha))$, and it follows that $Cn(Cn^\star(B - \alpha)) \subseteq Cn(B - \alpha)$. Hence, $\alpha \notin Cn(B - \alpha)$ implies $\alpha \notin Cn(Cn^\star(B - \alpha))$, and (*Lifted Success) is satisfied.

($\leftarrow$) Now suppose $-$ satisfies (*Lifted Success); that is, for any $B \subseteq \mathcal{L}$ and $\alpha \in \mathcal{L} \setminus Cn(\varnothing)$, $\alpha \notin Cn(Cn^\star(B - \alpha))$. By inclusion, $B \subseteq Cn^\star(B)$. As Cn is monotonic, $Cn(B) \subseteq Cn(Cn^\star(B))$. Finally, $\alpha \notin Cn(Cn^\star(B - \alpha))$ implies $\alpha \notin Cn(B - \alpha)$, and (Success) is satisfied. □

Proposition 5.3.6 . *Consider a consequence operation Cn that satisfies monotonicity and idempotence and a Tarskian consequence operation $Cn^\star$ satisfying Cn-dominance. The postulate of (Inclusion*) is equivalent to (*Lifted Inclusion). The postulate of (Relevance*) implies (*Lifted Relevance). The postulate of (Uniformity*) implies (*Lifted Uniformity). The postulate of (Fullness*) implies (*Lifted Fullness).*

Proof. Consider a base $B \subseteq \mathcal{L}$, a formula $\alpha \in \mathcal{L}$, and the contraction operators $-$ for B and $\overset{\star}{-}$ for $Cn^\star(B)$ such that $Cn^\star(B) \overset{\star}{-} \alpha = Cn^\star(B - \alpha)$.

Inclusion: If $B - \alpha \subseteq Cn^\star(B)$, then $Cn^\star(B - \alpha) \subseteq Cn^\star(B)$, by the monotonicity and the idempotence of $Cn^\star$, and (Inclusion*) implies (*Lifted Inclusion). If $Cn^\star(B-\alpha) \subseteq Cn^\star(B)$, then $B-\alpha \subseteq Cn^\star(B)$, for $Cn^\star$ satisfies inclusion ($B \subseteq Cn^\star(B)$), and (*Lifted Inclusion) implies (Inclusion*).

Relevance: Suppose $-$ for B satisfies (Relevance*). Since any β in $Cn^\star(B) \setminus Cn^\star(B-\alpha)$ is also in $Cn^\star(B) \setminus B-\alpha$, for $Cn^\star$ satisfies inclusion, there is, for every such β, a B' such that $B-\alpha \subseteq B' \subseteq Cn^\star(B)$, $\alpha \notin Cn(B')$ and $\alpha \in Cn(B' \cup \{\beta\})$. As $Cn^\star$ satisfies Cn-dominance, and Cn is monotonic, $Cn(Cn^\star(B')) \subseteq Cn(Cn(B'))$, and as Cn is idempotent, $Cn(Cn^\star(B')) \subseteq Cn(B')$. Therefore, $\alpha \notin Cn(Cn^\star(B'))$. As $Cn^\star$ satisfies inclusion, $B' \subseteq Cn^\star(B')$ and $B' \cup \{\beta\} \subseteq Cn^\star(B') \cup \{\beta\}$. Since Cn is monotonic, $Cn(B' \cup \{\beta\}) \subseteq Cn(Cn^\star(B') \cup \{\beta\})$. Consequently, $\alpha \in Cn(Cn^\star(B') \cup \{\beta\})$. Since $Cn^\star$ is monotonic and idempotent, $Cn^\star(B-\alpha) \subseteq Cn^\star(B') \subseteq Cn^\star(B)$ and $-$ for B satisfies (*Lifted Relevance).

Uniformity: Suppose $-$ for B satisfies (Uniformity*). Now suppose that, for any $B' \subseteq Cn^\star(B)$, it is the case that $\alpha \in Cn(B')$ iff $\beta \in Cn(B')$. By (Uniformity*), $B - \alpha = B - \beta$, so $Cn^\star(B - \alpha) = Cn^\star(B - \beta)$, and (*Lifted Uniformity) holds.

Fullness: Suppose $-$ for B satisfies (Fullness*). Since any β in $Cn^\star(B) \setminus Cn^\star(B - \alpha)$ is also in $Cn^\star(B) \setminus B-\alpha$, for $Cn^\star$ satisfies inclusion, it holds that $\alpha \notin Cn(B-\alpha)$ and $\alpha \in Cn(B-\alpha \cup \{\beta\})$ for any such β. As $Cn^\star$ satisfies Cn-dominance and Cn is monotonic, $Cn(Cn^\star(B - \alpha)) \subseteq Cn(Cn(B - \alpha))$, and as Cn is idempotent, $Cn(Cn^\star(B - \alpha)) \subseteq Cn(B - \alpha)$. Therefore, $\alpha \notin Cn(Cn^\star(B - \alpha))$. As $Cn^\star$ is satisfies inclusion, $B - \alpha \subseteq Cn^\star(B - \alpha)$ and $B - \alpha \cup \{\beta\} \subseteq Cn^\star(B - \alpha) \cup \{\beta\}$. Since Cn is monotonic, $Cn(B - \alpha \cup \{\beta\}) \subseteq Cn(Cn^\star(B - \alpha) \cup \{\beta\})$. Consequently, $\alpha \in Cn(Cn^\star(B - \alpha) \cup \{\beta\})$ and $-$ for B satisfies (*Lifted Fullness). □

Proposition 5.3.7 . *Consider a consequence operation Cn that satisfies monotonicity and idempotence and a Tarskian consequence operation $Cn^\star$ satisfying Cn-dominance. If (Enforced Closure*) is satisfied, then (*Lifted Relevance) is equivalent to (Relevance*), (*Lifted Uniformity) is equivalent to (Uniformity*) and (*Lifted Fullness) is equivalent to (Fullness*).*

Proof. Consider a base $B \subseteq \mathcal{L}$, a formula $\alpha \in \mathcal{L}$, and the contraction operators $-$ for B and $\overset{\star}{-}$ for $Cn^\star(B)$ such that $Cn^\star(B) \overset{\star}{-} \alpha = Cn^\star(B - \alpha)$. Note that the ($\leftarrow$)-direction of each equivalence is already proven in Proposition 5.3.6.

Relevance: Due to (Enforced Closure*), $B - \alpha = Cn^\star(B) \overset{\star}{-} \alpha = \Delta$ and both (*Lifted Relevance) and (Relevance*) are captured by the same proposition: if $\beta \in Cn^\star(B) \setminus \Delta$, then there is a B' such that $\Delta \subseteq B' \subseteq Cn^\star(B)$ such that $\alpha \notin Cn(B')$ and $\alpha \in Cn(B' \cup \{\beta\})$.

Uniformity: Suppose $-$ for B satisfies (*Lifted Uniformity). Now suppose that, for any $B' \subseteq Cn^\star(B)$, it is the case that $\alpha \in Cn(B')$ iff $\beta \in Cn(B')$. By (*Lifted Uniformity), $Cn^\star(B) \overset{\star}{-} \alpha = Cn^\star(B) \overset{\star}{-} -\beta$ and $Cn^\star(B - \alpha) = Cn^\star(B - \beta)$. Hence, by (Enforced Closure*), $B - \alpha = B - \beta$ and (Uniformity*) holds.

Fullness: Due to (Enforced Closure*), $B - \alpha = Cn^\star(B) \overset{\star}{-} \alpha = \Delta$ and both (*Lifted Fullness) and (Fullness*) are captured by the same proposition: if $\beta \in Cn^\star(B) \setminus \Delta$, then $\alpha \notin Cn(\Delta)$ and $\alpha \in Cn(\Delta \cup \{\beta\})$. □

Proposition 5.3.9 . *Consider a monotonic consequence operation Cn that satisfies the upper bound property and a Tarskian consequence operation $Cn^\star$ satisfying Cn-dominance. A contraction operator $- : \mathcal{L} \to 2^{\mathcal{L}}$ for a base $B \subseteq 2^{\mathcal{L}}$ satisfies the *lifted versions of the postulates of success, inclusion, relevance (fullness) and uniformity iff it is a liftable partial meet contraction (liftable maxichoice contraction). Furthermore, $-$ additionally satisfies (Enforced Closure*) iff $B - \alpha = \bigcap \gamma(Cn^\star(B) \bot \alpha)$ (and $\gamma(Cn^\star(B) \bot \alpha)$ is a singleton) for all $\alpha \in \mathcal{L}$ and for some selection function γ.*

Proof. A contraction $-$ for B satisfies the *lifted versions of the postulates of success, inclusion, relevance (fullness) and uniformity iff its closure, $\overset{\star}{-}$ for $Cn^\star(B)$ with $Cn^\star(B) \overset{\star}{-} \alpha = Cn^\star(B - \alpha)$, satisfies the non-lifted versions of the postulates. With relevance (or fullness), the latter statement is the case iff $\overset{\star}{-}$ is a partial meet contraction (maxichoice contraction), by Theorem 5.1.11 (or Theorem 5.1.13), which is itself the definition of $-$ as a liftable partial meet contraction (liftable maxichoice contraction). If $-$ additionally satisfies (Enforced Closure*), then $B - \alpha = Cn^\star(B - \alpha) = Cn^\star(B) \overset{\star}{-} \alpha$; now, if $\overset{\star}{-}$ is a partial meet contraction, $B - \alpha = Cn^\star(B) \overset{\star}{-} \alpha = \bigcap \gamma(Cn^\star(B) \bot \alpha)$, if $\overset{\star}{-}$ is a maxichoice contraction, $\gamma(Cn^\star(B) \bot \alpha)$ is a singleton. As $\bigcap \gamma(Cn^\star(B) \bot \alpha)$ is closed under $Cn^\star$ (see the proof of Proposition 5.3.12, $\leftarrow$-part), $B - \alpha = Cn^\star(B - \alpha)$ and (Enforced Closure*) is satisfied (whether or not γ returns a singleton). □

Proposition 5.3.10 . *Consider a monotonic consequence operation Cn that satisfies the upper bound property and a Tarskian consequence operation $Cn^\star$ satisfying Cn-dominance. A consolidation operation $B!$ satisfies the *lifted versions of the postulates of success, inclusion, and relevance (fullness) iff it is a liftable partial meet consolidation (liftable maxichoice consolidation). Furthermore, if $B!$ satisfies (Enforced Closure*), then $B! = \bigcap \gamma(Cn^\star(B) \bot^\bot)$ for some selection function γ ($B! \in (Cn^\star(B) \bot^\bot)$).*

Proof. A consolidation $B!$ satisfies the *lifted versions of the postulates of success, inclusion, and relevance (fullness) iff its closure $!^\star$, such that $Cn^\star(B)!^\star = Cn^\star(B!)$ satisfies the non-lifted versions of the postulates. With relevance (fullness), the latter statement is the case iff $!\star$ is a partial meet consolidation (partial maxichoice consolidation), by Theorem 5.1.15 (or 5.1.16), which is itself the definition of $-$ as a liftable partial meet consolidation (liftable maxichoice consolidation). If $B!$ additionally satisfies (Enforced Closure*), then $B! = Cn^\star(B!)$; now, if $Cn^\star(B) \overset{\star}{-} \bot = Cn^\star(B!)$ is a partial meet consolidation, $B! = \bigcap \gamma(Cn^\star(B) \bot^\bot)$, if $Cn^\star(B) \overset{\star}{-} \bot$ is a maxichoice contraction, $\gamma(Cn^\star(B) \bot^\bot)$ is a singleton. As $\bigcap \gamma(Cn^\star(B) \bot^\bot)$ is closed under $Cn^\star$ (see the proof of Proposition 5.3.12, $\leftarrow$-part), $B - \bot = Cn^\star(B - \bot)$ and (Enforced Closure*) is satisfied (whether or not γ returns a singleton). □

Proposition 5.3.12 . *Consider a consequence operation Cn that satisfies monotonicity, idempotence and the upper bound property and a Tarskian consequence operation $Cn^\star$ satisfying Cn-dominance. An operator $- : \mathcal{L} \to 2^{\mathcal{L}}$ for a base B is a liftable partial meet contraction (liftable maxichoice contraction) and satisfies ($f_\star$-Canonicity) iff, for any formula $\alpha \in \mathcal{L}$, $B - \alpha = f_\star(\bigcap \gamma(Cn^\star(B) \bot \alpha))$ for some selection function γ (and $\gamma(Cn^\star(B) \bot \alpha)$ is a singleton).*

Proof. ($\rightarrow$) The operator $-$ is a liftable partial meet contraction (or liftable maxichoice contraction) for B, iff $\overset{\star}{-}$, such that $Cn^\star(B) \overset{\star}{-} \alpha = Cn^\star(B - \alpha)$, is a partial meet contraction (or maxichoice contraction) for $Cn^\star(B)$: for every $\alpha \in \mathcal{L}$, $Cn^\star(B - \alpha) = \bigcap\gamma(Cn^\star(B)\bot\alpha)$ for some selection function γ (and $\gamma(Cn^\star(B)\bot\alpha)$ is a singleton). If $-$ additionally satisfies ($f_\star$-Canonicity), then $B-\alpha = f_\star(B-\alpha)$). As, by the idempotence of $Cn^\star$, $Cn^\star(B-\alpha) = Cn^\star(Cn^\star(B-\alpha))$, we have that $f_\star(B-\alpha) = f_\star\circ Cn^\star(B-\alpha)$, due to the fact that $f_\star$ is a canonical form, and $B-\alpha = f_\star\circ Cn^\star(B-\alpha)$. Hence, $B - \alpha = f_\star(\bigcap\gamma(Cn^\star(B)\bot\alpha))$ (and $\gamma(Cn^\star(B)\bot\alpha)$ is a singleton).

($\leftarrow$) If $B - \alpha = f_\star(\bigcap\gamma(Cn^\star(B)\bot\alpha))$ (and $\gamma(Cn^\star(B)\bot\alpha)$ is a singleton), then, as $f_\star$ is idempotent, $f_\star(B-\alpha) = B-\alpha$ and ($f_\star$-Canonicity) is satisfied. Moreover, $Cn^\star(B-\alpha) = Cn^\star(\bigcap\gamma(Cn^\star(B)\bot\alpha))$ (and $\gamma(Cn^\star(B)\bot\alpha)$ is a singleton). It remains to prove that $Cn^\star(\bigcap\gamma(Cn^\star(B)\bot\alpha)) = \bigcap\gamma(Cn^\star(B)\bot\alpha)$. Since $Cn^\star$ satisfies inclusion, $Cn^\star(\Delta) \supseteq \Delta$ for any Δ, so we just prove that $Cn^\star(\bigcap\gamma(Cn^\star(B)\bot\alpha)) \subseteq \bigcap\gamma(Cn^\star(B)\bot\alpha)$. We firstly prove that each $\Psi \in Cn^\star(B)\bot\alpha$ is such that $Cn^\star(\Psi) \subseteq \Psi$. To prove by contradiction, suppose that $Cn^\star(\Psi) \nsubseteq \Psi$. As $\Psi \subseteq Cn^\star(B)$ and $Cn^\star$ is monotonic and idempotent, $Cn^\star(\Psi) \subseteq Cn^\star(B)$. As $\alpha \notin Cn(\Psi)$, then $\alpha \notin Cn(Cn^\star(\Psi))$, as $Cn(\Psi) = Cn(Cn^\star(\Psi))$ for any idempotent, monotonic Cn and any $Cn^\star$ satisfying Cn-dominance[2]. Hence, $\Psi \subsetneq Cn^\star(\Psi)$ is not a maximal subset Δ of $Cn^\star(B)$ such that $\alpha \notin Cn(\Delta)$, which is a contradiction. To prove that $Cn^\star(\bigcap\gamma(Cn^\star(B)\bot\alpha)) = \bigcap\gamma(Cn^\star(B)\bot\alpha)$, we just note that each element of $\gamma(Cn^\star(B)\bot\alpha)$ is closed under $Cn^\star$, and so must be their intersection, for $Cn^\star$ is monotonic — see the proof of Observation 1.25 in (Hansson, 1999). Finally, $Cn^\star(B - \alpha) = Cn^\star(\bigcap\gamma(Cn^\star(B)\bot\alpha)) = \bigcap\gamma(Cn^\star(B)\bot\alpha)$ for any $\alpha \in \mathcal{L}$ and $\overset{\star}{-}$ for $Cn^\star(B)$ is a partial meet contraction (maxichoice contraction), and $-$ for B is a liftable partial meet contraction (liftable maxichoice contraction). □

Corollary 5.3.13 . *Consider a consequence operation Cn that satisfies monotonicity, idempotence and the upper bound property and a Tarskian consequence operation $Cn^\star$ satisfying Cn-dominance. An operation $B!$ is a liftable partial meet consolidation (liftable maxichoice consolidation) and satisfies ($f_\star$-Canonicity) iff, $B! = f_\star(\bigcap\gamma(Cn^\star(B)\bot^\bot))$ for some selection function γ (and $\gamma(Cn^\star(B)\bot^\bot)$ is a singleton).*

Proof. ($\rightarrow$) If $B!$ is a liftable partial meet consolidation (liftable maxichoice consolidation), then $Cn^\star(B)!^\star = Cn^\star(B!)$ is a partial meet consolidation (maxichoice consolidation): $Cn^\star(B - \bot) = \bigcap\gamma(Cn^\star(B)\bot^\bot)$ for some selection function γ (and $\gamma(Cn^\star(B)\bot^\bot)$ is a singleton). The rest of the proof follows that from Proposition 5.3.12, just assuming $\alpha = \bot$.

($\leftarrow$) See the ($\leftarrow$)-part of the proof of Proposition 5.3.12, taking $\alpha = \bot$. □

Lemma 5.3.14 . *Consider $Cn^\star = Cn_{ew}$ and $f_\star = Cn_{ew}^{-1}$. The operation $\Gamma!$ satisfies (*Lifted Success), (*Lifted Inclusion), (*Lifted Fullness) and ($f_\star$-Canonicity) for all $\Gamma \in \mathbb{K}_c$ iff $\mathcal{C} : \mathbb{K}_c \rightarrow \mathbb{K}_c$, with $\mathcal{C}(\Gamma) = \Gamma!$ for all $\Gamma \in \mathcal{L}$, is a maximal consolidation operator.*

Proof. ($\rightarrow$) By Propositions 5.3.10 and 5.3.13, if $\Gamma!$ satisfies (*Lifted Success), (*Lifted Inclusion), (*Lifted Fullness) and ($f_\star$-Canonicity), then $\Gamma! = Cn_{ew}^{-1}(\Psi)$, for some $\Psi \in Cn_{ew}(\Gamma)\bot^\bot$. As any $\Psi \in Cn_{ew}(\Gamma)\bot^\bot$ is closed under Cn_{ew}, it can be written as $\Psi = Cn_{ew}(\Delta)$, for a canonical $\Delta \in \mathbb{K}_c^\varepsilon$ (see the proof of Proposition 5.3.12, $\leftarrow$-part). By Lemma 5.2.9, Δ is a maximal consolidation of Γ. By the idempotence of Cn_{ew}, $Cn_{ew}(\Psi) = Cn_{ew}(\Delta)$. As Cn_{ew}^{-1} is a canonical form, $Cn_{ew}^{-1}(\Psi) = Cn_{ew}^{-1}(\Delta)$, and as $\Delta \in \mathbb{K}_c$, $Cn_{ew}^{-1}(\Delta) = \Delta = Cn_{ew}^{-1}(\Psi)$. Hence, $\Gamma! = \Delta$.

[2]See the proof of Proposition 5.3.6.

($\leftarrow$) By Lemma 5.2.9, $Cn_{ew}(\Gamma!) \in Cn_{ew}(\Gamma)\perp^{\perp}$ for all $\Gamma \in \mathbb{K}_c$, since $\mathcal{C}(\Gamma) = \Gamma!$ is a maximal consolidation operator. Therefore, for any $\Gamma \in \mathbb{K}_c$, $Cn_{ew}(\Gamma)!^{\star} = Cn_{ew}(\Gamma!)$ is a maxichoice consolidation, $\Gamma!$ is a liftable maxichoice consolidation and, by Proposition 5.3.10, $\Gamma!$ satisfies (*Lifted Success), (*Lifted Inclusion), (*Lifted Fullness). As $\Gamma! = \mathcal{C}(\Gamma)$ is a maximal consolidation for any $\Gamma \in \mathbb{K}_c$, it is in $\mathbb{K}_c$, and ($f_{\star}$-Canonicity) is satisfied. □

Lemma 5.3.16 . *Consider $Cn^{\star} = Cn_{ew}$, $f_{\star} = Cn_{ew}^{-1}$, a $\Gamma \in \mathbb{K}_c$. $\Gamma!$ satisfies (*Lifted Success), (*Lifted Inclusion), (*Lifted Relevance) and ($f_{\star}$-Canonicity) iff $\Gamma! = \inf M$, for a set M of maximal consolidations of Γ.*

Proof. ($\rightarrow$) By Propositions 5.3.10 and 5.3.13, if $\Gamma!$ satisfies (*Lifted Success), (*Lifted Inclusion), (*Lifted Relevance) and ($f_{\star}$-Canonicity), then $\Gamma! = Cn_{ew}^{-1}(\bigcap \gamma(Cn_{ew}(\Gamma)\perp^{\perp})$. In the proof of Theorem 5.4.5 (making $\alpha = \perp$), it is shown that $Cn_{ew}^{-1}(\bigcap \gamma(Cn_{ew}(\Gamma)\perp^{\perp}) = \inf M$, where M is a set of bases Δ such that $Cn_{ew}(\Delta) \in (Cn_{ew}(\Gamma)\perp^{\perp})$. By Proposition 5.2.9, M is a set of maximal consolidations.

($\leftarrow$) As M is a set of maximal consolidations of Γ, M is a set of bases Δ such that $Cn_{ew}(\Delta) \in Cn_{ew}(\Gamma)\perp^{\perp}$, by Lemma 5.2.9. Therefore, $\Gamma! = Cn_{ew}^{-1}(\bigcap \gamma(Cn_{ew}(\Gamma)\perp^{\perp}))$, for some selection function γ — see the proof of Theorem 5.4.5, with $\alpha = \perp$. Hence, $Cn_{ew}(\Gamma!) = Cn_{ew}(\bigcap \gamma(Cn_{ew}(\Gamma)\perp^{\perp}))$ and $Cn_{ew}(\Gamma!) = (\bigcap \gamma(Cn_{ew}(\Gamma)\perp^{\perp})$ (see the proof of Proposition 5.3.12, $\leftarrow$-part) — a partial meet consolidation. Therefore $\Gamma!$ is a liftable partial meet consolidation and, by Proposition 5.3.10, $\Gamma!$ satisfies (*Success), (*Lifted Inclusion), (*Lifted Relevance). As $\Gamma!$ has only one minimum lower bound for each $\varphi_{|}\psi$, $\Gamma \in \mathbb{K}_c$, and ($f_{\star}$-Canonicity) is satisfied. □

Proposition 5.4.1 . *Consider the base $\Gamma \in \mathbb{K}$ such that $\Gamma = \{P(\varphi) \geq q | q \in [0,1]$ and $q < 1\}$, for a non-tautological, consistent φ. $P(\varphi) \geq 1 \in Cn_{Pr}(\Gamma)$ but $P(\varphi) \geq 1 \notin Cn_{Pr}(\Psi)$ for every finite $\Psi \subseteq \Gamma$.*

Proof. Note that any probabilistic interpretation π such that $P_{\pi}(\varphi) = q < 1$ does not satisfy Γ, for there would be a $q' \in (q,1)$ such that $P(\varphi) \geq q' \in \Gamma$ is not satisfied by π. Nonetheless, any finite Ψ has a canonical form $Cn_{ew}^{-1}(\Psi) = \{P(\varphi) \geq q'\}$, where $q' < 1$ is the maximum lower bound in Ψ, such that there is a π with $P_{\pi}(\varphi) = q' < 1$ that satisfies Ψ without satisfying $P(\varphi) \geq 1$. □

Proposition 5.4.2 . *Let $\Gamma = Cn_{ew}\{P(\varphi) \geq q\}$, for some $q \in (0,1]$ and a consistent φ. There is no $\Psi \subseteq \Gamma$ such that $P(\varphi) \geq q \notin Cn_{Pr}(\Psi)$ and $P(\varphi) \geq q \in Cn_{Pr}(\Psi')$ for any Ψ' such that $\Psi \subsetneq \Psi' \subseteq \Gamma$.*

Proof. To prove by contradiction, suppose there is a $\Psi \subseteq \Gamma$ such that $P(\varphi) \geq q \notin Cn_{Pr}(\Psi)$ and $P(\varphi) \geq q \in Cn_{Pr}(\Psi')$ for any Ψ' such that $\Psi \subsetneq \Psi' \subseteq \Gamma$. Define $\alpha = P(\varphi) \geq q$. If, for every $q' < q$, $P(\varphi) \geq q' \in \Psi$, then $\alpha \in Cn_{Pr}(\Psi)$ (see the proof of Proposition 5.4.1), a contradiction. Therefore, there is a $q^{\star} < q$ such that $P(\varphi) \geq q^{\star} \in \Gamma \setminus \Psi$. Let Ψ' be $\Psi \cup \{P(\varphi) \geq q^{\star}\}$. If $q^{\star} > r$ for every r such that $P(\varphi) \geq r \in \Psi$, then a probability mass π such that $P_{\pi}(\varphi) = q^{\star}$ satisfies Ψ' but not $P(\varphi) \geq q$, a contradiction. Hence, there must be some $P(\varphi) \geq r^{\star} \in \Psi$ such that $r^{\star} \geq q^{\star}$. Thus, any probabilistic interpretation π that satisfies ψ also satisfies Ψ'. As $\alpha \notin Cn_{Pr}(\Psi)$, there is a probabilistic interpretation π satisfying Ψ that does not satisfy α. Hence, such π also satisfies Ψ' without satisfying α. Finally, $\alpha \notin Cn_{Pr}(\Psi')$, a contradiction, finishing the proof. □

Theorem 5.4.3 (Upper Bound Property)**.** *Let $\Gamma = Cn_{ew}(\Gamma)$ be a probabilistic knowledge base in $\mathbb{K}$, and $\alpha \in \bar{\mathcal{L}}_{X_n}^{P}$, a negated conditional. For every $\Delta \subseteq \Gamma$ such that $\alpha \notin Cn_{Pr}(\Delta)$, there is a $\Psi \in (\Gamma\perp\alpha)$ such that $\Delta \subseteq \Psi$.*

Proof. Consider a $\Delta \subseteq Cn_{ew}(\Gamma)$ and an $\alpha \in \bar{\mathcal{L}}^P_{X_n}$ such that $\alpha \notin Cn_{Pr}(\Delta)$. Now construct a canonical Δ_c from Δ, adding conditionals $P(\varphi_i|\psi_i) \geq 0$ for every $P(\varphi_i|\psi) \geq q_i \in \Gamma$ and then taking only the highest lower bound assigned to each $(\varphi_i|\psi_i)$. To construct a Δ' such that $\Delta_c, \Delta \subseteq Cn_{ew}(\Delta')$ and $Cn_{ew}(\Delta') \in \Gamma \perp \alpha$, we increase as much as possible the sum of the lower bounds in Δ_c, while consistent with $\neg\alpha$ through solving a quadratic program. Consider $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq m\}$ and $\Delta_c = \{P(\varphi_i|\psi) \geq q'_i | 1 \leq i \leq m\}$, with $q_i \geq q'_i$ for all $1 \leq i \leq m$ (recall that all q'_i are such that $P(\varphi_i|\psi_i) \geq q'_i \in Cn_{ew}(\beta)$ for a $\beta \in \Gamma$). Suppose also that $\alpha = P(\varphi_\alpha|\psi_\alpha) < q_\alpha$, so that $\neg\alpha = P(\neg\varphi_\alpha|\psi_\alpha) \geq 1 - q_\alpha$

Let π be a vector of 2^n variables, corresponding to a probability mass $\pi : W_{X_n} \to [0,1]$ over the possible worlds $w_1, \ldots, w_{2^n}$, *i.e.*, $\pi_j = \pi(w_j)$. Let $r = \langle r_1, \ldots, r_m \rangle$ be a vector of real variables denoting the lower bounds whose sum is being maximized. The first set of restrictions forces these bounds to be consistent, that is, for each $1 \leq i \leq m$, $P_\pi(\varphi_i \wedge \psi_i) - r_i P_\pi(\psi_i) \geq 0$, where $P_\pi(\varphi) = \sum\{\pi_j | w_j \models \varphi\}$. As π must be a probability mass, we have the constraints $\sum \pi = 1$ and $\pi \geq 0$. To not imply α, these probability bounds must be consistent — jointly satisfiable — with $\neg\alpha$, which is encoded into the restriction $P_\pi(\neg\varphi_\alpha \wedge \psi_\alpha) - (1 - q_\alpha)P_\pi(\psi_\alpha) \geq 0$. Finally, we want the bounds r_i to be no lower than those in Δ_c being increased and no greater than the bounds in the base Γ being consolidated: $q_i \geq r_i \geq q'_i$. Wrapping up, we have the following program:

$$
\begin{array}{lll}
\max & r_1 + \cdots + r_m \quad \text{subject to:} & \\
& P_\pi(\varphi_i \wedge \psi_i) - r_i P_\pi(\psi_i) \geq 0 & \text{for all } 1 \leq i \leq m \\
& P_\pi(\neg\varphi_\alpha \wedge \psi_\alpha) - (1 - q_\alpha)P_\pi(\psi_\alpha) \geq 0 & \\
& \sum_{i=1}^{2^n} \pi = 1 & \\
& \pi_1, \ldots, \pi_{2^n} \geq 0 & \\
& q_i \geq r_i \geq q'_i & \text{for all } 1 \leq i \leq m
\end{array}
$$

The program above is quadratic due to the multiplication between the lower bounds r_i and the variables π_j that assign probability mass to the worlds w_j satisfying ψ_i. The program is feasible, for taking $r = q'$ yields a solution, since Δ_c is consistent with $\neg\alpha$; the program is bounded, due to the fact that $r \leq q$ and $\sum r_i \leq \sum q_i$. To prove that the maximisation above converges — the maximum is well-defined —, it suffices to note that the set S of pairs $\langle \pi, r \rangle$ that are feasible solutions to the program is closed, and so is its projection on the values of r. [3]. Let $\langle \pi_i, r^i \rangle \in S$ for $i \in \mathbb{N}$ be such that the following limit exists:

$$\lim_{i \to \infty} \langle \pi^i, r^i \rangle = \langle \pi^\star, r^\star \rangle$$

As each π^i is non-negative, so is the limit $\lim\limits_{i \to \infty} \pi^i = \pi^\star$. Furthermore, as $\sum\limits_{j=1}^{2^n} \pi^i_j = 1$ for all $i \in \mathbb{N}$, and each $\lim\limits_{i \to \infty} \pi^i_j$ exists:

$$\sum_{j=1}^{2^n} \pi^\star_j = \sum_{j=1}^{2^n} \lim_{i \to \infty} \pi^i_j = \lim_{i \to \infty} \sum_{j=1}^{2^n} \pi^i_j = 1$$

Moreover, each r^j is such that $q \geq r^j \geq q'$, and $r^\star = \lim\limits_{i \to \infty} r^i$ will be such that $q \geq r^\star \geq q'$. Any $\langle \pi^i, r^i \rangle \in S$ must satisfy $P_{\pi^i}(\varphi_j \wedge \psi_j) \geq r^i_j P_{\pi^i}(\psi_j)$ for all $1 \leq j \leq m$. As the limits respects non-strict

[3] Our proof is an adaptation from the proof given by Thimm (2013) in his Theorem 1.

inequalities, and the limits of r^i_j and $P_{\pi^i}(\psi_j)$ both exist, we have that:

$$\lim_{i\to\infty} P_{\pi^i}(\varphi_j \wedge \psi_j) \geq \lim_{i\to\infty} r^i_j P_{\pi^i}(\psi_j) = \lim_{i\to\infty} r^i_j \lim_{k\to\infty} P_{\pi^k}(\psi_j)\,,$$

or, equivalently, $P_{\pi^\star}(\varphi_j \wedge \psi_j) \geq r^\star_j P_{\pi^\star}(\psi_j)$ for any $1 \leq j \leq m$. For similar reasons, $\pi^\star$ is such that $P_{\pi^\star}(\neg\varphi_\alpha \wedge \psi_\alpha) \geq (1 - q_\alpha)P_{\pi^\star}(\psi_\alpha)$. Thus, the pair $\langle \pi^\star, r^\star\rangle$ is a solution to the program above, and the set S is closed. As any probabilistic interpretation π and any vector of lower bounds r that are solutions to the program above have all elements within the interval $[0,1]$, $S \subseteq [0,1]^{2^n} \times [0,1]^m$. Since since $[0,1]^{2^n}$ is compact, the projection $\rho : [0,1]^{2^n} \times [0,1]^m \to [0,1]^m$, defined as $\rho(\langle \pi, r\rangle) = r$, is a closed map, by the Tube Lemma (see e.g. (Munkres, 1999)). Hence, the projection $S_r = \{\rho(\langle \pi, r\rangle) | \langle \pi, r\rangle \in S\}$ is closed as well, and the maximum of the program above is well-defined.

Let Δ' be $\{P(\varphi_i|\psi) \geq r_i | 1 \leq i \leq m\}$ where $r_1, \ldots, r_m$ is a solution to the program above. We prove by contradiction that $Cn_{ew}(\Delta')$ is an element of $Cn_{ew}(\Gamma)\bot\alpha$. Suppose $Cn_{ew}(\Delta')$ is not in $Cn_{ew}(\Gamma)\bot\alpha$. That is, $Cn_{ew}(\Delta')$ is not a maximal subset of $Cn_{ew}(\Gamma)$ consistent with $\neg\alpha$. Since, for all $1 \leq i \leq m$, every $s \leq r_i$ is such that $P(\varphi_i|\psi_i) \geq s \in Cn_{ew}(\Delta')$, there must be some $P(\varphi_{i\star}|\psi_{i\star}) \geq s^\star \in Cn_{ew}(\Gamma) \setminus Cn_{ew}(\Delta')$ such that $s^\star > r_{i\star}$ and $Cn_{ew}(\Delta') \cup \{P(\varphi_{i\star}|\psi_i\star) \geq s^\star\}$ is consistent with $\neg\alpha$. Now define the vector $r' = \langle r'_1, \ldots, r'_m\rangle$ such that $r'_i = r_i$ for all $1 \leq i \leq m$ if $i \neq i\star$, and $r'_i = s^\star$ if $i = i\star$. Finally, such r' would yield a feasible solution to the program above, and $\sum_{i=1}^m r'_i > \sum_{i=1}^m r_i$, a contradiction, since $\sum_{i=1}^m r_i$ is maximum. □

Corollary 5.4.4 . *The operator* $- : \bar{\mathcal{L}}^P_{X_n} \to \mathbb{K}$ *for a* $\Gamma = Cn_{ew}(\Gamma)$ *in* $\mathbb{K}_c$ *satisfies (Success), (Inclusion), (Relevance) and (Uniformity) iff* $\Gamma - \alpha = \bigcap \gamma(\Gamma\bot\alpha)$ *for all* $\alpha \in \bar{\mathcal{L}}^P_{X_n}$ *and for some selection function* γ.

Proof. It follows from the Upper bound property, from Theorem 5.4.3, and from the representation result, Theorem 5.1.11. □

Theorem 5.4.5 . *The operator* $- : \bar{\mathcal{L}}^P_{X_n} \to \mathbb{K}$ *for a* $\Gamma \in \mathbb{K}_c$ *satisfies the *lifted versions of success, inclusion, relevance, uniformity and (*$f_\star$*-Canonicity), for* $Cn^\star = Cn_{ew}$ *and* $f_\star = Cn^{-1}_{ew}$*, iff, for all* $\alpha \in \bar{\mathcal{L}}^P_{X_n}$, $\Gamma - \alpha = \inf M$*, where* $M = \{\Psi \in \mathbb{K}_c | Cn_{ew}(\Psi) \in \gamma(Cn_{ew}(\Gamma)\bot\alpha)\}$ *for some selection function* γ.

Proof. ($\to$) By Propositions 5.3.9 and 5.3.12, if $-$ for Γ satisfies (⋆Lifted Success), (⋆Lifted Inclusion), (⋆Lifted Relevance), (⋆Lifted Uniformity) and ($f_\star$-Canonicity), then $\Gamma - \alpha = Cn^{-1}_{ew}(\bigcap \gamma(Cn_{ew}(\Gamma)\bot\alpha)$ for any $\alpha \in \bar{\mathcal{L}}^P_{X_n}$. Every element in $Cn_{ew}(\Gamma)\bot\alpha$ is closed under Cn_{ew} (see the proof of Proposition 5.3.12, $\leftarrow$-part) and can be written as $Cn_{ew}(\Delta)$, for a canonical Δ. Hence, $\gamma(Cn_{ew}(\Gamma)\bot\alpha)$ can be written as $\{Cn_{ew}(\Delta) | \Delta \in M\}$, where M is a (possibly infinite) set bases $M = \{\Delta \in \mathbb{K}_c | Cn_{ew}(\Delta) \in \gamma(Cn_{ew}(\Gamma)\bot\alpha)\}$. Suppose the base Γ is given by $\{P(\psi_i|\varphi_i) \geq q_i | 1 \leq i \leq m\}$. Let each $\Delta \in M$ be equal to $\Lambda_\Gamma(r)$ for some vector $r \in [0,1]^m$. Let $q^{\inf}$ be the infimum of $\{r | \Lambda_\Gamma(r) \in M\}$, so that $\inf M = \Lambda_\Gamma(q^{\inf})$. To prove that $\bigcap \gamma(Cn_{ew}(\Gamma)\bot\alpha) = Cn_{ew}(\inf M)$, suppose $P(\varphi_i|\psi_i) \geq q'_i \in \bigcap \gamma(Cn_{ew}(\Gamma)\bot\alpha)$. Then, $P(\varphi_i|\psi_i) \geq q'_i \in Cn_{ew}(\Delta)$ for all $\Delta \in M$ and $q^{\inf}$ is such that $q^{\inf}_i \geq q'_i$; hence, $P(\varphi_i|\psi_i) \geq q'_i \in Cn_{ew}(\inf M)$. Now suppose $P(\varphi_i|\psi_i) \geq q'_i \in Cn_{ew}(\inf M)$. Consequently, $q^{\inf}$ is such that $q^{\inf}_i \geq q'_i$ and every base $\Delta = \Lambda_\Gamma(r)$ must be such that $r_i \geq q^{\inf}_i$. Hence, $P(\varphi_i|\psi_i) \geq q'_i \in Cn_{ew}(\Delta)$ for all $\Delta \in M$ and $P(\varphi_i|\psi_i) \geq q'_i \in \bigcap \gamma(Cn_{ew}(\Gamma)\bot\alpha)$. Therefore, $\bigcap \gamma(Cn_{ew}(\Gamma)\bot\alpha) = Cn_{ew}(\inf M)$, and $\Gamma! = Cn^{-1}_{ew}(Cn_{ew}(\inf M))$, by Proposition 5.3.12. Finally,

for Cn_{ew}^{-1} is a canonical form and Cn_{ew} is idempotent, $Cn_{ew}^{-1}(Cn_{ew}(\inf M)) = Cn_{ew}^{-1}(\inf M)$, and since $\inf M$ is canonical, $Cn_{ew}^{-1}(\inf M) = \inf M = \Gamma - \alpha$, finishing the proof.

($\leftarrow$) If $\Gamma - \alpha = \inf M$, for a set M of bases $\Delta \in \mathbb{K}_c$ such that $Cn_{ew}(\Delta) \in Cn_{ew}(\Gamma)\perp\alpha$, we have just proven that $\Gamma - \alpha = Cn_{ew}^{-1}(\bigcap \gamma(Cn_{ew}(\Gamma)\perp\alpha))$, for some selection function γ. Hence, $Cn_{ew}(\Gamma - \alpha) = Cn_{ew}(\bigcap \gamma(Cn_{ew}(\Gamma)\perp^{\alpha}))$ and $Cn_{ew}(\Gamma - \alpha) = (\bigcap \gamma(Cn_{ew}(\Gamma)\perp\alpha))$ (see the proof of Proposition 5.3.12, $\leftarrow$-part) — and $\overset{\star}{-}$ for $Cn_{ew}(\Gamma)$ a partial meet contraction. Therefore $-$ for Γ is a liftable partial meet contraction and, by Proposition 5.3.9, the operator $-$ satisfies (*Lifted Success), (*Lifted Inclusion), (*Lifted Relevance) and (*Lifted Uniformity). As $\Gamma - \alpha$ has only one minimum lower bound for each $\varphi_i|\psi_i$, $\Gamma \in \mathbb{K}_c$, and ($f_\star$-Canonicity) is satisfied. □

Proposition 5.5.3 . *Cn_ε is Tarskian and satisfies Cn_{Pr}-dominance.*

Proof. **Inclusion**: Consider a base $\Gamma \in \mathbb{K}^\varepsilon$ and a conditional $\alpha \in \Gamma$. Note that $\alpha \in Cn_\varepsilon(\{\alpha\}) \subseteq Cn_\varepsilon(\Gamma)$, thus inclusion is satisfied.

Monotonicity: Consider a $\Psi \subseteq \Gamma \in \mathbb{K}^\varepsilon$. If $\beta \in Cn_\varepsilon(\Psi)$, there must be an $\alpha \in \Psi$ such that $\beta \in Cn_\varepsilon(\{\alpha\})$. As $\alpha \in \Psi$ implies $\alpha \in \Gamma$, $Cn_\varepsilon(\{\alpha\}) \subseteq Cn_\varepsilon(\Gamma)$ and monotonicity is satisfied.

Idempotence: Consider a conditional α in a base $\Gamma \in \mathbb{K}^\varepsilon$. For every $\gamma \in Cn_\varepsilon(Cn_\varepsilon(\Gamma))$, there is a $\beta \in Cn_\varepsilon(\Gamma)$ such that $\gamma \in Cn_\varepsilon(\{\beta\})$. Similarly, $\beta \in Cn_\varepsilon(\Gamma)$ implies there is some $\alpha \in \Gamma$ such that $\beta \in Cn_\varepsilon(\{\alpha\})$. But note that, if $\beta \in Cn_\varepsilon(\{\alpha\})$, then $Cn_\varepsilon(\{\beta\}) \subseteq Cn_\varepsilon(\{\alpha\})$ and $\gamma \in Cn_\varepsilon(\{\alpha\}) \subseteq Cn_\varepsilon(\Gamma)$. Therefore, $Cn_\varepsilon(Cn_\varepsilon(\Gamma)) \subseteq Cn_\varepsilon(\Gamma)$ and, by inclusion, $Cn_\varepsilon(Cn_\varepsilon(\Gamma)) = Cn_\varepsilon(\Gamma)$.

Cn_{Pr}-dominance: Consider a base $\Gamma \in \mathbb{K}^\varepsilon$ and a conditional $\beta = P(\varphi|\psi) \geq_\varepsilon q$ in $Cn_\varepsilon(\Gamma)$. There must be an $\alpha \in \Gamma$ such that $\beta \in Cn_\varepsilon(\{\alpha\})$, so that $\alpha = P(\varphi|\psi) \geq_{\varepsilon'} q$, for some $\varepsilon' \leq \varepsilon$. Note that, if some probabilistic interpretation π satisfies α, then π also satisfies β and $\beta \in Cn_{Pr}(\{\alpha\})$. As Cn_{Pr} satisfies monotonicity, $Cn_{Pr}(\{\alpha\}) \subseteq Cn_{Pr}(\Gamma)$ and $\beta \in Cn_{Pr}(\Gamma)$, finishing the proof. □

Lemma 5.5.4 (Upper bound property for probabilistic ε-consolidation)**.** *Let $\Gamma \in \mathbb{K}_c^\varepsilon$ be a canonical ε-base. For every consistent $\Delta \subseteq Cn_\varepsilon(\Gamma)$, there is a Δ' such that $\Delta \subseteq \Delta' \subseteq Cn_{ew}(\Gamma)$ and $\Delta' \in Cn_\varepsilon(\Gamma)\perp^\perp$.*

Proof. Consider a $\Delta \subseteq Cn_\varepsilon(\Gamma)$ and an $\alpha \in \bar{\mathcal{L}}_{X_n}^\varepsilon$ such that $\perp \notin Cn_{Pr}(\Delta)$. Now construct a canonical Δ_c from Δ, adding conditionals $P(\varphi_i|\psi_i) \geq_1 q_i$ for every $P(\varphi_i|\psi) \geq q_i \in \Gamma$ and then taking only the ε-conditional with the highest violation for each conditional $P(\varphi_i|\psi_i) \geq q_i$. To construct a Δ' such that $\Delta_c, \Delta \subseteq Cn_\varepsilon(\Delta')$ and $Cn_\varepsilon(\Delta') \in \Gamma\perp\alpha$, we decrease as much as possible the sum of the violations in Δ_c, while consistent with $\neg\alpha$ through solving a linear program. Consider $\Gamma = \{P(\varphi_i|\psi_i) \geq_\varepsilon q_i | 1 \leq i \leq m\}$ and $\Delta_c = \{P(\varphi_i|\psi) \geq_{\varepsilon_i'} q_i | 1 \leq i \leq m\}$, with $\varepsilon_i \leq \varepsilon_i'$ for all $1 \leq i \leq m$ (recall that all ε_i' are such that $P(\varphi_i|\psi_i) \geq_{\varepsilon_i'} q_i \in Cn_\varepsilon(\beta)$ for a $\beta \in \Gamma$).

Let π be a vector of 2^n variables, corresponding to a probability mass $\pi : W_{X_n} \to [0,1]$ over the possible worlds $w_1, \ldots, w_{2^n}$, *i.e.*, $\pi_j = \pi(w_j)$. Let $r = \langle r_1, \ldots, r_m \rangle$ be a vector of real variables denoting the violations whose sum is being minimised. The first set of restrictions forces these violations to be consistent — to not imply $\perp$ —, that is, for each $1 \leq i \leq m$, $P_\pi(\varphi_i \wedge \psi_i) - q_i P_\pi(\psi_i) \geq -r_i$, where $P_\pi(\varphi) = \sum\{\pi_j | w_j \models \varphi\}$. As π must be a probability mass, we have the constraints $\sum \pi = 1$ and $\pi \geq 0$. Finally, we want the violations r_i to be no greater than those in Δ_c being minimised and no lesser than the bounds in the base Γ being consolidated: $\varepsilon_i \leq r_i \leq \varepsilon_i'$. Wrapping

up, we have the following program:

$$\begin{aligned} \min \quad & r_1 + \cdots + r_m \quad \text{subject to:} \\ & P_\pi(\varphi_i \wedge \psi_i) - q_i P_\pi(\psi_i) \geq -\varepsilon_i \quad \text{for all } 1 \leq i \leq m \\ & \textstyle\sum_{i=1}^{2^n} \pi = 1 \\ & \pi_1, \ldots, \pi_{2^n} \geq 0 \\ & \varepsilon_i \leq r_i \leq \varepsilon'_i \qquad\qquad \text{for all } 1 \leq i \leq m \end{aligned}$$

The program is always feasible, for taking $r = \varepsilon'$ yields a solution, since Δ_c is consistent; the program is bounded, due to the fact that $r \geq \varepsilon$ and $\sum r_i \geq \sum \varepsilon_i$. The that the minimisation above converges — the minimum is well-defined — trivially follows from the fact that the program above is linear, defining a Simplex of feasible solutions.

Let Δ' be $\{P(\varphi_i|\psi) \geq_{r_i} q_i | 1 \leq i \leq m\}$ where $r_1, \ldots, r_m$ is a solution to the program above. We prove by contradiction that $Cn_\varepsilon(\Delta')$ is an element of $Cn_\varepsilon(\Gamma)\perp^\perp$. Suppose $Cn_\varepsilon(\Delta')$ is not in $Cn_\varepsilon(\Gamma)\perp^\perp$. That is, $Cn_\varepsilon(\Delta')$ is not a maximal consistent subset of $Cn_\varepsilon(\Gamma)$. Since, for all $1 \leq i \leq m$, every $s \leq r_i$ is such that $P(\varphi_i|\psi_i) \geq_s q_i \in Cn_\varepsilon(\Delta')$, there must be some $P(\varphi_{i\star}|\psi_{i\star}) \geq_{s^\star} q_{i\star} \in Cn_\varepsilon(\Gamma) \setminus Cn_\varepsilon(\Delta')$ such that $s^\star < r_{i\star}$ and $Cn_\varepsilon(\Delta') \cup \{P(\varphi_{i\star}|\psi_i\star) \geq_{s^\star} q_{i\star}\}$ is consistent. Now define the vector $r' = \langle r'_1, \ldots, r'_m \rangle$ such that $r'_i = r_i$ for all $1 \leq i \leq m$ if $i \neq i\star$, and $r'_i = s^\star$ if $i = i\star$. Finally, such r' would yield a feasible solution to the program above, and $\sum_{i=1}^m r'_i < \sum_{i=1}^m r_i$, a contradiction, since $\sum_{i=1}^m r_i$ is minimum. □

Lemma 5.5.5 . *Let $\Gamma \in \mathbb{K}_c^\varepsilon$ be a canonical ε-base. $\Psi \in \mathbb{K}_c^\varepsilon$ is a maximal ε-consolidation of Γ iff $Cn_\varepsilon(\Psi) \in Cn_\varepsilon(\Gamma)\perp^\perp$.*

Proof. The proof of the Proposition is broken into two parts, corresponding to the directions of the bi-implication, where $r = \langle r_1, \ldots, r_m \rangle$, $\varepsilon = \langle \varepsilon_1, \ldots, \varepsilon_m \rangle$ and $q = \langle q_1, \ldots, q_m \rangle$ are vectors in $[0,1]^m$ and $m = |\Gamma|$:

($\rightarrow$) Let $\Psi = \{P(\varphi_i|\psi) \geq_{r_i} q_i | 1 \leq i \leq m\}$ be a maximal ε-consolidation of $\Gamma = \{P(\varphi_i|\psi) \geq_{\varepsilon_i} q_i | 1 \leq i \leq m\}$. As Ψ is consistent, so is $Cn_\varepsilon(\Psi)$, for each of its conditionals is implied by some $\alpha \in \Psi$. If Γ and $Cn_\varepsilon(\Gamma)$ are consistent, $\Psi = \Gamma$ is a maximal ε-consolidation, and $Cn_\varepsilon(\Psi) = Cn_\varepsilon(\Gamma)$ is maximal consistent subset of $Cn_\varepsilon(\Gamma)$; thus it is in $Cn_\varepsilon(\Gamma)\perp^\perp$. So we consider an inconsistent Γ, which must be different from Ψ, which is consistent. Any $\alpha \in Cn_\varepsilon(\Gamma) \setminus Cn_\varepsilon(\Psi)$ is such that $\alpha = P(\varphi|\psi) \geq_{r'_{i\star}} q_{i\star}$ with $r_{i\star} > r'_{i\star} \geq \varepsilon_{i\star}$ for some $1 \leq i\star \leq m$. As Ψ is a maximal ε-consolidation, $\Psi' = \Psi \setminus \{P(\varphi_i|\psi) \geq_{r_{i\star}} q_{i\star}\} \cup \{P(\varphi_i|\psi) \geq_{r'_{i\star}} q_{i\star}\}$ is inconsistent and so is $Cn_\varepsilon(\Psi) \cup \{\alpha\} \supset \Psi'$. Hence, $Cn_\varepsilon(\Psi)$ is a maximal consistent subset of $Cn_\varepsilon(\Gamma)$ and $Cn_\varepsilon(\Psi) \in Cn_\varepsilon(\Gamma)\perp^\perp$.

($\leftarrow$) Consider a $\Psi \in \mathbb{K}_c^\varepsilon$ such that $Cn_\varepsilon(\Psi) \in Cn_\varepsilon(\Gamma)\perp^\perp$. Consider $\Gamma = \{P(\varphi_i|\psi) \geq_{\varepsilon_i} q_i | 1 \leq i \leq m\}$. Since $Cn_\varepsilon(\Psi) \subseteq Cn_\varepsilon(\Gamma)$, $\Psi = \{P(\varphi_i|\psi) \geq_{r_i} q_i | 1 \leq i \leq m\}$ for some $r \geq \varepsilon$. To prove by the contrapositive, suppose Ψ is not a maximal ε-consolidation of Γ. As Ψ is consistent — for $Cn_\varepsilon(\Psi) \supset \Psi$ also is —, there must be a $r' \in [0,1]^m$ such that $r > r' \geq \varepsilon$ and $\Psi' = \{P(\varphi_i|\psi) \geq_{r'_i} q_i | 1 \leq i \leq m\}$ is consistent. Furthermore, $Cn_\varepsilon(\Psi')$ is consistent, since Cn_{Pr} satisfies idempotence and monotonicity and Cn_ε satisfies Cn_{Pr}-dominance. . Finally, as $Cn_\varepsilon(\Psi) \subsetneq Cn_\varepsilon(\Psi') \subseteq Cn_\varepsilon(\Gamma)$, $Cn_\varepsilon(\Psi)$ is not a maximal consistent subset of $Cn_\varepsilon(\Gamma)$. □

Lemma 5.5.6 . *Consider $Cn^\star = Cn_\varepsilon$ and $f_\star = Cn_\varepsilon^{-1}$. The operation $\Gamma!$ satisfies (*Lifted Success),*

(⋆Lifted Inclusion), (⋆Lifted Fullness) and ($f_\star$-Canonicity) for all $\Gamma \in \mathbb{K}_c$ iff $\mathcal{V} : \mathbb{K}_c \to \mathbb{K}_c^\varepsilon$, with $\mathcal{V}(\Gamma) = \Gamma!$ for all $\Gamma \in \mathbb{K}_c$, is a maximal ε-consolidation operator.

Proof. ($\to$) By Propositions 5.3.10 and 5.3.13, if $\Gamma!$ satisfies (⋆Lifted Success), (⋆Lifted Inclusion), (⋆Lifted Fullness) and ($f_\star$-Canonicity), then $\Gamma! = Cn_\varepsilon^{-1}(\Psi)$, for some $\Psi \in Cn_\varepsilon(\Gamma)\perp^\perp$. As any $\Psi \in Cn_\varepsilon(\Gamma)\perp^\perp$ is closed under Cn_ε, it can be written as $\Psi = Cn_\varepsilon(\Delta)$, for a canonical $\Delta \in \mathbb{K}_c^\varepsilon$ (see the proof of Proposition 5.3.12, $\leftarrow$-part). By Lemma 5.5.5, Δ is a maximal ε-consolidation of Γ. By the idempotence of Cn_ε, $Cn_\varepsilon(\Psi) = Cn_\varepsilon(\Delta)$. As Cn_ε^{-1} is a canonical form, $Cn_\varepsilon^{-1}(\Psi) = Cn_\varepsilon^{-1}(\Delta)$, and as $\Delta \in \mathbb{K}_c$, $Cn_\varepsilon^{-1}(\Delta) = \Delta = Cn_\varepsilon^{-1}(\Psi)$. Hence, $\Gamma! = \Delta$.

($\leftarrow$) By Lemma 5.5.5, $Cn_\varepsilon(\Gamma!) \in Cn_\varepsilon(\Gamma)\perp^\perp$ for all $\Gamma \in \mathbb{K}_c$, since $\mathcal{V}(\Gamma) = \Gamma!$ is a maximal ε-consolidation operator. Therefore, for any $\Gamma \in \mathbb{K}_c$, $Cn_\varepsilon(\Gamma)!^\star = Cn_\varepsilon(\Gamma!)$ is a maxichoice consolidation, $\Gamma!$ is a liftable maxichoice consolidation and, by Proposition 5.3.10, $\Gamma!$ satisfies (⋆Lifted Success), (⋆Lifted Inclusion), (⋆Lifted Fullness). As $\Gamma! = \mathcal{V}(\Gamma)$ is a maximal ε-consolidation for any $\Gamma \in \mathbb{K}_c$, it is in $\mathbb{K}_c^\varepsilon$, and ($f_\star$-Canonicity) is satisfied. □

A.4 Technical Results from Chapter 6

Proposition 6.2.2 . *$\mathcal{I}_{CRV}$ satisfies (Weak Independence).*

Proof. That $\mathcal{I}_{CRV}$ satisfies (Weak Independence) follows from the fact that so does $\mathcal{I}_{CRV}^P$. Consider a precise base $\Psi = \Gamma \cup \{P(\varphi|\psi) = q\}$ in $\mathbb{K}_c$, where $P(\varphi|\psi) = q$ is safe. Let Ψ be built over the set of atoms $X_n = \{x_1, \ldots, x_n\}$. Suppose the set of atoms appearing in $\alpha = P(\varphi|\psi) = q$ is $X_\alpha = \{x_1, \ldots, x_m\}$, for some $m < n$. As α is satisfiable (for it is safe), there is a probability mass $\pi_\alpha : W_{X_\alpha} \to [0, 1]$ satisfying it, where W_{X_α} is the set containing the 2^m possible worlds with atoms from X_α. The base $\Gamma = \Psi \setminus \{\alpha\}$ is built over the set of atoms $X_\Gamma = X_n \setminus X_\alpha$. Let $\pi_\Gamma : W_{X_\Gamma} \to [0, 1]$ be a probability mass such that $\mathcal{I}_{CRV}^P(\Gamma) = d_{CRV}(\Gamma, \pi_\Gamma)$, where W_{X_Γ} is the set containing the 2^{n-m} possible worlds with atoms from X_Γ. Consider the probability mass $\pi : W_{X_n} \to [0, 1]$ such that $\pi(w_i \wedge w_j) = \pi_\alpha(w_i) \times \pi_\Gamma(w_j)$ for any pair $(w_i, w_j) \in W_{X_\alpha} \times W_{X_\Gamma}$. Note that π satisfies α, thus $d_{CRV}(\{\alpha\}, \pi) = 0$. Furthermore, $d_{CRV}(\Gamma, \pi) = d_{CRV}(\Gamma, \pi_\Gamma)$, therefore $\mathcal{I}_{CRV}^P(\Psi) \leq d_{CRV}(\Gamma, \pi) + d_{CRV}(\{\alpha\}, \pi) = \mathcal{I}_{CRV}^P(\Gamma)$. Finally, since $\mathcal{I}_{CRV}^P$ is monotonic, $\mathcal{I}_{CRV}^P(\Psi) = \mathcal{I}_{CRV}^P(\Gamma)$, and (Weak Independence) is satisfied. □

Theorem 6.2.4 . *Consider a knowledge base $\Gamma \in \mathbb{K}$ and a probabilistic conditional $\alpha \in \Gamma$. The following statements are equivalent:*

1. *There is no minimal abrupt repair Δ of Γ such that $\alpha \in \Delta$.*
2. *For all maximal abrupt consolidation Γ' of Γ, $\alpha \in \Gamma'$.*
3. *If $\Gamma' = \Gamma \setminus \Delta$ is an abrupt consolidation of Γ (equivalently, Δ is an abrupt repair of Γ), then α is consistent with Γ'.*
4. *There is no minimal inconsistent set $\Delta \subseteq \Gamma$ such that $\alpha \in \Delta$.*

Proof. The first two items are clearly dual, and the fourth one is the definition of free conditional. Suppose α is free in Γ. Note that all abrupt consolidations Γ' of Γ are consistent with α. As Γ' is consistent, it has no MIS, and adding α cannot create a MIS, for it is free. Thus, if an

abrupt consolidation does not contain α, it is not maximal. Now suppose there is a maximal abrupt consolidation Γ' such that $\alpha \notin \Gamma'$. For Γ' is maximal, α cannot be consistent with it. As Γ' is consistent, it has no MIS, and adding α creates a MIS (that contains α), which also is a MIS of Γ — hence, α cannot be free. □

Lemma 6.2.5 . *Consider a knowledge base $\Gamma \in \mathbb{K}$ and a conditional $\alpha \in \Gamma$. The following statements are equivalent:*

1. *For all maximal Γ' of Γ, $\alpha \in \Gamma'$.*
2. *If Γ' is a natural consolidation, then α is consistent with Γ'.*

Proof. Suppose all maximal consolidations of Γ contain α. For any consolidation Ψ, there is a maximal consolidation Ψ' such that, for each $\beta' \in \Psi'$, there is a $\beta \in \Psi$ such that $\beta' \geq \beta$. Therefore, any probability mass π satisfying Ψ' must also satisfies Ψ, and $\alpha \in \Psi'$ implies π satisfies α as well. Now suppose there is a maximal consolidation Ψ that does not contain α. As $\alpha \in \Gamma$, there is a $\beta \in \Psi$ such that $\beta \leq \alpha$. For Ψ is a maximal consolidation, $(\Psi \setminus \{\beta\}) \cup \{\alpha\}$ cannot be a consolidation and thus is inconsistent. Finally, α is not consistent with Ψ. □

Proposition 6.2.8 . *Consider a probabilistic $\alpha \in \Gamma$. If α is safe, it is innocuous; if α is innocuous, it is free.*

Proof. If $\Gamma = \{\alpha\}$, then α is safe, innocuous and free iff it is satisfiable, thus we focus on $\Gamma \neq \{\alpha\}$. Let Γ be built over the set of atoms $X_n = \{x_1, \ldots, x_n\}$. Suppose α is safe and, without loss of generality, the set of atoms appearing in α is $X_\alpha = \{x_1, \ldots, x_m\}$, for some $m < n$. As α is satisfiable, there is a probability mass $\pi_\alpha : W_{X_\alpha} \to [0, 1]$ satisfying it, where W_{X_α} is the set containing the 2^m possible worlds with atoms from X_α. The base $\Gamma' = \Gamma \setminus \{\alpha\}$ is built over the set of atoms $X_{\Gamma'} = X_n \setminus X_\alpha$. Any natural consolidation Ψ of Γ' must also be formed by atoms in $X_{\Gamma'}$. If Δ is a natural consolidation of Γ, there is a natural consolidation Ψ of Γ' such that $\Delta = \Psi \cup \{\beta\}$, for some β such that $\beta \leq \alpha$. Let $\pi_\Psi : W_{X_{\Gamma'}} \to [0, 1]$ be the probability mass satisfying Ψ, where $W_{X_{\Gamma'}}$ is the set containing the 2^{n-m} possible worlds with atoms from $X_{\Gamma'}$. Consider the probability mass $\pi : W_{X_n} \to [0, 1]$ such that $\pi(w_i \wedge w_j) = \pi_\alpha(w_i) \times \pi_\Psi(w_j)$ for any pair $(w_i, w_j) \in W_{X_\alpha} \times W_{X_{\Gamma'}}$. Note that π satisfies Ψ and α, thus π satisfies Ψ and β. Therefore, α is consistent with any natural consolidation $\Delta = \Psi \cup \{\beta\}$ of Γ and is innocuous by Lemma 6.2.5.

Now suppose α is innocuous. Any abrupt consolidation $\Delta \subseteq \Gamma$ is equivalent (and equisatisfiable) to a natural consolidation $\Delta' \in \Gamma$ such that $\Delta' = \Delta \cup \{P(\varphi|\psi) \geq 0 | P(\varphi|\psi) \geq q \in \Gamma \setminus \Delta\}$. By Lemma 6.2.5, as α is innocuous, it is consistent with any natural consolidation Δ' and, consequently, any abrupt consolidation Δ. Finally, by Theorem 6.2.4, α is free. □

Proposition 6.3.1 . *A knowledge base Γ is a minimal inconsistent set iff Γ is inconsistent and there are no $\Delta_1, \ldots, \Delta_k \subsetneq \Gamma$, with $k \geq 1$, such that:*

1. $\bigcup_{i=1}^{k} \Delta_i = \Gamma$;
2. *For every $\Gamma' \subseteq \Gamma$ if $\Gamma' \cap \Delta_i$ is an abrupt consolidation of Δ_i for all $1 \leq i \leq k$, then Γ' is an abrupt consolidation of Γ.*

Proof. ($\rightarrow$) Suppose Γ is a MIS and there are $\Delta_1, \ldots, \Delta_k \subsetneq \Gamma$ satisfying both items. For any $1 \leq i \leq k$, as $\Delta_i \subsetneq \Gamma$ is consistent, $\Gamma \cap \Delta_i$ is an abrupt consolidation of Δ_i. Thus, by the second item, Γ is an abrupt consolidation of itself, which contradicts the fact that Γ is inconsistent.

($\leftarrow$) Now suppose Γ is inconsistent but not a MIS. Let $MIS(\Gamma) = \{\Delta_1, \ldots, \Delta_m\}$ be the set of MISes in Γ, for some $m \geq 1$. Let Δ_{m+1} denote the set of free formulas in Γ. Clearly, $\bigcup_{i=1}^{m+1} \Delta_i = \Gamma$. Now consider a set $\Gamma' \subseteq \Gamma$ such that $\Gamma' \cap \Delta_i$ is consistent for any $1 \leq i \leq m+1$. If Γ' was inconsistent, it would contain a MIS $\Delta_i \in MIS(\Gamma)$ and $\Gamma' \cap \Delta_i$ would be inconsistent — a contradiction. Thus Γ' is an abrupt consolidation of Γ. □

Lemma 6.3.3 . *A knowledge base Γ is an inescapable conflict iff there is a weakening Γ' of Γ such that Γ' is a minimal inconsistent set.*

Proof. ($\leftarrow$) Consider a minimal inconsistent set Γ' that is a weakening of Γ. To prove by contradiction, suppose Γ is not an inescapable conflict. As its weakening Γ' is inconsistent, Γ also is, for each conditional in Γ' is implied by a conditional in Γ. Hence, as Γ is not an inescapable conflict, there must be $\Delta_1, \ldots, \Delta_k \subsetneq \Gamma$ such that $\bigcup_{i=1}^{k} \Delta_i = \Gamma$ and, if Δ'_i is a natural consolidation of Δ_i, for all $1 \leq i \leq k$, and $\bigcup_{i=1}^{k} \Delta'_i$ is a weakening of Γ, then $\bigcup_{i=1}^{k} \Delta'_i$ is a natural consolidation of Γ. Consider such collection $\Delta_1, \ldots, \Delta_k \subsetneq \Gamma$. Note that, for each Δ_i, there is a weakening $\Psi_i \subsetneq \Gamma'$, defined via $\Psi_i = \{\beta \in \Gamma' | \alpha \in \Delta_i \text{ and } \alpha \subseteq \beta\}$, for $1 \leq i \leq k$. Furthermore any $\Psi_i \subsetneq \Gamma'$ is consistent, for Γ' is a minimal inconsistent set. As $\bigcup_{i=1}^{k} \Psi_i$ is equal to Γ', it is a weakening of Γ. As each Ψ_i is consistent, Γ' must be a natural consolidation of Γ and, thus, Γ' is consistent. This is a contradiction, which proves that Γ is an inescapable conflict.

($\rightarrow$) Let $\Gamma = \{\alpha_1, \alpha_2, \ldots, \alpha_m\}$ be an inescapable conflict and define $\Delta_i = \Gamma \setminus \{\alpha_i\}$, for $1 \leq i \leq m$. Note that $\Delta_1, \ldots, \Delta_m \subsetneq \Gamma$ such that $\bigcup_{i=1}^{m} \Delta_i = \Gamma$. Let Δ'_i denote an arbitrary natural consolidation of Δ_i, for $1 \leq i \leq m$. If every set $\{\Delta'_1, \Delta'_2, \ldots, \Delta'_m | \bigcup_{i=1}^{m} \Delta'_i \text{ is a weakening of } |\Gamma\}$ is such that $\bigcup_{i=1}^{m} \Delta'_i$ is a natural consolidation of Γ, Γ would not be an inescapable conflict. So, there are natural consolidations Δ'_i for each Δ_i ($1 \leq i \leq m$) such that $\bigcup_{i=1}^{m} \Delta'_i = \Gamma'$ is a weakening of Γ but is not consistent. As Γ' is a weakening of Γ, $\Gamma' = \{\alpha'_1, \alpha'_2, \ldots, \alpha'_m\}$, where $\alpha'_i \leq \alpha_i$ for $1 \leq i \leq m$. As $\Delta_i = \Gamma \setminus \{\alpha_i\}$, $\Delta'_i = \Gamma' \setminus \{\alpha'_i\}$, for $1 \leq i \leq m$. Hence, $\Delta'_1, \ldots, \Delta'_m$ are the maximal proper subsets of Γ', and every proper subset of Γ' is consistent. Thus, Γ' is a minimal inconsistent set. □

Corollary 6.3.4 . *Consider two knowledge bases $\Gamma, \Gamma' \in \mathbb{K}$ such that Γ' is a weakening of Γ. If for every inescapable conflict $\Delta \subseteq \Gamma$ its weakening $\{\beta \in \Gamma' \,|\, \alpha \in \Delta \text{ and } \beta \leq \alpha\}$ is consistent, then Γ' is a natural consolidation of Γ.*

Proof. We will prove via the contrapositive: given Γ and a weakening Γ', if Γ' is not a natural consolidation of Γ, then there is an inescapable conflict $\Delta \subseteq \Gamma$ such that the set $\{\beta \in \Gamma' \,|\, \alpha \in \Delta \text{ and } \beta \leq \alpha\}$ is inconsistent.

If Γ' is not a natural consolidation of Γ, Γ' is inconsistent and must contain at least one minimal inconsistent set, that we denote by Δ'. Let Δ be the set $\{\alpha \in \Gamma | \beta \in \Delta' \text{ and } \beta \leq \alpha\}$ — that is, $\Delta' \subseteq \Gamma'$ is a weakening of $\Delta \subseteq \Gamma$. By Lemma 6.3.3, Δ is an inescapable conflict. □

Corollary 6.3.5 . *If Δ is a minimal inconsistent set, then Δ is an inescapable conflict.*

Proof. Just note that Δ is a weakening of itself. By Lemma 6.3.3, Δ is an inescapable conflict. □

Theorem 6.3.7 . *The following statements are equivalent:*

1. *For all maximal consolidation* Γ' *of* Γ, $\alpha \in \Gamma'$.

2. *If* Γ' *is a natural consolidation of* Γ, *then* α *is consistent with* Γ'.

3. *There is no inescapable conflict* Δ *in* Γ *such that* $\alpha \in \Delta$.

4. α *is an innocuous conditional in* Γ.

Proof. By the definition of innocuous conditionals and Lemma 6.2.5, the first, the second and the fourth statements are equivalent. It remains to prove that α is innocuous iff there is no inescapable conflict Δ in Γ such that $\alpha \in \Delta$. As any inconsistent $\{\alpha\}$ is a MIS — thus inconsistent with any natural consolidation of Γ —, it is also an inescapable conflict, by Corollary 6.3.5. Therefore, we focus on satisfiable α.

($\rightarrow$) Let α be innocuous in Γ. Suppose there is an inescapable conflict $\Delta \subseteq \Gamma$ such that $\alpha \in \Delta$. Consider the base $\Psi = \Delta \setminus \{\alpha\}$. Let Ψ' be a consolidation of Ψ. Thus, $\Gamma' = \Psi' \cup \{P(\varphi|\psi) \geq 0 | P(\varphi|\psi) \geq q \in \Gamma \setminus \Psi\}$ is consistent and it is a natural consolidation of Γ. Due to the fact that α is innocuous, α is consistent with Γ' (by Lemma 6.2.5) and, therefore, with Ψ'. Consequently, $\Psi' \cup \{\alpha\}$ is a consolidation of Δ for any natural consolidation Ψ' of Ψ. Furthermore, if $\{\beta\}$ is a natural consolidation of $\{\alpha\}$ (*i.e.*, $\beta \leq \alpha$), $\Psi' \cup \{\beta\}$ is a natural consolidation of Δ. As $\Psi, \{\alpha\} \subsetneq \Delta$ are such that $\Psi \cup \{\alpha\} = \Delta$, and any natural consolidation Ψ' and $\{\beta\}$ of theirs are such that $\Psi' \cup \{\beta\}$ is a natural consolidation of Δ, Δ is not an inescapable conflict, which is a contradiction.

($\leftarrow$) Suppose there is no inescapable conflict Δ in Γ such that $\alpha \in \Delta$. Consider the base $\Psi = \Gamma \setminus \{\alpha\}$. Every natural consolidation Γ' of Γ can be written as $\Gamma' = \Psi' \cup \{\beta\}$, where Ψ' is a natural consolidation of Ψ and $\alpha \subseteq \beta$. As all inescapable conflicts of Γ are in Ψ, by Corollary 6.3.4, $\Psi' \cup \{\alpha\}$ is consistent. Hence, α is consistent with any natural consolidation $\Gamma' = \Psi' \cup \{\beta\}$ and α is innocuous by Lemma 6.2.5. □

Lemma 6.4.2 . *Consider a knowledge base* $\Gamma \in \mathbb{K}$ *and a conditional* $\alpha \in \Gamma$. *The following statements are equivalent:*

1. *For all maximal* ε*-consolidations* Γ' *of* Γ, $\alpha \in \Gamma'$.

2. *If* Γ' *is an* ε*-consolidation of* Γ, *then* α *is consistent with* Γ'.

Proof. Suppose all maximal ε-consolidations of Γ contain α. For any ε-consolidation Ψ, there is a maximal ε-consolidation Ψ' such that, for each $\beta' \in \Psi'$, there is a $\beta \in \Psi$ such that $\beta' \leq_\varepsilon \beta$. Therefore, any probability mass π satisfying Ψ' must also satisfies Ψ, and $\alpha \in \Psi'$ implies π satisfies α as well. Now suppose there is a maximal ε-consolidation Ψ that does not contain α. As $\alpha \in \Gamma$, there is a $\beta \in \Psi$ such that $\beta \geq_\varepsilon \alpha$. For Ψ is a maximal ε-consolidation, $(\Psi \setminus \{\beta\}) \cup \{\alpha\}$ cannot be an ε-consolidation and thus is inconsistent. Finally, α is not consistent with Ψ. □

Proposition 6.4.3 . *Consider a probabilistic conditional* $\alpha \in \Gamma$. *If* α *is safe, it is* ε*-innocuous; if* α *is* ε*-innocuous, it is free.*

Proof. If $\Gamma = \{\alpha\}$, then α is safe, ε-innocuous and free iff it is satisfiable, thus we focus on $\Gamma \neq \{\alpha\}$. Let Γ be built over the set of atoms $X_n = \{x_1, \ldots, x_n\}$. Suppose α is safe and, without loss of generality, the set of atoms appearing in α is $X_\alpha = \{x_1, \ldots, x_m\}$, for some $m < n$. As α is satisfiable, there is a probability mass $\pi_\alpha : W_{X_\alpha} \rightarrow [0, 1]$ satisfying it, where W_{X_α} is the set

containing the 2^m possible worlds with atoms from X_α. The base $\Gamma' = \Gamma \setminus \{\alpha\}$ is built over the set of atoms $X_{\Gamma'} = X_n \setminus X_\alpha$. Any ε-consolidation Ψ of Γ' must also be formed by atoms in $X_{\Gamma'}$. If Δ is an ε-consolidation of Γ, there is an ε-consolidation Ψ of Γ' such that $\Delta = \Psi \cup \{\beta\}$, for some β such that $\beta \geq_\varepsilon \alpha$. Let $\pi_\Psi : W_{X_{\Gamma'}} \to [0,1]$ be the probability mass satisfying Ψ, where $W_{X_{\Gamma'}}$ is the set containing the 2^{n-m} possible worlds with atoms from $X_{\Gamma'}$. Consider the probability mass $\pi : W_{X_n} \to [0,1]$ such that $\pi(w_i \wedge w_j) = \pi_\alpha(w_i) \times \pi_\Psi(w_j)$ for any pair $(w_i, w_j) \in W_{X_\alpha} \times W_{X_{\Gamma'}}$. Note that π satisfies Ψ and α, thus π satisfies Ψ and β. Therefore, α is consistent with any ε-consolidation $\Delta = \Psi \cup \{\beta\}$ of Γ and is ε-innocuous by Lemma 6.4.2.

Now suppose α is ε-innocuous. Any abrupt consolidation $\Delta \subseteq \Gamma$ is equivalent (and equisatisfiable) to an ε-consolidation $\Delta' \in \Gamma$ such that $\Delta' = \Delta \cup \{P(\varphi|\psi) \geq_1 q_i | P(\varphi|\psi) \geq q \in \Gamma \setminus \Delta\}$. By Lemma 6.4.2, as α is ε-innocuous, it is consistent with any ε-consolidation Δ' and, consequently, with any abrupt consolidation Δ. Finally, by Theorem 6.2.4, α is free. □

Lemma 6.4.5 . *A knowledge base Γ is an ε-inescapable conflict iff there is a ε-weakening Γ' of Γ such that Γ' is a minimal inconsistent set.*

Proof. ($\leftarrow$) Consider a minimal inconsistent set Γ' that is a ε-weakening of Γ. To prove by contradiction, suppose Γ is not an ε-inescapable conflict. As its ε-weakening Γ' is inconsistent, Γ also is, for each ε-conditional in Γ' is implied by a ε-conditional in Γ. Hence, as Γ is not an ε-inescapable conflict, there must be $\Delta_1, \ldots, \Delta_k \subsetneq \Gamma$ such that $\bigcup_{i=1}^k \Delta_i = \Gamma$ and, if Δ'_i is an ε-consolidation of Δ_i, for all $1 \leq i \leq k$, and $\bigcup_{i=1}^k \Delta'_i$ is a ε-weakening of Γ, then $\bigcup_{i=1}^k \Delta'_i$ is an ε-consolidation of Γ. Consider such collection $\Delta_1, \ldots, \Delta_k \subsetneq \Gamma$. Note that, for each Δ_i, there is an ε-weakening $\Psi_i \subsetneq \Gamma'$, defined via $\Psi_i = \{\beta \in \Gamma' | \alpha \in \Delta_i \text{ and } \alpha \leq_\varepsilon \beta\}$, with $1 \leq i \leq k$. Furthermore any $\Psi_i \subsetneq \Gamma'$ is consistent, for Γ' is a minimal inconsistent set. As $\bigcup_{i=1}^k \Psi_i$ is equal to Γ', it is an ε-weakening of Γ. As each Ψ_i is consistent, Γ' must be an ε-consolidation of Γ and, thus, Γ' is consistent. This is a contradiction, which proves that Γ is an ε-inescapable conflict.

($\rightarrow$) Let $\Gamma = \{\alpha_1, \alpha_2, \ldots, \alpha_m\}$ be an ε-inescapable conflict and define $\Delta_i = \Gamma \setminus \{\alpha_i\}$, for $1 \leq i \leq m$. Note that $\Delta_1, \ldots, \Delta_m \subsetneq \Gamma$ such that $\bigcup_{i=1}^m \Delta_i = \Gamma$. Let Δ'_i denote an arbitrary ε-consolidation of Δ_i, for $1 \leq i \leq m$. If every set $\{\Delta'_1, \Delta'_2, \ldots, \Delta'_m | \bigcup_{i=1}^m \Delta'_i$ is a ε-weakening of $|\Gamma\}$ is such that $\bigcup_{i=1}^m \Delta'_i$ is an ε-consolidation of Γ, Γ would not be an ε-inescapable conflict. So, there are ε-consolidations Δ'_i for each Δ_i $(1 \leq i \leq m)$ such that $\bigcup_{i=1}^m \Delta'_i = \Gamma'$ is an ε-weakening of Γ but is not consistent. As Γ' is an ε-weakening of Γ, $\Gamma' = \{\alpha'_1, \alpha'_2, \ldots, \alpha'_m\}$, where $\alpha_i \leq_\varepsilon \alpha'_i$ for $1 \leq i \leq m$. As $\Delta_i = \Gamma \setminus \{\alpha_i\}$, $\Delta'_i = \Gamma' \setminus \{\alpha'_i\}$, for $1 \leq i \leq m$. Hence, $\Delta'_1, \ldots, \Delta'_m$ are the maximal proper subsets of Γ', and every proper subset of Γ' is consistent. Thus, Γ' is a minimal inconsistent set. □

Corollary 6.4.6 . *Consider two knowledge bases $\Gamma, \Gamma' \in \mathbb{K}$ such that Γ' is a ε-weakening of Γ. If, for every ε-inescapable conflict $\Delta \subseteq \Gamma$, the corresponding ε-weakening $\{\beta \in \Gamma' \,|\, \alpha \in \Delta$ and $\alpha \leq_\varepsilon \beta\}$ is consistent, then Γ' is an ε-consolidation of Γ.*

Proof. We will prove via the contrapositive: given Γ and an ε-weakening Γ', if Γ' is not a natural consolidation of Γ, then there is an ε-inescapable conflict $\Delta \subseteq \Gamma$ such that the set $\{\beta \in \Gamma' \,|\, \alpha \in \Delta$ and $\alpha \leq_\varepsilon \beta\}$ is inconsistent.

If Γ' is not an ε-consolidation of Γ, Γ' is inconsistent and must contain at least one minimal inconsistent set, that we denote by Δ'. Let Δ be the set $\{\alpha \in \Gamma | \beta \in \Delta' \text{ and } \beta \leq \alpha\}$ — that is, $\Delta' \subseteq \Gamma'$ is an ε-weakening of $\Delta \subseteq \Gamma$. By Lemma 6.4.5, Δ is an ε-inescapable conflict. □

Corollary 6.4.7 . *If Δ is a minimal inconsistent set, then Δ is an ε-inescapable conflict.*

Proof. Just note that Δ is a ε-weakening of itself. By Lemma 6.4.5, Δ is an ε-inescapable conflict. □

Theorem 6.4.8 . *The following statements are equivalent:*

1. *For all maximal ε-consolidations Γ' of Γ, $\alpha \in \Gamma'$.*
2. *If Γ' is an ε-consolidation of Γ, then α is consistent with Γ'.*
3. *There is no ε-inescapable conflict Δ in Γ such that $\alpha \in \Delta$.*
4. *α is an ε-innocuous conditional in Γ.*

Proof. By the definition of ε-innocuous conditionals and Lemma 6.4.2, the first, the second and the fourth statements are equivalent. It remains to prove that α is ε-innocuous iff there is no ε-inescapable conflict Δ in Γ such that $\alpha \in \Delta$. As any inconsistent $\{\alpha\}$ is a MIS — thus inconsistent with any ε-consolidation of Γ —, it is also an ε-inescapable conflict, by Corollary 6.4.7. Therefore, we focus on satisfiable α.

($\rightarrow$) Let α be ε-innocuous in Γ. To prove by contradiction, suppose there is an ε-inescapable conflict $\Delta \subseteq \Gamma$ such that $\alpha \in \Delta$. Consider the base $\Psi = \Delta \setminus \{\alpha\}$. Let Ψ' be an ε-consolidation of Ψ. Thus, $\Gamma' = \Psi' \cup \{P(\varphi|\psi) \geq_1 q | P(\varphi|\psi) \geq q \in \Gamma \setminus \Psi\}$ is consistent and it is an ε-consolidation of Γ. Due to the fact that α is ε-innocuous, α is consistent with Γ' (by Lemma 6.4.2) and, therefore, with Ψ'. Consequently, $\Psi' \cup \{\alpha\}$ is an ε-consolidation of Δ for any ε-consolidation Ψ' of Ψ. Furthermore, if $\{\beta\}$ is an ε-consolidation of $\{\alpha\}$ (*i.e.*, $\beta \geq_\varepsilon \alpha$), $\Psi' \cup \{\beta\}$ is an ε-consolidation of Δ. As $\Psi, \{\alpha\} \subsetneq \Delta$ are such that $\Psi \cup \{\alpha\} = \Delta$, and any ε-consolidation Ψ' and $\{\beta\}$ of theirs are such that $\Psi' \cup \{\beta\}$ is an ε-consolidation of Δ, Δ is not an ε-inescapable conflict, which is a contradiction.

($\leftarrow$) Suppose there is no ε-inescapable conflict Δ in Γ such that $\alpha \in \Delta$. Consider the base $\Psi = \Gamma \setminus \{\alpha\}$. Every ε-consolidation Γ' of Γ can be written as $\Gamma' = \Psi' \cup \{\beta\}$, where Ψ' is an ε-consolidation of Ψ and $\alpha \leq_\varepsilon \beta$. As all ε-inescapable conflicts of Γ are in Ψ, by Corollary 6.4.6, $\Psi' \cup \{\alpha\}$ is consistent. Hence, α is consistent with any ε-consolidation $\Gamma' = \Psi' \cup \{\beta\}$ and α is ε-innocuous by Lemma 6.4.2. □

Lemma 6.4.9 . *For any unconditional probabilistic knowledge base $\Gamma \in \mathbb{K}_c$, Γ is an inescapable conflict iff Γ is an ε-inescapable conflict.*

Proof. Consider an arbitrary unconditional $\Gamma = \{P(\varphi_i) \geq q_i | 1 \leq i \leq m\}$ in $\mathbb{K}_c$. By Lemma 6.4.5, Γ is an ε-inescapable conflict iff there is an ε-weakening $\Gamma' = \{P(\varphi_i) \geq_{\varepsilon_i} q_i | 1 \leq i \leq m\}$, with $0 \leq \varepsilon_i \leq q_1$ for all $1 \leq i \leq m$, that is a MIS. Note that Γ' is a MIS iff $\Gamma'' = \{P(\varphi_i) \geq q_i - \varepsilon_i | 1 \leq i \leq m\}$ is a MIS as well. Finally, by Lemma 6.3.3, Γ is an inescapable conflict iff there is a weakening $\Gamma'' = \{P(\varphi_i) \geq q_i - \varepsilon_i | 1 \leq i \leq m\}$ that is MIS. □

Corollary 6.4.10 . *For any unconditional probabilistic knowledge base $\Gamma \in \mathbb{K}_c$ and conditional $\alpha \in \Gamma$, α is innocuous in Γ iff α is ε-innocuous in Γ.*

Proof. As all subsets of Γ are unconditional bases, each subset $\Psi \subseteq \Gamma$ is an inescapable conflicts iff it is an ε-inescapable conflict, by Lemma 6.4.9. Hence, α does not belong to any inescapable conflict in Γ iff it is not in any ε-inescapable conflict, and the result follows from Theorems 6.3.7 and 6.4.8. □

Corollary 6.5.2 . *If $\mathcal{I}$ satisfies (Independence), then $\mathcal{I}$ satisfies (i-Independence). If $\mathcal{I}$ satisfies (i-Independence), then $\mathcal{I}$ satisfies (Weak Independence).*

Proof. It follows directly from the definitions and Proposition 6.2.8. □

Corollary 6.5.4 . *If $\mathcal{I}$ satisfies (MIS-separability), then $\mathcal{I}$ satisfies (IC-separability).*

Proof. If follows directly from the definitions and Corollary 6.3.5. □

Corollary 6.5.5 . *If $\mathcal{I}$ satisfies (IC-separability) and (Consistency), then $\mathcal{I}$ satisfies (i-Independence).*

Proof. Let Γ be a knowledge base and $\alpha \in \Gamma$ an innocuous conditional. As α is innocuous, all inescapable conflicts of Γ are in $\Gamma \setminus \{\alpha\}$ by Lemma 6.3.7. By (IC-separability), we have $\mathcal{I}(\Gamma) = \mathcal{I}(\Gamma \setminus \{\alpha\}) + \mathcal{I}(\{\alpha\})$. As $\{\alpha\}$ is not an inescapable conflict, it is not a MIS, so it is consistent; and, by (Consistency), $\mathcal{I}(\{\alpha\}) = 0$. Finally, $\mathcal{I}(\Gamma) = \mathcal{I}(\Gamma \setminus \{\alpha\})$. □

Lemma 6.5.6 . *For any $p \in \overline{\mathbb{N}}_{>0}$, $\mathcal{I}_p$ satisfies (i-Independence). Furthermore, $\mathcal{I}_p$ satisfies (IC-Separability) iff $p = 1$.*

Proof. **(i-Independence):** Consider the bases $\Gamma = \Lambda_\Gamma(r)$ and $\Psi = \Gamma \setminus \{\alpha\}$ in $\mathbb{K}_c$, where $\alpha = P(\varphi|\psi) \geq q$ is innocuous in Γ. We are going to prove that $\mathcal{I}_p(\Gamma) \leq \mathcal{I}_p(\Psi)$, and the desired result follows from (Monotonicity), by Theorem 3.2.3. Let $\Psi' = \Lambda_\Psi(q')$ be a natural consolidation of $\Psi = \Lambda_\Psi(q)$ such that $\|q' - q\|_p$ is minimised, for a $p \in \overline{\mathbb{N}}_{>0}$, and $\mathcal{I}_p(\Psi) = \|q' - q\|_p$. Note that $\Gamma' = \Psi' \cup \{P(\varphi|\psi) \geq 0\}$ is a natural consolidation of Γ. As $\alpha = P(\varphi|\psi) \geq q$ is innocuous, α is consistent with Γ' and Ψ'. Hence, $\Psi' \cup \{\alpha\} = \Lambda_\Gamma(r')$ is a natural consolidation of Γ. Note that $r' - r$ is $q' - q$ with an extra 0's (from α). Finally, $\mathcal{I}_p(\Gamma) \leq \|r' - r\|_p = \|q' - q\|_p = \mathcal{I}_p(\Psi)$.

(IC-Separability): That (IC-Separability) is violated by $\mathcal{I}_p$ when $p > 2$ can be proven by the counter-example for (Super-Additivity) given in the proof of Lemma 3.2.4; just note there that Δ and Ψ are the only inescapable conflicts in Γ. To prove that (IC-separability) holds for $\mathcal{I}_1$, suppose there are bases $\Psi, \Delta, \Gamma = \Psi \cup \Delta$ in $\mathbb{K}_c$ such that $\Psi \cap \Delta = \varnothing$, $\mathrm{IC}(\Gamma) = \mathrm{IC}(\Psi) \cup \mathrm{IC}(\Delta)$. Let $\Psi' = \Lambda_\Psi(q'), \Delta' = \Lambda_\Delta(r')$, be natural consolidations of $\Psi = \Lambda_\Psi(q), \Delta = \Lambda_\Delta(r)$ that minimise $\|q' - q\|_1, \|r' - r\|_1$, corresponding to $\mathcal{I}_1(\Psi), \mathcal{I}_1(\Delta)$. As $\Gamma' = \Psi' \cup \Delta' = \Lambda_\Gamma(s')$ is a weakening of $\Gamma = \Lambda_\Gamma(s)$ such that, for each $\Phi \in \mathrm{IC}(\Gamma) = \mathrm{IC}(\Psi) \cup \mathrm{IC}(\Delta)$, the base $\{\beta \in \Gamma' \,|\, \alpha \in \Phi \text{ and } \alpha \geq \beta\}$ is consistent (all inescapable conflicts are solved), Γ' is a natural consolidation of Γ by Corollary 6.3.4. As $\|s' - s\|_1 = \|q' - q\|_1 + \|r' - r\|_1 = \mathcal{I}_1(\Psi) + \mathcal{I}_1(\Delta)$, it follows that $\mathcal{I}_1(\Gamma) \leq \mathcal{I}_1(\Psi) + \mathcal{I}_1(\Delta)$. By (Super-additivity), $\mathcal{I}_1(\Gamma) \geq \mathcal{I}_1(\Psi) + \mathcal{I}_1(\Delta)$, thus $\mathcal{I}_1(\Gamma) = \mathcal{I}_1(\Psi) + \mathcal{I}_1(\Delta)$. □

Corollary 6.5.7 . *There is an inconsistency measure $\mathcal{I} : \mathbb{K}_c \to [0, \infty)$ that satisfies (Consistency), (Continuity), (i-Independence) and (IC-Separability).*

Proof. From Theorem 3.2.3 and Lemma 3.2.4, it follows that $\mathcal{I}_1$ satisfies (Consistency), (Continuity), (i-Independence) and (IC-Separability). □

Corollary 6.5.9 . *If $\mathcal{I}$ satisfies (Independence), then $\mathcal{I}$ satisfies (ε-Independence). If $\mathcal{I}$ satisfies (ε-Independence), then $\mathcal{I}$ satisfies (Weak Independence).*

Proof. It follows directly from the definitions and Proposition 6.4.3. □

Corollary 6.5.11 . *If $\mathcal{I}$ satisfies (MIS-separability), then $\mathcal{I}$ satisfies (IC-separability).*

Proof. If follows directly from the definitions and Corollary 6.3.5. □

Corollary 6.5.12 . *If $\mathcal{I}$ satisfies (ε-Separability) and (Consistency), then $\mathcal{I}$ satisfies (ε-Independence).*

Proof. Let Γ be a knowledge base and $\alpha \in \Gamma$ an ε-innocuous conditional. As α is ε-innocuous, all ε-inescapable conflicts of Γ are in $\Gamma \setminus \{\alpha\}$ by Lemma 6.4.8. By (ε-separability), we have $\mathcal{I}(\Gamma) = \mathcal{I}(\Gamma \setminus \{\alpha\}) + \mathcal{I}(\{\alpha\})$. As $\{\alpha\}$ is not an ε-inescapable conflict, it is not a MIS, so it is consistent; and, by (Consistency), $\mathcal{I}(\{\alpha\}) = 0$. Finally, $\mathcal{I}(\Gamma) = \mathcal{I}(\Gamma \setminus \{\alpha\})$. □

Lemma 6.5.13 . *$\mathcal{I}_p^\varepsilon$ satisfies (ε-Independence) for any $p \in \mathbb{N}_{>0} \cup \{\infty\}$. $\mathcal{I}_p^\varepsilon$ satisfies (ε-Separability) iff $p = 1$.*

Proof. See the proof of Theorem 7.4.3, taking $\delta_i = 1$ for all confidence factors δ_i, due to Proposition 7.4.2. It remains to prove that $\mathcal{I}_p^\varepsilon$ violates (ε-Separability) for any $p > 1$. Just look at the counterexample for (Super-Additivity) given in the proof of Lemma 3.2.4. Note there that Δ and Ψ are the only ε-inescapable conflicts in Γ and that, by Proposition 3.2.7, $\mathcal{I}_p(\Gamma) = \mathcal{I}_p^\varepsilon(\Gamma)$ for any unconditional $\Gamma \in \mathbb{K}_c$. □

Corollary 6.5.14 . *There is an inconsistency measure $\mathcal{I} : \mathbb{K}_c \to [0, \infty)$ that satisfies (Consistency), (Continuity), (ε-Independence) and (ε-Separability).*

Proof. By Lemma 6.5.13 and Theorem 3.2.8, $\mathcal{I}_1^\varepsilon$ satisfies (Consistency), (Continuity), (ε-Independence) and (ε-Separability). □

A.5 Technical Results from Chapter 7

Lemma 7.3.3 . *$\mathcal{I}_{SSK}^{a,sum}$, $\mathcal{I}_{SSK}^{b,sum}$, $\mathcal{I}_{SSK}^{a,max}$ and $\mathcal{I}_{SSK}^{b,max}$ are well-defined and satisfy (Consistency), (ε-Independence) and (Monotonicity). $\mathcal{I}_{SSK}^{a,max}$ and $\mathcal{I}_{SSK}^{b,max}$ also satisfy (Super-additivity) and (ε-separability).*

Proof. It follows directly from Theorem 7.4.3 together with Lemma 7.4.7, 7.4.8 or Corollary 7.4.10. □

Lemma 7.3.4 . *$\mathcal{I}_{SSK}^{a,sum}$, $\mathcal{I}_{SSK}^{a,max}$, $\mathcal{I}_{SSK}^{b,sum}$ and $\mathcal{I}_{SSK}^{b,max}$ are continuous for probabilities within $(0, 1)$.*

Proof. Consider the knowledge base $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i' | 1 \leq i \leq m\}$ and the vector $q = \langle q_1, \ldots, q_m \rangle$. By Lemmas 7.4.7, 7.4.8 and Corollary 7.4.10, if $\mathcal{I}$ is a measure in $\{\mathcal{I}_{SSK}^{a,sum}, \mathcal{I}_{SSK}^{a,max}, \mathcal{I}_{SSK}^{b,sum}, \mathcal{I}_{SSK}^{b,max}\}$, there is a weighted base (see Section 7.4) $\Gamma' = \{P(\varphi_i|\psi_i) \geq q_i]\langle\delta_i\rangle | 1 \leq i \leq m\} \in \mathbb{K}_c^w$ and a $p \in \{1, \infty\}$ such that $\mathcal{I}(\Gamma) = \mathcal{I}_p^\delta(\Gamma')$. For $\mathcal{I}_{SSK}^{a,sum}$ and $\mathcal{I}_{SSK}^{b,sum}$, $p = \infty$, and, for all $1 \leq i \leq m$, $\delta_i = 1/q_i$ or $\delta_i = 1/(1 - q_i)$, respectively. For $\mathcal{I}_{SSK}^{a,max}$ and $\mathcal{I}_{SSK}^{b,max}$, $p = 1$, and, for all $1 \leq i \leq m$, $\delta_i = 1/q_i$ or $\delta_i = 1/(1 - q_i)$, respectively.

Note that, for any such measure $\mathcal{I}_p^\delta$, the confidence factors $\delta_1, \ldots, \delta_m$ are finite and positive when $q \in (0, 1)^m$. When these confidence factors are finite, every probability mass $\pi : W_{X_n} \to [0, 1]$ defines a vector $\varepsilon_\pi(q) = \langle \varepsilon_1, \ldots, \varepsilon_m \rangle$ the following way: $\varepsilon_i = -\min\{0, \delta_i(P_\pi(\varphi_i \wedge \psi_i) - q_i P_\pi(\psi_i))\}$ for every $1 \leq i \leq m$ and each $q \in (0, 1)^m$. Note that, for any $q \in (0, 1)^m$, the pair $\pi, \varepsilon_\pi(q)$ is a feasible solution to the program (7.4)–(7.7) that computes $\mathcal{I}_p^\delta(\Lambda_\Gamma(q))$, since $P_\pi(\varphi_i \wedge \psi_i) - q_i P_\pi(\psi_i) \geq -\varepsilon_i/\delta_i$ for all $1 \leq i \leq m$. Thus, every π yields a value for the objective function $h_\pi(q) = \|\varepsilon_\pi(q)\|_p$ of the program (7.4)–(7.7), for any $q \in (0, 1)^m$ — we point out that each δ_i is also a function of q. Let $\mathcal{I}$

be a measure in $\{\mathcal{I}_{SSK}^{a,sum}, \mathcal{I}_{SSK}^{a,max}, \mathcal{I}_{SSK}^{b,sum}, \mathcal{I}_{SSK}^{b,max}\}$. To compute $\mathcal{I}(\Lambda_\Gamma(q))$ for a particular q, one needs to take the minimum in π of the corresponding $\{h_\pi(q)|\pi : W_{X_n} \to [0,1]$ is a probability mass$\}$. As $q \in (0,1)^m$ implies finite $\delta_1, \ldots, \delta_m$, such confidence factors change continuously with $q \in (0,1)^m$, since either $\delta_i = 1/q_i$ or $\delta_i = 1/(1-q_i)$ for any $1 \leq i \leq m$. Hence, $\varepsilon_\pi(q)$ is continuous in $q \in (0,1)^m$, and as any p-norm is a continuous function, $h_\pi : (0,1)^m \to [0,\infty)$ also is for any π. Thus, $h_\pi : (0,1)^m \to [0,\infty)$ is continuous on $q \in (0,1)^m$. As the minimum of continuous functions is continuous, $\mathcal{I} \circ \Lambda_\Gamma : (0,1)^m \to [0,\infty) \cup \{\infty\}$ is continuous. □

Lemma 7.3.5 . *$\mathcal{I}_{SSK}^{a,sum}$ satisfy (Normalisation).*

Proof. When we are computing $\mathcal{I}_{SSK}^{a,sum}$, the maximum sure loss when limiting the agent's total escrows to one. As the agent cannot lose more her total escrow in a Dutch book, $\mathcal{I}_{SSK}^{a,sum}$ is trivially normalized. □

Proposition 7.4.2 . *For any $p \in \overline{\mathbb{N}}_{>0}$ and $\Gamma \in \mathbb{K}_c$, $\mathcal{I}_p^\delta(\Gamma) = \mathcal{I}_p^\varepsilon(\Gamma)$.*

Proof. Just note that, when $\delta_i = 1$ for each confidence factor δ_i in Γ, $\varepsilon_i\delta_i = \varepsilon_i$ for all i. □

Theorem 7.4.3 . *For any $p \in \overline{\mathbb{N}}_{>0}$, $\mathcal{I}_p^\delta : \mathbb{K} \to [0,\infty) \cup \{\infty\}$ is well-defined and satisfies (Consistency), (ε-Independence) and (Monotonicity). $\mathcal{I}_1^\delta$ also satisfies (Super-additivity) and (ε-Separability).*

Proof. Let $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i\langle\delta_i\rangle | 1 \leq i \leq m\}$ be an arbitrary knowledge base in $\mathbb{K}_c^w$. Consider the program from lines (7.4)–(7.7) that compute $\mathcal{I}_p^\delta$, assuming $\varepsilon_i\delta_i = 0$ when $\delta_i = \infty$.

Note that the linear restrictions in the program (7.4)–(7.7), when it is feasible, define a convex, closed region of feasible points (a simplex). The p-norm is a continuous function, so the minimum of the objective function in (7.4) is well-defined for any $p \in \overline{\mathbb{N}}_{>0}$. If the program (7.4)–(7.7) is infeasible for some $\Gamma \in \mathbb{K}_c^w$, $\mathcal{I}_p^\delta(\Gamma)$ is (well-)defined as ∞.

(Consistency): Note that a p-norm is never negative. The base Γ is consistent iff the corresponding program (2.4)–(2.5) is feasible; and such program is feasible iff the program (7.4)–(7.7) has a feasible solution with $\langle\varepsilon_1, \ldots, \varepsilon_m\rangle = \langle 0, \ldots, 0\rangle$; which is the case iff $\|\langle\varepsilon_1, \ldots, \varepsilon_m\rangle\|_p = 0$ is the minimum of the objective function in (7.4).

(Monotonicity): Consider the program $\mathcal{P}$ from lines (7.4)–(7.7), corresponding to the computation of $\mathcal{I}_p^\delta(\Gamma)$, for some $\Gamma \in \mathbb{K}$. Let $\Psi = \Gamma \cup \{\alpha\}$ be a knowledge base. For any $p \in \overline{\mathbb{N}}_{>0}$ and confidence factors $\delta_1, \ldots, \gamma_m > 0$, the program (7.4)–(7.7) whose solution gives $\mathcal{I}_p^\delta(\Psi)$ has two extra constraints in comparison with $\mathcal{P}$. Thus, the program that computes $\mathcal{I}_p^\delta(\Psi)$ cannot reach a smaller value for $\|\langle\varepsilon_1, \ldots, \varepsilon_m\rangle\|_p$, the objective function being minimized by $\mathcal{P}$. Furthermore, $\|\langle\varepsilon_1, \ldots, \varepsilon_{m+1}\rangle\|_p \geq \|\langle\varepsilon_1, \ldots, \varepsilon_m\rangle\|_p$, for any $p \in \overline{\mathbb{N}}_{>0}$ and confidence factors $\delta_1, \ldots, \delta_m \geq 0$. Hence, $\mathcal{I}_p^\delta(\Gamma \cup \{\alpha\}) \geq \mathcal{I}_p^\gamma(\Gamma)$, for any $p \in \overline{\mathbb{N}}_{>0}$.

(ε-Independence): Let $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i\langle\delta_i\rangle | 1 \leq i \leq m\}$ be a knowledge base in $\mathbb{K}_c^w$ and $\alpha = P(\varphi_m|\psi_m) \geq q_m\langle\delta_i\rangle$ be an ε-innocuous conditional in Γ, and define $\Psi = \Gamma \setminus \{\alpha\}$. Suppose $\mathcal{I}_p^\delta(\Psi)$ is finite. The solution on $\langle\varepsilon_1, \ldots, \varepsilon_{m-1}\rangle$ to the program (7.4)–(7.7) that computes $\mathcal{I}_p^\delta(\Psi)$ corresponds to an ε-consolidation of Ψ given by $\Psi' = \{P(\varphi_i|\psi_i) \geq_{\varepsilon_i} | 1 \leq i \leq m-1\}$. For α is ε-innocuous in Γ, it is consistent with $\Psi' \cup \{P(\varphi_m|\psi_m) \geq_1 q_m\}$ (an ε-consolidation of Γ) and $\Psi' \cup \{\alpha\}$ is an ε-consolidation of Γ. Hence, $\langle\varepsilon_1, \ldots, \varepsilon_{m-1}, 0\rangle$ corresponds to a feasible solution to the program (7.4)–(7.7) computing $\mathcal{I}_p^\delta(\Gamma)$. As $\|\langle\varepsilon_1, \ldots, \varepsilon_{m-1}\rangle\|_p$ is equal to $\|\langle\varepsilon_1, \ldots, \varepsilon_{m-1}, 0\rangle\|_p$ for any $p \in \overline{\mathbb{N}}_{>0}$, $\mathcal{I}_p^\gamma(\Gamma) \leq \mathcal{I}_p^\delta(\Psi)$. By (Monotonicity), $\mathcal{I}_p^\delta(\Gamma) = \mathcal{I}_p^\delta(\Psi)$.

Now suppose $\mathcal{I}_p^\delta(\Psi)$ is infinite. Thus, the program (7.4)–(7.7) that computes $\mathcal{I}_p^\delta(\Psi)$ is infeasible. Constraints in such program are inherited by the program that computes $\mathcal{I}_p^\delta(\Gamma) = \mathcal{I}_p^\delta(\Psi \cup \{\alpha\})$ together with the infeasibility, hence $\mathcal{I}_p^\delta(\Gamma) = \infty$ by definition.

(Super-additivity): Suppose there are bases $\Psi, \Delta, \Gamma = \Psi \cup \Delta$ in $\mathbb{K}_c^w$ such that $\Psi \cap \Delta = \varnothing$. Without loss of generality, consider $\Psi = \{P(\varphi_i|\psi_i) \geq q_i\langle\delta_i\rangle | 1 \leq i \leq k\}$, $\Delta = \{P(\varphi_i|\psi_i) \geq q_i\langle\delta_i\rangle | k+1 \leq i \leq m\}$ and $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i\langle\delta_i\rangle | 1 \leq i \leq m\}$. If $\mathcal{I}_1^\gamma(\Gamma) = \infty$, (Super-additivity) trivially holds, then consider $\mathcal{I}_1^\delta(\Gamma)$ is finite. Let $\langle\varepsilon_1, \ldots, \varepsilon_m\rangle$ be part of a solution (that includes π) to the program (7.4)–(7.7) that computes $\mathcal{I}_1^\delta(\Gamma)$, minimizing the objective function. As $\Gamma' = \{P(\varphi_i|\psi_i) \geq_{\varepsilon_i} q_i\langle\delta_i\rangle | 1 \leq i \leq m\}$ is consistent, so are $\Psi' = \{P(\varphi_i|\psi_i) \geq_{\varepsilon_i} q_i\langle\delta_i\rangle | 1 \leq i \leq k\}$ and $\Delta' = \{P(\varphi_i|\psi_i) \geq_{\varepsilon_i} q_i\langle\delta_i\rangle | k+1 \leq i \leq m\}$, which are ε-consolidations of Ψ and Δ (ignoring the confidence factors). Thus, $\langle\varepsilon_1, \ldots, \varepsilon_k\rangle$ and $\langle\varepsilon_{k+1}, \ldots, \varepsilon_m\rangle$ correspond to feasible solutions to the programs that compute $\mathcal{I}_1^\delta(\Psi)$ and $\mathcal{I}_1^\delta(\Delta)$, respectively. It follows that $\mathcal{I}_1^\delta(\Psi) \leq \|\langle\varepsilon_1, \ldots, \varepsilon_k\rangle\|_1$ and $\mathcal{I}_1^\delta(\Delta) \leq \|\langle\varepsilon_{k+1}, \ldots, \varepsilon_m\rangle\|_1$. Finally, $\mathcal{I}_1^\delta(\Delta) + \mathcal{I}_1^\delta(\Psi) \leq \sum_{i=1}^k \varepsilon_i + \sum_{i=k+1}^m \varepsilon_i = \sum_{i=1}^m \varepsilon_i = \mathcal{I}_1^\delta(\Gamma)$.

(ε-separability): To prove that (ε-Separability) holds, suppose there are bases $\Psi, \Delta, \Gamma = \Psi \cup \Delta$ in $\mathbb{K}_c$ such that $\Psi \cap \Delta = \varnothing$, $\mathrm{IC}_\varepsilon(\Gamma) = \mathrm{IC}_\varepsilon(\Psi) \cup \mathrm{IC}_\varepsilon(\Delta)$. Without loss of generality, consider $\Psi = \{P(\varphi_i|\psi_i) \geq q_i\delta_i | 1 \leq i \leq k\}$, $\Delta = \{P(\varphi_i|\psi_i) \geq q_i\delta_i | k+1 \leq i \leq m\}$ and $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i\delta_i | 1 \leq i \leq m\}$. If $\mathcal{I}_1^\delta(\Psi) = \infty$ or $\mathcal{I}_1^\delta(\Delta) = \infty$, then $\mathcal{I}_1^\delta(\Gamma) = \infty$ by (Monotonicity), and (ε-Separability) holds, considering that ∞ plus any non-negative number yields ∞; thus, we assume $\mathcal{I}_1^\delta(\Psi), \mathcal{I}_1^\delta(\Delta) < \infty$. Let $\langle\varepsilon_1, \ldots, \varepsilon_k\rangle$ and $\langle\varepsilon_{k+1}, \ldots, \varepsilon_m\rangle$ be solutions (on ε) to the programs in the form (7.4)–(7.7) that compute $\mathcal{I}_1^\delta(\Psi)$ and $\mathcal{I}_1^\delta(\Delta)$, respectively, minimizing their objective functions. As all ε-inescapable conflicts of Γ are either in Ψ or in Δ, the union of ε-consolidations of Ψ and Δ is an ε-consolidation of Γ, by Corollary 6.4.6. Hence, $\langle\varepsilon_1 \ldots, \varepsilon_m\rangle$ corresponds to a feasible solution to the program in the form (7.4)–(7.7) that computes $\mathcal{I}_1^\delta(\Gamma)$ and $\mathcal{I}_1^\delta(\Gamma) \leq \|\langle\varepsilon_1, \ldots, \varepsilon_m\rangle\|_1 = \sum_{i=1}^k \varepsilon_i + \sum_{i=k+1}^m \varepsilon_i = \mathcal{I}_1^\delta(\Psi) + \mathcal{I}_1^\delta(\Delta)$. By (Super-additivity), $\mathcal{I}_1^\delta(\Gamma) = \mathcal{I}_1^\delta(\Psi) + \mathcal{I}_1^\delta(\Delta)$. □

Lemma 7.4.4 . *If for any weighted conditional $P(\varphi_i|\psi_i) \geq q_i\langle\delta_i\rangle$ in a base $\Gamma \in \mathbb{K}_c^w$, $\delta_i = \infty$ implies $q_i = 0$, then $\mathcal{I}_p^\delta(\Gamma)$ is finite.*

Proof. Consider the knowledge base $\Gamma = \{(\varphi_i|\psi_i) \geq q_i\langle\delta_i\rangle | 1 \leq i \leq m\}$ and the vector $q = \langle q_1, \ldots, q_m\rangle$. Suppose $\delta_i < \infty$ for all $1 \leq i \leq m$. Consider the restrictions $P_\pi(\varphi_i \wedge \psi_i) - q_i P_\pi(\psi_i) \geq -\varepsilon_i/\delta_i$ in the linear program (7.4)–(7.7) that computes $\mathcal{I}_p^\delta(\Gamma)$. As both left-hand sides have values in $[-1, 1]$, taking $\varepsilon_i > \delta_i$ for $1 \leq i \leq m$ relaxes all such constraints. Thus, with any probability mass π we have a feasible solution with the (finite) value for the objective function $\|\langle\varepsilon_1, \ldots, \varepsilon_m\rangle\|_p$, and $\mathcal{I}_p^\delta(\Gamma)$ is finite for any $p \in \overline{\mathbb{N}}_{>0}$.

Now suppose some δ_i are infinite. The corresponding weighted conditionals $P(\varphi_i|\psi_i) \geq q_i\langle\delta_i\rangle \in \Gamma$ are such that $q_i = 0$. Hence, the corresponding linear inequalities are trivially satisfied, and can be ignored. The remaining constraints can be relaxed as explained above, the linear program becomes feasible, and $\mathcal{I}_p^\delta(\Gamma)$ is finite. □

Lemma 7.4.5 . *For any weighted knowledge base $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i'\langle\delta_i\rangle | 1 \leq i \leq m\}$ in $\mathbb{K}_c^w$, if δ_i is finite for any $1 \leq i \leq m$, then $\mathcal{I}_p^\delta \circ \Lambda_\Gamma(q) : [0,1]^m \to [0, \infty)$ is continuous.*

Proof. Consider the weighted knowledge base $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i'\langle\delta_i\rangle | 1 \leq i \leq m\}$ and a vector $q = \langle q_1, \ldots, q_m\rangle$. Every probability mass $\pi : W_{X_n} \to [0, 1]$ defines a vector $\varepsilon_\pi(q) = \langle\varepsilon_1, \ldots, \varepsilon_m\rangle$ for each q in the following way: $\varepsilon_i = -\min\{0, \delta_i(P_\pi(\varphi_i \wedge \psi_i) - q_i P_\pi(\psi_i))\}$ for every $1 \leq i \leq m$ and

$q \in [0,1]^m$. Note that the pair $\pi, \varepsilon_\pi(q)$ is a feasible solution to the program (7.4)–(7.7) that computes $\mathcal{I}_p^\delta(\Lambda_\Gamma(q))$ for any $q \in [0,1]^m$, since $P_\pi(\varphi_i \wedge \psi_i) - q_i P_\pi(\psi_i) \geq -\varepsilon_i/\delta_i$ and $P_\pi(\varphi_i \wedge \psi_i) - q_i P_\pi(\psi_i) \leq \varepsilon_i/\delta_i$ for all $1 \leq i \leq m$. Thus, every π yields a value for the objective function $h_\pi(q) = \|\varepsilon_\pi(q)\|_p$ of the program (7.4)–(7.7), for any $q \in [0,1]^m$. Since for each δ_i is finite, each ε_i is continuous in $q \in [0,1]^m$. Therefore, $\varepsilon_\pi(q)$ is continuous in $q \in [0,1]^m$, and as any p-norm is a continuous function, $h_\pi : [0,1]^m \to [0,\infty)$ also is for any π. To compute $\mathcal{I}_p^\delta(\Lambda_\Gamma(q))$ for a particular q, one needs to take the minimum in π of $\{h_\pi(q) | \pi : W_{X_n} \to [0,1]$ is a probability mass$\}$. As the minimum of continuous functions is continuous, $\mathcal{I}_p^\delta \circ \Lambda_\Gamma : [0,1]^m \to [0,\infty) \cup \{\infty\}$ is continuous for any $p \in \overline{\mathbb{N}}_{>0}$. □

Lemma 7.4.6 . *For all* $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i]\langle\delta_i\rangle | 1 \leq i \leq m\} \in \mathbb{K}_c^w$*, if* $\delta_i \leq 1/q_i$ *for every* $1 \leq i \leq m$*, then* $\mathcal{I}_\infty^\delta(\Gamma) \in [0,1]$.

Proof. Note that when $\delta_i = 1/q_i$, we are limiting the agent's escrows to one, computing $\mathcal{I}_{SSK}^{a,sum}$ by Theorem 7.4.7. As the agent cannot lose more her total escrow in a Dutch book, $\mathcal{I}_{SSK}^{a,sum}$ is normalized. But when $\delta_i \leq 1/q_i$, we are strengthening the restriction over the stakes in the program that maximises sure loss (see Theorem 7.4.7), and it cannot have a higher maximum. Formally, when $\delta_i \leq 1/q_i$, $\sum_{i=1}^m \lambda_i/\delta_i \leq 1$ implies $\sum_{i=1}^m q_i\lambda_i \leq 1$, thus $\mathcal{I}_\infty^\delta \leq \mathcal{I}_{SSK}^a$ in this case, and both are normalized. □

Lemma 7.4.7 . *Consider the knowlegde base* $\Gamma = \{P(\varphi_i|\psi_i)q_i | 1 \leq i \leq m\} \in \mathbb{K}_c$ *and the weighted knowledge base* $\Gamma' = \{P(\varphi_i|\psi_i)q_i\langle\delta_i\rangle | 1 \leq i \leq m\} \in \mathbb{K}_c^w$*. If* $1/\delta_i = q_i$ *for every* $1 \leq i \leq m$*, then* $\mathcal{I}_{SSK}^{a,sum}(\Gamma) = \mathcal{I}_\infty^\delta(\Gamma')$ *and both are finite.*

Proof. To compute $\mathcal{I}_{SSK}^{a,sum}(\Gamma)$, we need to maximize sure loss with the agent's resources limited up to one. Remember that the agent's escrow is how much she can lose in a gamble, which is $\lambda_i q_i$ for a gamble on $P(\varphi_i|\psi_i) \geq q_i$ with stake λ_i. Consider the linear program (7.1)–(7.3) that maximizes sure loss via Dutch books when the agent's beliefs are represented by Γ. Instead of adding the constraint $\lambda_1 + \cdots + \lambda_m \leq 1$ to limit the stakes sum, we limit the agent's total escrow with the restriction $q_1\lambda_1 + \cdots + q_m\lambda_m \leq 1$. Now, as $1/\delta_i = q_i$ for all $1 \leq i \leq m$, such constraint is equal to $\sum_{i=1}^m \lambda_i/\delta_i \leq 1$. Taking the dual of this linear program we recover exactly the program (7.8)–(7.11), which computes $\mathcal{I}_\infty^\delta(\Gamma')$. Note that the sure loss cannot be greater than the agent's total escrow, thus it is always finite in this setting. Hence, by the strong duality theorem, $\mathcal{I}_{SSK}^{a,sum}(\Gamma) = \mathcal{I}_\infty^\delta(\Gamma')$ and both are finite. □

Lemma 7.4.8 . *Consider the knowlegde base* $\Gamma = \{P(\varphi_i|\psi_i)q_i | 1 \leq i \leq m\} \in \mathbb{K}_c$ *and the weighted knowledge base* $\Gamma' = \{P(\varphi_i|\psi_i)q_i\langle\delta_i\rangle | 1 \leq i \leq m\} \in \mathbb{K}_c^w$*. If* $\delta_i = 1/(1-q_i)$ *for every* $1 \leq i \leq m$*, then* $\mathcal{I}_{SSK}^{b,sum}(\Gamma) = \mathcal{I}_\infty^\delta(\Gamma')$.

Proof. To compute $\mathcal{I}_{SSK}^{b,sum}(\Gamma)$, we need to add a restriction to the program (7.1)–(7.3) (which maximises sure loss) in order to limit the gambler's total escrow up to one. Remember that the bettor's escrow is how much he can lose in a gamble, which is $\lambda_i(1-q_i)$ for a gamble on $P(\varphi_i|\psi_i) \geq q_i$ with stake λ_i . With $1/\delta_i = 1 - q_i$ for all $1 \leq i \leq m$, such constraint is equal to $\sum_{i=1}^m \gamma_i/\lambda_i \leq 1$. Once again, the dual of this linear program is the program (7.8)–(7.11), which computes $\mathcal{I}_\infty^\delta(\Gamma')$. By the strong duality theorem, $\mathcal{I}_{SSK}^{b,sum}(\Gamma) = \mathcal{I}_\infty^\delta(\Gamma')$ if both are finite. When $\mathcal{I}_{SSK}^{b,sum}(\Gamma)$ is unbounded, the program (7.8)–(7.11) is infeasible, and, by our definition, $\mathcal{I}_\infty^\delta(\Gamma') = \infty = \mathcal{I}_{SSK}^{g,sum}(\Gamma)$. □

Theorem 7.4.9 . *Consider the knowlegde base* $\Gamma = \{P(\varphi_i|\psi_i)q_i|1 \leq i \leq m\} \in \mathbb{K}_c$ *and the weighted knowledge base* $\Gamma' = \{P(\varphi_i|\psi_i)q_i\langle\delta_i\rangle|1 \leq i \leq m\} \in \mathbb{K}_c^w$. $\mathcal{I}_{SSK}^{\delta,max}(\Gamma) = \mathcal{I}_1^{\delta}(\Gamma')$.

Proof. Consider the linear program (7.1)–(7.3) that maximises sure loss via Dutch books when the agent's beliefs are represented by Γ. We can limit the stakes $\lambda_1, \ldots, \lambda_m$ through restrictions $\lambda_i/\delta_i \leq 1$ for all $1 \leq i \leq m$ in order to compute $\mathcal{I}_{SSK}^{\delta,max}(\Gamma)$. But then the resulting program would be exactly the dual of that in lines (7.4)–(7.7) when $p = 1$. By the strong duality theorem, if the program in lines (7.4)–(7.7) is feasible, both dual programs are feasible, and $\mathcal{I}_{SSK}^{\delta,max}(\Gamma) = \mathcal{I}_1^{\delta}(\Gamma') < \infty$. If (7.4)–(7.7) is infeasible, the program that computes $\mathcal{I}_{SSK}^{\delta,max}(\Gamma)$ is unbounded, by the strong duality theorem, and $\mathcal{I}_{SSK}^{\delta,max}(\Gamma) = \mathcal{I}_1^{\delta}(\Gamma') = \infty$, by definition. □

Corollary 7.4.10 . *Consider a weighted knowledge base* $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i\langle\delta_i\rangle|1 \leq i \leq m\} \in \mathbb{K}_c^w$ *and a knowledge base* $\Gamma' = \{P(\varphi_i|\psi_i) \geq q_i|1 \leq i \leq m\} \in \mathbb{K}_c$. *If* $\delta_i = 1/q_i$ *for all* $1 \leq i \leq m$, *then* $\mathcal{I}_{SSK}^{\delta,max}(\Gamma) = \mathcal{I}_{SSK}^{a,max}(\Gamma') = \mathcal{I}_1^{\delta}(\Gamma)$. *If* $\delta_i = 1/(1-q_i)$ *for all* $1 \leq i \leq m$, *then* $\mathcal{I}_{SSK}^{\delta,max}(\Gamma) = \mathcal{I}_{SSK}^{b,max}(\Gamma') = \mathcal{I}_1^{\delta}(\Gamma)$.

Proof. Just note that when $1/\delta_i = q_i$ for all $1 \leq i \leq m$, $\mathcal{I}_{SSK}^{\delta,max}(\Gamma)$ is the maximum sure loss with $q_i\lambda_i \leq 1$ for $1 \leq i \leq m$, that is, with the agent's escrows limited to one; hence $\mathcal{I}_{SSK}^{\delta,max}(\Gamma) = \mathcal{I}_{SSK}^{a,max}(\Gamma')$. Analogously, $1/\delta_i = -1q_i$ for all $1 \leq i \leq m$, $\mathcal{I}_{SSK}^{\delta,max}(\Gamma)$ is the maximum sure loss with the bettor's escrows limited to one; hence $\mathcal{I}_{SSK}^{\delta,max}(\Gamma) = \mathcal{I}_{SSK}^{g,max}(\Gamma')$. The result follows from Theorem 7.4.9. □

Bibliography

Adlassnig and Kolarz(1986) K.P. Adlassnig and G. Kolarz. Representation and semiautomatic acquisition of medical knowledge in cadiag-1 and cadiag-2. *Computers and Biomedical Research*, 19(1):63–79. Cited in p. 2, 10, 11

Alchourrón *et al.***(1985)** C.E. Alchourrón, P. Gärdenfors and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Symbolic Logic*, 50(02): 510–530. Cited in p. 4, 5, 35, 51, 52

Baioletti and Capotorti() M. Baioletti and A. Capotorti. Efficient l1-based probability assessments correction: Algorithms and applications to belief merging and revision. To appear. Cited in p. 41

Batsell *et al.***(2002)** R. Batsell, L. Brenner, D. Osherson, S. Tsavachidis and M.Y. Vardi. Eliminating incoherence from subjective estimates of chance. In *KR*, pages 353–364. Cited in p. 25, 41

Bertsekas(1999) D.P. Bertsekas. Nonlinear programming. Cited in p. 25

Besnard(2014) P. Besnard. Revisiting postulates for inconsistency measures. In *Logics in Artificial Intelligence*, pages 383–396. Springer. Cited in p. 97

Boole(1854) G. Boole. *An Investigation of the Laws of Thought: on which are Founded the Mathematical Theories of Logic and Probabilities.* Walton and Maberly. Cited in p. 1

Boyd and Vandenberghe(2004) S. Boyd and L. Vandenberghe. *Convex Optimization.* Cambridge university press. Cited in p. 25, 27

Capotorti and Regoli(2008) A. Capotorti and G. Regoli. Coherent correction of inconsistent conditional probability assessments. In *Proceedings of IPMU'08*, pages 891–898. Cited in p. 28, 30, 50

Capotorti and Vantaggi(1998) A. Capotorti and B. Vantaggi. An algorithm for coherent conditional probability assessments. In *Proceedings IV Congresso Nazionale SIMAI*, volume 2, pages 144–148. Cited in p. 16

Capotorti *et al.***(2009)** A. Capotorti, G. Regoli and F. Vattari. On the use of a new discrepancy measure to correct incoherent assessments and to aggregate conflicting opinions based on imprecise conditional probabilities. In *Proceedings of the Sixth International Symposium on Imprecise Probability: Theories and Applications, ISIPTA*. Cited in p. 28, 29, 48, 49

Capotorti *et al.***(2010)** A. Capotorti, G. Regoli and F. Vattari. Correction of incoherent conditional probability assessments. *International Journal of Approximate Reasoning*, 51(6):718–727. Cited in p. 5, 17, 28, 30, 47, 50, 102, 103

Coletti(2002) G. Coletti. Coherent numerical and ordinal probabilistic assessments. *Systems, Man and Cybernetics, IEEE Transactions on*, 24(12):1747–1754. ISSN 0018-9472. Cited in p. 16

Coletti and Scozzafava(1996) G. Coletti and R. Scozzafava. Characterization of coherent conditional probabilities as a tool for their assessment and extension. *International Journal of Uncertainty, Fuzziness and Knowledge Based Systems*, 4:103–128. ISSN 0218-4885. Cited in p. 16

Coletti and Scozzafava(2002) G. Coletti and R. Scozzafava. *Probabilistic Logic in a Coherent Setting*. Kluwer Academic Pub. ISBN 1402009178. Cited in p. 1, 16

Cook(1971) S.A. Cook. The complexity of theorem-proving procedures. In *Proceedings of the third annual ACM symposium on Theory of Computing*, pages 158. ACM. Cited in p. 10

Cozman(2000) F.G. Cozman. Credal networks. *Artificial Intelligence*, 120(2):199–233. Cited in p. 11

Cozman and di Ianni(2013) F.G. Cozman and L.F. di Ianni. Probabilistic satisfiability and coherence checking through integer programming. In *Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, pages 145–156. Springer. Cited in p. 41

De Bona and Finger(2015) G. De Bona and M. Finger. Measuring inconsistency in probabilistic logic: rationality postulates and Dutch book interpretation. *Artificial Intelligence*, 227:140–164. Cited in p. 17, 30, 32, 67, 71, 76, 77, 79

De Bona *et al.***(2013)** G. De Bona, F.G. Cozman and M. Finger. Generalized probabilistic satisfiability. In *Intelligent Systems (BRACIS), 2013 Brazilian Conference on*, pages 182–188. IEEE. Cited in p. 20

De Bona *et al.***(2014)** G. De Bona, F.G. Cozman and M. Finger. Towards classifying propositional probabilistic logics. *Journal of Applied Logic*, 12(3):349–368. Cited in p. 11, 16

De Bona *et al.***(2015)** G. De Bona, F.G. Cozman and M. Finger. Generalized probabilistic satisfiability through integer programming. *Journal of the Brazilian Computer Society*, 21(1):1–14. Cited in p. 20

de Finetti(1930) B. de Finetti. Problemi determinati e indeterminati nel calcolo delle probabilità. *Rendiconti Reale Accademia dei Lincei*, 6:367–373. Cited in p. 1

de Finetti(1949) B. de Finetti. Sull'impostazione Assiomatica del Calcolo Delle Probabilità (Ch. 5 in Probability, Induction, Statistics, Wiley, London). *Annali Univ. Trieste*, 19:3–55. Cited in p. 16

de Finetti(1974) B. de Finetti. Theory of Probability, 1974. Cited in p. 1, 10, 16, 83

de Morais *et al.***(2015)** E.M. de Morais, G. De Bona and M. Finger. Non-markovian logic-probabilistic modeling and inference. In *IJCAI-2015 Workshop on Weighted Logics for Artificial Intelligence (WL4AI-2015)*, pages 48–55. Cited in p. 98

Dubois and Prade(1997) D. Dubois and H. Prade. Focusing vs. belief revision: A fundamental distinction when dealing with generic knowledge. In *Qualitative and Quantitative Practical Reasoning*, pages 96–107. Springer. Cited in p. 97

Finger and De Bona(2015) M. Finger and G. De Bona. Probabilistic satisfiability: algorithms with the presence and absence of a phase transition. *Annals of Mathematics and Artificial Intelligence*, pages 1–39. ISSN 1012-2443. doi: 10.1007/s10472-015-9466-6. URL http://dx.doi.org/10.1007/s10472-015-9466-6. Cited in p. 12

Finger *et al.***(2013)** M. Finger, R. Le Bras, C.P. Gomes and B. Selman. Solutions for hard and soft constraints using optimized probabilistic satisfiability. In *Proceedings of SAT*. Cited in p. 15

Finger and De Bona(2011) Marcelo Finger and Glauber De Bona. Probabilistic satisfiability: Logic-based algorithms and phase transition. In *IJCAI*, pages 528–533. Cited in p. 12, 15

Finthammer *et al.***(2007)** M. Finthammer, G. Kern-Isberner and M. Ritterskamp. Resolving inconsistencies in probabilistic knowledge bases. In *KI 2007: Advances in Artificial Intelligence*, pages 114–128. Springer. Cited in p. 34

Frisch and Haddawy(1994) A.M. Frisch and P. Haddawy. Anytime deduction for probabilistic logic. *Artificial Intelligence*, 69(1):93–122. Cited in p. 11

Garey and Johnson(1979) M.R. Garey and D.S. Johnson. *Computers and Intractability. A Guide to the Theory of NP-Completeness. A Series of Books in the Mathematical Sciences*. WH Freeman and Company, San Francisco, Calif. Cited in p. ix

Georgakopoulos *et al.***(1988)** G. Georgakopoulos, D. Kavvadias and C.H. Papadimitriou. Probabilistic satisfiability. *Journal of Complexity*, 4(1):1–11. Cited in p. 14, 15

Gneiting and Raftery(2007) T. Gneiting and A.E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378. Cited in p. 30

Grant and Hunter(2011) J. Grant and A. Hunter. Measuring the good and the bad in inconsistent information. In *IJCAI Proceedings-International Joint Conference on Artificial Intelligence*, volume 22, pages 2632–2637. Citeseer. Cited in p. 2

Grant and Hunter(2013) J. Grant and A. Hunter. Distance-based measures of inconsistency. In *Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, pages 230–241. Springer. Cited in p. 2

Haenni *et al.***(2011)** R. Haenni, J.W. Romeijn, G. Wheeler and J. Williamson. Probabilistic logic and probabilistic networks, 2011. Cited in p. 1, 11

Hailperin(1965) T. Hailperin. Best possible inequalities for the probability of a logical function of events. *American Mathematical Monthly*, 72(4):343–359. Cited in p. 1

Hailperin(1986) T. Hailperin. *Boole's Logic and Probability*. North-Holland Pub. Co. Cited in p. 1

Halpern(1990) J.Y. Halpern. An analysis of first-order logics of probability. *Artificial Intelligence*, 46(3):311–350. Cited in p. 11

Halpern(2003) J.Y. Halpern. Reasoning about uncertainty. Cited in p. 1

Hansen and Jaumard(2000) P. Hansen and B. Jaumard. Probabilistic satisfiability. *Handbook of Defeasible Reasoning and Uncertainty Management Systems: Algorithms for uncertainty and defeasible reasoning*, pages 321. Cited in p. 1

Hansen and Perron(2008) P. Hansen and S. Perron. Merging the local and global approaches to probabilistic satisfiability. *International Journal of Approximate Reasoning*, 47(2):125–140. Cited in p. 15

Hansen *et al.***(1998)** P. Hansen, B. Jaumard and M.P. de Aragão. Mixed-integer column generation algorithms and the probabilistic maximum satisfiability problem. *European Journal of Operational Research*, 108(3):671–683. Cited in p. 37

Hansson(1992) S.O. Hansson. A dyadic representation of belief. *Belief Revision*, 29:89–121. Cited in p. 52

Hansson(1997) S.O. Hansson. Semi-revision. *Journal of Applied Non-Classical Logics*, 7(1-2): 151–175. Cited in p. 35, 36, 51, 54, 69

Hansson(1999) S.O. Hansson. *A Textbook of Belief Dynamics*. Springer. Cited in p. 5, 35, 51, 52, 53, 54, 55, 69, 107, 113

Howson and Urbach(1989) C. Howson and P. Urbach. *Scientific Reasoning: The Bayesian Approach.* Open Court Publishing Co. Cited in p. 10

Hunter and Konieczny(2005) A. Hunter and S. Konieczny. Approaches to measuring inconsistent information. In *Inconsistency Tolerance*, volume 3300 of *Lecture Notes in Computer Science*, pages 191–236. Springer Berlin Heidelberg. ISBN 978-3-540-24260-4. doi: 10.1007/978-3-540-30597-2_7. URL http://dx.doi.org/10.1007/978-3-540-30597-2_7. Cited in p. 68

Hunter and Konieczny(2006) A. Hunter and S. Konieczny. Shapley inconsistency values. In *10th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pages 249–259. Cited in p. 2, 3, 17, 18, 19, 37, 68

Hunter and Konieczny(2008) A. Hunter and S. Konieczny. Measuring inconsistency through minimal inconsistent sets. In *11th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pages 358–366. Cited in p. 2, 18, 37, 68

Hunter and Konieczny(2010) A. Hunter and S. Konieczny. On the measure of conflicts: Shapley inconsistency values. *Artificial Intelligence*, 174(14):1007–1026. Cited in p. 2, 32, 68, 70, 97

Jaumard *et al.***(1991)** B. Jaumard, P. Hansen and M. Poggi de Aragão. Column generation methods for probabilistic logic. *INFORMS Journal on Computing*, 3(2):135. Cited in p. 15, 20, 41

Joyce(1998) J.M. Joyce. A nonpragmatic vindication of probabilism. *Philosophy of Science*, pages 575–603. Cited in p. 10

Karvetski *et al.***(2013)** C.W. Karvetski, K.C. Olson, D.R. Mandel and C.R. Twardy. Probabilistic coherence weighting for optimizing expert forecasts. *Decision Analysis*, 10(4):305–326. Cited in p. 2

Kavvadias and Papadimitriou(1990) D. Kavvadias and C.H. Papadimitriou. A linear programming approach to reasoning about probabilities. *Annals of Mathematics and Artificial Intelligence*, 1(1):189–205. Cited in p. 1, 15, 20

Kern-Isberner and Lukasiewicz(2004) G. Kern-Isberner and T. Lukasiewicz. Combining probabilistic logic programming with the power of maximum entropy. *Artificial Intelligence*, 157(1): 139–202. Cited in p. 1, 12

Klinov(2011) P. Klinov. *Practical Reasoning in Probabilistic Description Logic.* PhD Thesis, The University of Manchester, Manchester, UK. Cited in p. 19, 37

Knight(2002) K. Knight. Measuring inconsistency. *Journal of Philosophical Logic*, 31(1):77–98. Cited in p. 2, 19, 20, 87

Kullback(1997) S. Kullback. *Information Theory and Statistics.* Courier Dover Publications. Cited in p. 28

Kyburg Jr(1961) H.E. Kyburg Jr. Probability and the logic of rational belief. Cited in p. 89

Lad(1996) F. Lad. *Operational Subjective Statistical Methods: a mathematical, philosophical, and historical introduction*, volume 315. Wiley-Interscience. Cited in p. 29

Liffiton and Sakallah(2005) M.H. Liffiton and K.A. Sakallah. On finding all minimally unsatisfiable subformulas. In *Theory and Applications of Satisfiability Testing*, pages 173–186. Springer. Cited in p. 37

Lukasiewicz(1999) T. Lukasiewicz. Probabilistic deduction with conditional constraints over basic events. *Journal of Artificial Intelligence Research*, 10:199–241. Cited in p. 11

Marques-Silva(2010) J. Marques-Silva. Minimal unsatisfiability: Models, algorithms and applications. In *Multiple-Valued Logic (ISMVL), 2010 40th IEEE International Symposium on*, pages 9–14. IEEE. Cited in p. 37

Muiño(2011) D.P. Muiño. Measuring and repairing inconsistency in probabilistic knowledge bases. *International Journal of Approximate Reasoning*, 52(6):828–840. Cited in p. 2, 5, 15, 17, 22, 23, 24, 39, 41, 52, 93

Munkres(1999) J. Munkres. Topology, 1999. Cited in p. 116

Nau(1981) R.F. Nau. Coherent assessment of subjective probability. Technical report, DTIC Document. Cited in p. 25, 79, 82, 84, 87, 88, 92

Nilsson(1986) N.J. Nilsson. Probabilistic logic. *Artificial Intelligence*, 28(1):71–87. Cited in p. 1, 13, 15

Ognjanovic and Raškovic(2000) Z. Ognjanovic and M. Raškovic. Some first-order probability logics. *Theoretical Computer Science*, 247(1):191–212. Cited in p. 62

Ognjanovic *et al.***(2009)** Z. Ognjanovic, M. Raškovic and Z. Markovic. Probability logics. *Logic in Computer Science, Mathematical Institute of Serbian Academy of Sciences and Arts*, pages 35–111. Cited in p. 1

Papadimitriou and Steiglitz(1998) C.H. Papadimitriou and K. Steiglitz. *Combinatorial Optimization: Algorithms and Complexity*. Dover Pubns. Cited in p. 16

Pearl(2009) J. Pearl. *Causality*. Cambridge university press. Cited in p. 11

Potyka(2014) N. Potyka. Linear programs for measuring inconsistency in probabilistic logics. In *Fourteenth International Conference on Principles of Knowledge Representation and Reasoning (KR-14)*. AAAI. Cited in p. 2, 3, 5, 6, 17, 25, 26, 27, 43, 44, 45, 70, 84, 102, 104

Potyka and Thimm(2014) N. Potyka and M. Thimm. Consolidation of probabilistic knowledge bases by inconsistency minimization. In *ECAI 2014*. Cited in p. 34, 38, 42, 44, 45, 46

Potyka and Thimm(2015) N. Potyka and M. Thimm. Probabilistic reasoning with inconsistent beliefs using inconsistency measures. In *IJCAI 2015*. Cited in p. 89, 98

Prasolov and Tikhomirov(2001) V.V. Prasolov and V.M. Tikhomirov. *Geometry*, volume 200 of *Translations of Mathematical Monographs*. American Mathematical Society. ISBN 0821802364. Cited in p. 14

Reiter(1987) R. Reiter. A theory of diagnosis from first principles. *Artificial intelligence*, 32(1): 57–95. Cited in p. 35, 37, 68, 69

Rödder and Meyer(1996) W. Rödder and C. Meyer. Coherent knowledge processing at maximum entropy by spirit. In *Proceedings of the Twelfth international conference on Uncertainty in artificial intelligence*, pages 470–476. Morgan Kaufmann Publishers Inc. Cited in p. 10, 11

Santos *et al.***(2015)** Y.D. Santos, M.M. Ribeiro and R. Wassermann. Between Belief Bases and Belief Sets: Partial Meet Contraction. In *Proceedings of the International Workshop on Defeasible and Ampliative Reasoning (DARe-15)*. Cited in p. 5, 51, 52, 55, 56, 109

Schervish *et al.***(1998)** M.J. Schervish, T. Seidenfeld and J.B. Kadane. Two measures of incoherence: How not to gamble if you must. Technical report, Department of Statistics, Carnegie Mellon University. Cited in p. 79, 85, 86, 87

Schervish *et al.***(2002a)** M.J. Schervish, T. Seidenfeld and J.B. Kadane. A rate of incoherence applied to fixed-level testing. *Proceedings of the Philosophy of Science Association*, 2002(3): 248–264. Cited in p. 2

Schervish *et al.***(2002b)** M.J. Schervish, T. Seidenfeld and J.B. Kadane. Measuring incoherence. *Sankhyā: The Indian Journal of Statistics, Series A*, pages 561–587. Cited in p. 4, 6, 28, 79, 84, 86

Schervish *et al.***(2003)** M.J. Schervish, J.B. Kadane and T. Seidenfeld. Measures of incoherence: How not to gamble if you must. In *Bayesian Statistics 7: Proceedings of the 7th Valencia Conference on Bayesian Statistics*, pages 385–402. Cited in p. 86

Shannon and Weaver(1949) C.E. Shannon and W. Weaver. The mathematical theory of communication, 1949. Cited in p. 28

Shimony(1955) A. Shimony. Coherence and the axioms of confirmation. *The Journal of Symbolic Logic*, 20(01):1–28. Cited in p. 83

Staffel(2015) J. Staffel. Measuring the overall incoherence of credence functions. *Synthese*, pages 1–27. Cited in p. 4, 79

Thimm(2011) M. Thimm. Probabilistic reasoning with incomplete and inconsistent beliefs. volume 331 of *Dissertations in Artificial Intelligence*. IOS Press, December. Cited in p. 22, 25, 38, 52

Thimm(2013) M. Thimm. Inconsistency measures for probabilistic logics. *Artificial Intelligence*, 197:1–24. Cited in p. 2, 3, 5, 17, 18, 19, 20, 21, 22, 23, 30, 32, 68, 70, 101, 115

Thomson(1891) W. Thomson. *Popular lectures and addresses*, volume 1. Macmillan and Co. Available at https://archive.org/details/popularlectures10kelvgoog. Cited in p. 1

Uchii(1973) S. Uchii. Higher order probabilities and coherence. *Philosophy of Science*, pages 373–381. Cited in p. 11

van Maaren and Wieringa(2008) H. van Maaren and S. Wieringa. Finding guaranteed muses fast. In *Theory and Applications of Satisfiability Testing–SAT 2008*, pages 291–304. Springer. Cited in p. 37

Vanderbei(1996) R.J. Vanderbei. *Linear Programming: Foundations and Extensions*. Cited in p. 83

Wang *et al.***(2011)** G. Wang, S.R. Kulkarni, H.V. Poor and D.N. Osherson. Aggregating large sets of probabilistic forecasts by weighted coherent adjustment. *Decision Analysis*, 8(2):128–144. Cited in p. 2

Wassermann(2000) R. Wassermann. *Resource-Bounded Belief Revision*. PhD Thesis, Institute for Logic, Language and Computation — University of Amsterdam. Cited in p. 54, 55, 103, 107, 108

Wheeler(2012) G. Wheeler. Formal epistemology. In A. Cullison, editor, *The Continuum Companion to Epistemology*. Bloomsbury Publishing. Cited in p. 10

Williamson(2010) J. Williamson. *In Defence of Objective Bayesianism*. Oxford University Press. Cited in p. 12