Informatics

Integrated Preferences in Logic and Abstract Argumentation

DISSERTATION

submitted in partial fulfillment of the requirements for the degree of

Doktor der Technischen Wissenschaften

by

Dipl.-Ing. Michael Bernreiter, BSc
Registration Number 01307069

to the Faculty of Informatics
at the TU Wien

Advisor: Dr. Stefan Woltran
Co-advisor: Dr. Wolfgang Dvořák
Co-advisor: Dr. Jan Maly

The dissertation has been reviewed by:

Gabriele Kern-Isberner

Leon van der Torre

Vienna, June 13, 2024

Michael Bernreiter

Erklärung zur Verfassung der Arbeit

Dipl.-Ing. Michael Bernreiter, BSc

Hiermit erkläre ich, dass ich diese Arbeit selbständig verfasst habe, dass ich die verwendeten Quellen und Hilfsmittel vollständig angegeben habe und dass ich die Stellen der Arbeit – einschließlich Tabellen, Karten und Abbildungen –, die anderen Werken oder dem Internet im Wortlaut oder dem Sinn nach entnommen sind, auf jeden Fall unter Angabe der Quelle als Entlehnung kenntlich gemacht habe.

Wien, 13. Juni 2024

Michael Bernreiter

Acknowledgements

Firstly, I want to thank my brilliant supervisors Stefan Woltran, Wolfgang Dvořák, and Jan Maly. I could not have hoped for better guidance, sturdier support, or more understanding than I received.

I would like to also thank all of my colleagues at the Institute of Logic and Computation for the wonderful last four years I spent with them. Special thank goes to Matthias König and Anna Rapberger for being truly amazing office mates, co-authors, and friends. I must also not forget my other, equally wonderful, co-authors Robert Freiman, Anela Lolic, and Oliviero Nardi. I want to also especially thank Juliane Auerböck for taking all the bureaucratic loads of an academic job off my shoulders, Martin Kronegger for the smooth organization of the Algodat course, and Sanja Lukumbuzya for helping me without hesitation in a time of need.

I want to thank my family and friends for their unwavering support, not only in the past four years but throughout my entire life. I will never be able to properly express the gratitude I have for my parents who were always there for me, accepted all of my decisions, and never asked for anything in return. I want to thank my lovely girlfriend Yol who has always believed in me and accompanied me through many hard times. A special mention also goes to my dear friends Mathias and Sam, my brothers Benjamin and Tobias, and my deceased grandmother Gertrude who, in her down-to-earth way, was one of the most inspiring people I ever knew personally.

Last but not least, I would like to thank the reviewers of my thesis, Gabriele Kern-Isberner and Leon van der Torre, for carefully reading my thesis and providing me with valuable suggestions and feedback.

Kurzfassung

Präferenzen sind ein wichtiges Konzept im Bereich der Künstlichen Intelligenz sowie dem Teilbereich Knowledge Representation and Reasoning (KR). Eine wichtige Forschungsfrage in KR ist wie Präferenzen zusammen mit Informationen über objektive Fakten repräsentiert und verarbeitet werden können. In dieser Arbeit werden zwei Formalismen aus dem Bereich KR welche Präferenzen unterschiedlich integrieren untersucht. Zum einen Choice Logics, welche klassische Aussagenlogik um nicht-klassische Präferenz-Konnektive erweitern, zum anderen Abstract Argumentation, wo Konflikte zwischen Argumenten von Präferenzen beeinflusst werden können. In beiden Formalismen sind die Konzepte von Präferenzen und Fakten eng miteinander verwoben, sei es, weil Präferenzen und Fakten zusammen repräsentiert werden, wie in Choice Logics, oder weil die zwei Konzepte zusammen ausgewertet werden, wie in Abstract Argumentation. Um diese Art von integrierten Präferenzen im Bereich KR besser zu verstehen, untersuchen wir Choice Logics und Abstract Argumentation bezüglich ihrer syntaktischen, semantischen, und komplexitätstheoretischen Eigenschaften.

In Abstract Argumentation verwenden wir vier sogenannte Reduktionen aus der Literatur um Präferenzen aufzulösen, und untersuchen den Einfluss dieser Reduktionen in zwei verschiedenen Situationen. Erstens führen wir Conditional Preference-based Argumentation Frameworks (CPAFs) ein, ein neuer Formalismus mit dessen Hilfe man bedingte Präferenzen darstellen kann. Wir erforschen die Eigenschaften von CPAFs, und zeigen dass die Wahl der Reduktion sowohl semantische Eigenschaften als auch die Komplexität von wichtigen Entscheidungsproblemen beeinflusst. Zweitens verallgemeinern wir Claim-augmented Argumentation Frameworks (CAFs) indem wir Preference-based CAFs (PCAFs) einführen. Das Anwenden der Reduktionen führt dazu, dass die in CAFs zentrale syntaktische Eigenschaft der "well-formedness" nicht mehr garantiert werden kann. Wir zeigen allerdings, dass der Einfluss der Reduktionen nicht willkürlich ist, sondern Teile der Struktur, welche mit well-formedness einhergeht, erhalten bleiben. Des Weiteren zeigen wir, dass manche Reduktionen vorteilhafte semantische und komplexitätstheoretische Eigenschaften von well-formedness erhalten.

Bezüglich Choice Logics untersuchen wir Preferred Model Entailment, ein zentrales Konzept in Choice Logics welches nichtmonotones Schließen ermöglicht. Wir berücksichtigen dabei insbesondere Qualitative Choice Logic (QCL), Conjunctive Choice Logic (CCL), und Lexicographic Choice Logic (LCL), sowie mehrere Methoden um die präferierten

Modelle einer Theorie zu berechnen. Wir beweisen, dass Preferred Model Entailment wichtige logische Eigenschaften für nichtmonotones Schließen erfüllt. Des Weiteren zeigen wir, dass die Komplexität von Preferred Model Entailment auf der zweiten Stufe der polynomiellen Hierarchie liegt, wobei die exakte Komplexität sowohl von der verwendeten Logik (QCL, CCL, LCL,...) als auch der verwendeten Methode zum Berechnen der präferierten Modelle abhängt. Außerdem führen wir Sequenzkalküle für Preferred Model Entailment ein, und zeigen deren Korrektheit und Vollständigkeit.

Schlussendlich beleuchten wir Zusammenhänge zwischen Choice Logics und Abstract Argumentation indem wir QCL-Theorien zu SETAFs (Argumentation Frameworks with Collective Attacks) transformieren. Wir zeigen, dass die ursprüngliche QCL-Theorie und das konstruierte SETAF semantisch äquivalent sind, und dass das SETAF verwendet werden kann um Preferred Model Entailment in der ursprünglichen QCL-Theorie zu entscheiden. Im Gegensatz zu bereits bestehenden Transformationen von Choice Logics zu Abstract Argumentation ist unsere Konstruktion rein syntaktischer Natur und polynomiell in Größe sowie Laufzeit.

Abstract

Preferences are an important notion in Artificial Intelligence and many of its subfields such as Knowledge Representation and Reasoning (KR). A key challenge when it comes to preferences (or soft-constraints) in KR is how they can be best represented alongside knowledge about truth (hard-constraints), and what effect they should have in view of the given hard-constraints. In this thesis, we study two KR-formalisms featuring preferences, namely choice logics, which extend classical propositional logic with additional non-classical choice connectives, and abstract argumentation with preferences, where the attack relation between arguments is influenced by a given preference ordering. While choice logics and abstract argumentation are quite different from each other, they have in common that hard- and soft-constraints are tightly interlinked. This motivates us to identify the notion of integrated preferences, where hard- and soft-contraints are represented and/or resolved jointly instead of separately. To better understand integrated preferences in KR, we examine the syntactic, semantic, and computational properties of choice logics and abstract argumentation with preferences.

Regarding argumentation, we consider four so-called preference reductions from the literature and study their effects in two settings. Firstly, we introduce Conditional Preference-based Argumentation Frameworks (CPAFs), a novel formalism capable of expressing and reasoning with conditional preferences in abstract argumentation. We formally study CPAFs, and show that the choice of preference reduction has an impact on the behavior of semantics and the computational complexity of main reasoning tasks. Secondly, we generalize Claim-augmented Argumentation Frameworks (CAFs) by introducing Preference-based CAFs (PCAFs). Since the introduction of preferences to CAFs means that the important property of well-formedness can not be guaranteed, we analyze PCAFs from a syntactic, semantic, and computational perspective to better understand the impact of preferences in claim-based argumentation. Our syntactic analysis shows that some of the structure associated with well-formedness remains intact even after preferences have been resolved. Moreover, our semantic and computational analysis shows that, for some of the preference reductions, advantageous properties associated with well-formedness can still be guaranteed in view of preferences.

In choice logics, we study the important notion of preferred model entailment with regards to logical, computational, and proof-theoretic properties. To this end, we consider Qualitative Choice Logic (QCL), Conjunctive Choice Logic (CCL), and Lexicographic

Choice Logic (LCL), as well as several preferred model semantics, i.e., methods of determining the preferred models of a choice logic theory. We prove that preferred model entailment for choice logics satisfies key logical properties for non-montonic entailment. Our results also show that the computational complexity of preferred model entailment is located on the second level of the poylnomial hierarchy, with the exact complexity depending both on the choice logic and preferred model semantics. Moreover, we introduce the first sequent calculi for preferred model entailment in choice logics and prove soundness and completeness.

Finally, we investigate the relationship between choice logics and abstract argumentation by translating QCL-theories to Argumentation Frameworks with Collective Attacks (SETAFs). We prove that the original QCL-theory is in semantic correspondence to the constructed SETAF, and we show that our translation can be used to decide preferred model entailment in QCL. Moreover, we argue that our construction has advantages compared to an already existing translation from Prioritized QCL-theories to Value-based AFs since it is purely syntactic and polynomial in both size and runtime.

Contents

CHAPTER 1

Introduction

Storing knowledge in computer systems and letting computers draw conclusions from available knowledge is a key challenge in modern artificial intelligence (AI) (Russell and Norvig 2021). Indeed, *knowledge representation and reasoning* (KR) is an important subfield of AI. In KR, knowledge is represented explicitly and "in a symbolic way" (Flasiński 2016, p. 15), e.g., using a logical language or other formal objects such as graphs, and a method of reasoning with this knowledge is described formally. The study of KR has lead to the development of many influential concepts and formalisms including non-monotonic reasoning (Kraus, Lehmann, and Magidor 1990), answer set programming (Lifschitz 2019), description logics (Baader et al. 2017), and formal argumentation (Baroni et al. 2018). Moreover, KR has contributed to other fields of AI such as natural language processing (Bhattarai, Granmo, and Jiao 2023) or cognitive robotics (Paulius and Sun 2019). Understanding and advancing KR is crucial for designing robust AI systems that are predictable, explainable, and capable of logical reasoning (Sheth, Roy, and Gaur 2023). This thesis contributes hereto with the study of two KR-formalisms, namely choice logics (Brewka, Benferhat, and Berre 2004) and abstract argumentation with preferences (Kaci et al. 2021). Both formalisms are designed to represent preferences and reason with them, and moreover share the commonality that they closely interlink the concept of *preferences* with that of *truth*.

Preferences are ubiquitous in our everyday lives and guide human decision making in various situations (Hausman 2011). Both small decisions, such as which movie to watch, and big decisions, such as which career path to pursue, require us to make a choice between several alternatives based on our preferences. Consequently, the notion of preferences has been studied in many research areas including philosophy, economics, sociology, but also computer science and AI (Pigozzi, Tsoukiàs, and Viappiani 2016). For instance, recommender systems fulfill tasks such as suggesting songs or movies to users based on their explicit or implicit preferences (Bobadilla et al. 2013) and computational social choice considers various ways of aggregating the conflicting preferences of several individuals into a single fair outcome (Brandt et al. 2016). In KR, preferences often occur naturally as part of the knowledge to be represented. For example, we may know that a certain user prefers horror movies to documentaries, or that an afternoon kick-off for a football game is preferable to a morning kick-off, or that a certain argument in a debate is considered stronger than another argument. Of course, knowledge about preferences is different in character from knowledge about truth, i.e., the hard facts about the world. To distinguish these concepts, we will therefore refer to knowledge that deals with preferences as soft-constraints, while knowledge about truth will be referred to as hard-constraints. Given the ubiquity of preferences, it is important to develop and study systems that are capable of faithfully representing and resolving not only hard-constraints but also soft-constraints.

How exactly preferences are represented and resolved differs from formalism to formalism. We identify two main paradigms of preference handling in KR, namely those of *separated* and *integrated* preferences.

In systems adhering to the paradigm of separated preferences, hard- and soft-constraints are specified and resolved independently of each other (Alfano et al. 2022; Brewka, Niemelä, and Truszczynski 2003; Faber, Truszczyński, and Woltran 2013). Preferences in such systems are typically defined "on top" of an underlying formalism used for representing hard-constraints, such as classical logic. Reasoning about the given knowledge then occurs in two steps: first, the possible world views are established in accordance with the given hard-constraints; then, the most preferred world views are filtered out in accordance with the given soft-constraints. In these systems, hard- and soft-constraints do not interact directly and are clearly separated. While this allows for a high degree of modularity, it means that knowledge about truth and preferences are not represented in a single language.

In contrast to this are formalisms adhering to the paradigm of integrated preferences, where hard- and soft-constraints are jointly represented and/or resolved. Both choice logics and abstract argumentation with preferences, the two formalisms that are the focus of this thesis, belong to this paradigm. In choice logics, propositional logic is extended with an additional connective with which preferences can be expressed. In abstract argumentation, reasoning with preferences is often not carried out in two separate steps as described above. Instead, resolving preferences has a direct influence on the given hard-constraints. While integrated preferences allow us to talk about hard- and soft-constraints in a single unified language, the interactions between hard- and soft-constraints can impact syntactic, semantic, and computational properties in different ways than separated preferences. In this thesis, we study choice logics and preferences in abstract argumentation in order to better understand the notion of integrated preferences, the challenges they bring with them, and how they can be overcome.

The remainder of this chapter is organized as follows: in Section 1.1 we explore the concept of integrated preferences in more detail. In Section 1.2 we list the main contributions of this thesis, and in Section 1.3 we discuss important related work. Section 1.4 gives a brief outline and reading guide to the thesis. Section 1.5 lists the publications of the author.

1.1 Integrated Preferences

As discussed above, in formalisms featuring integrated preferences, hard-constraints (truth) and soft-constraints (preferences) are represented and/or resolved together, not separately. We now explore two formalisms that clearly demonstrate this concept, namely choice logics, where hard- and soft- constraints are jointly *represented*, and abstract argumentation with preferences, where hard- and soft-constraints are jointly *resolved*.

Qualitative Choice Logic (QCL) (Brewka, Benferhat, and Berre 2004) extends classical propositional logic with a non-classical connective $\overrightarrow{\times}$ called ordered disjunction. Intuitively, $F \overrightarrow{\times} G$ means that it is preferable to satisfy F but, if that is not possible, satisfying G is also acceptable. Thus, QCL allows to express both truth (F or G must be satisfied) and preferences (satisfying F is better than G) in the same language. As an example for how QCL can model situations that naturally occur in a knowledge representation context, consider a product configuration system for cars (Gençay, Schüller, and Erdem 2019). Assume a user of the system is interested in the sports version of the car and wants an automatic transmission. Moreover, the user would like cruise control or lane assist, but cruise control is the preferred option. Such a query could be formalized in QCL:

$$\mathit{sport} \land \mathit{automatic} \land (\mathit{cruise} \overrightarrow{\times} \mathit{lane}).$$

There are three interpretations that satisfy the hard-constraints of the above formula, i.e., there are three models of this formula, namely $M_1 = \{\mathit{sport}, \mathit{automatic}, \mathit{cruise}, \mathit{lane}\}$, $M_2 = \{\mathit{sport}, \mathit{automatic}, \mathit{cruise}\}$, and $M_3 = \{\mathit{sport}, \mathit{automatic}, \mathit{lane}\}$. According to QCL-semantics, the models M_1 and M_2 containing the first option *cruise* in the ordered $\mathit{cruise} \overrightarrow{\times} \mathit{lane}$ are the preferred models of this formula. Suppose now that we obtain the additional information that it is not possible to have the sport version of the car with cruise control. This situation can be modeled by adapting the above formula as follows:

$$\mathit{sport} \land \mathit{automatic} \land (\mathit{cruise} \overrightarrow{\times} \mathit{lane}) \land \neg(\mathit{sport} \land \mathit{cruise}).$$

M_3 is the only model, and therefore also the preferred model, of the adapted formula. This example shows that the notion of preferred models in QCL is non-monotonic, as adding new information resulted in a completely new preferred model. Another interesting feature of QCL is that it allows us to express conditional preferences with as little as a single formula. For instance, we can formalize that we prefer cruise control over a lane assistant, but only if we choose an automatic transmission.

$$\mathit{automatic} \rightarrow (\mathit{cruise} \overrightarrow{\times} \mathit{lane}).$$

The preference expressed by $\mathit{cruise} \overrightarrow{\times} \mathit{lane}$ now only applies to cars with automatic transmission. The ability of QCL to express conditional preferences brings with it many possibilities when attempting to model real world situations accurately. On the other hand, the interplay between hard- and soft-constraints means that the properties of QCL, whether they are semantic, computational, or proof-theoretic, must be investigated with care and by taking the unique characteristics of choice logics into account.

(a) Example AF. (b) The preference $b \succ c$ causes an attack reversal.

Figure 1.1: Resolving preferences in abstract argumentation.

We have seen above how choice logics allow us jointly represent hard- and soft-constraints in a single unified language. We now shift our attention to another KR-formalism featuring integrated preferences, namely formal argumentation. Abstract Argumentation Frameworks (AFs) (Dung 1995) are a popular formalism used to find justifiable, consistent world views when facing conflicting or inconsistent information. As the name suggests, arguments in AFs are abstract, which means that we are not particularly interested in the content or internal structure of arguments, but rather in the relationship between the arguments. Specifically, an argument can *attack* another argument, resulting in a conflict between the two. As an example, consider the AF shown in Figure 1.1a, with arguments a, b, and c, as well as an attack from a to b, from b to a, and from c to b. Note that b is attacked by both a and c. While b attacks a and thus defends itself against a, it does not defend itself against c. Thus, in most argumentation semantics, $\{b\}$ is not an acceptable world-view. Indeed, the only stable argument set[1] in this framework is $\{a, c\}$.

AFs have been extended in several ways in the literature, including with preferences. Specifically, preference-based AFs (PAFs) (Kaci et al. 2021) consist of an AF and a preference order $\succ$ over arguments. If $b \succ c$ holds we say that the argument b is stronger than the argument c. Notice that, in contrast to choice logics, hard- and soft-constraints are represented completely separately in this case: hard-constraints are represented by the AF and soft-constraints by the preference order $\succ$. However, hard- and soft-constraints interact when preferences are *resolved*. Specifically, the semantics of PAFs are given relative to so-called preference reductions which modify the attack relation based on the preference ordering. One such method of resolving the preference $b \succ c$ in the AF of Figure 1.1a is to revert the attack from the weaker argument c to the stronger argument b, i.e., to delete the attack from c to b and add an attack from b to c. The AF that results from this procedure is shown in Figure 1.1b. The stable argument sets of this AF are $\{a, c\}$ and $\{b\}$, which is different from before. One can see that, while hard- and soft-constraints are represented separately in PAFs, they are not independent of each other when it comes to *reasoning*. Thus, the approach of using preference reductions to resolve preferences is another example of integrated preferences in KR.

Note that the interplay of hard- and soft-constraints in abstract argumentation brings some challenges with it. For instance, the exact method of how preferences should be resolved has to be chosen with care. In the literature, three other preference reductions besides reverting attacks have been described, with one option being to simply delete attacks that go against the preference order. The choice of method has implications on

[1] A set S of arguments is considered stable if it is conflict-free and every argument in the framework is either contained in S or attacked by S.

the reasoning outcome. Another challenge is that resolving preferences by deleting or reverting attacks changes the *structure* of the underlying AF, therefore possibly impacting more general properties, be they semantic or computational. Since many advantageous properties in argumentation hold only for frameworks with a certain structure (Dvořák and Woltran 2020; König 2020), this means that we need to carefully investigate the impact of preferences and whether advantageous properties can still be guaranteed in view of preferences.

1.2 Contributions

From a high level perspective, this thesis contributes to the field of KR and therefore AI by studying KR-formalisms that deal with integrated preferences. Specifically, we investigate abstract argumentation using preference-reductions and choice logics with their joint representation of hard- and soft-constraints. In abstract argumentation we examine conditional preferences as well as the effect of preferences on claim-based reasoning in argumentation. For choice logics we consider the crucial notion of preferred model entailment and study it in detail. We investigate these formalisms with regards to semantic, computational, and syntactic properties. Our results provide novel insights into these formalisms, show the various impacts that integrated preferences can have, and allow for informed decisions when designing systems based on these formalisms. Moreover, we formally study the relationship between choice logics and abstract argumentation, expanding on previous work in this direction from the literature. We now present our contributions in more detail, chapter by chapter.

In Chapter 3 we introduce Conditional PAFs (CPAFs), a novel formalism that can deal with conditional preferences in abstract argumentation. CPAFs allow not only to incorporate but also to *talk about* preferences, and therefore enable us to directly represent conditional preferences. We study CPAFs with regards to semantic principles that have been investigated for PAFs (with unconditional preferences) in the literature, and find that most, but not all, of the principles satisfied in the unconditional case are also satisfied for CPAFs. Moreover, we study the complexity of CPAFs and find that, in contrast to the unconditional preferences in PAFs, conditional preferences can lead to an increase in computational complexity. Lastly, we compare our CPAFs with related formalisms from the literature, namely Value-based AFs (Atkinson and Bench-Capon 2021; Bench-Capon, Doutre, and Dunne 2007) and Extended AFs (Modgil 2009).

In Chapter 4 we investigate the effect of preferences on Claim-based AFs (CAFs), a formalism from the literature that associates each argument in a framework with a claim/conclusion and therefore allows for claim-centric reasoning in abstract argumentation (Dvořák and Woltran 2020). To this end, we introduce Preference-based CAFs (PCAFs) and study them with regards to syntactic, semantic, and computational properties. Our syntactic analysis shows that resolving preferences changes the underlying structure of the given framework in such a way that the central syntactic property of well-formedness (arguments with the same claim have the same outgoing attacks) can no longer be guaranteed. Since well-formedness is associated with advantageous semantic and computational properties, we study the frameworks that result by resolving preferences and investigate whether these advantageous properties can still be guaranteed. Regarding semantic properties, we find that whether a property is preserved in view of preferences largely depends on how exactly preferences are resolved, i.e., on which preference reduction is chosen. Regarding computational properties, we show that the low complexity of well-formed CAFs is preserved in most, but not all, cases.

In Chapter 5 we examine the notion of preferred model entailment in choice logics, where a choice logic theory (a set of choice logic formulas) T entails a classical formula F if and only if F is true in all preferred models of T. When doing so, we have to consider two axes of generalization. The first axis is the considered choice logic: while QCL is the most prominent choice logic in the literature, it is not the only one; other examples are Conjunctive Choice Logic (CCL) (Boudjelida and Benferhat 2016) and Lexicographic Choice Logic (LCL) (Bernreiter 2020), both of which replace the ordered disjunction of QCL with alternative choice connectives. The second axis is the preferred model semantics, i.e., the method of determining the preferred models of a choice logic theory. Here we mainly consider three approaches from the literature, namely the lexicographic, inclusion-based, and minmax approaches. We show that preferred model entailment under the lexicographic and inclusion-based approaches satisfy key principles for entailment in non-monotonic logics introduced by Kraus, Lehmann, and Magidor (1990), while the inclusion-based approach does not satisfy the property of rational monotony. Moreover, we study the complexity of preferred model entailment and show that it (1) lies on the second level of the polynomial hierarchy and (2) depends both on which choice logic is used (i.e. QCL, CCL, LCL,. . .) and on which preferred model semantics is considered (i.e. lexicographic, inclusion-based, minmax,. . .). Lastly, we introduce a sound and complete sequent calculus for the three preferred model semantics mentioned above, and for the logics QCL, CCL, and LCL.

Finally, in Chapter 6 we study the relationship between choice logics and abstract argumentation by providing a translation from QCL-theories to AFs with collective attacks (SETAFs) (Nielsen and Parsons 2006). To this end we consider both the inclusion-based and minmax approaches of determining preferred models. In contrast to an already existing translation (Sedki 2015) from choice logic theories (using the lexicographic approach) to Value-based AFs (Atkinson and Bench-Capon 2021; Bench-Capon, Doutre, and Dunne 2007), our translation is purely syntactic and polynomial in size.

1.3 Related Work

In this section we outline important generally related work to this thesis. Note that a more detailed discussion on related work specific to each chapter are contained in the chapters themselves.

In abstract argumentation, the reduction-based approach to preferences that we use in this thesis has been studied before for standard AFs (Kaci et al. 2021). Moreover, it has been employed in several extensions of AFs, most notably Value-based AFs (Atkinson and Bench-Capon 2021; Bench-Capon, Doutre, and Dunne 2007) and Extended AFs (Baroni et al. 2009; Modgil 2009). In Section 3.4 we investigate the relationship between these formalisms and our CPAFs in detail.

An alternative approach to preferences in argumentation, which belongs to the paradigm of separated preferences, is the lifting-based approach in which the given preference ordering over arguments is lifted to an ordering over argument sets (Alfano et al. 2022, 2023; Amgoud and Vesic 2014; Brewka, Truszczynski, and Woltran 2010; Kaci, van der Torre, and Villata 2018). The properties that this alternative approach to preferences induces are quite different to the properties of the reduction-based approach we use in this thesis. This concerns both semantic (Kaci, van der Torre, and Villata 2018) as well as computational properties (Alfano et al. 2022), with the complexity of reasoning tasks being generally higher under the lifting-based approach.

Preferences have also been used extensively in structured argumentation, where in contrast to abstract argumentation the internal structure of individual arguments is taken into consideration. Popular formalisms such as ABA+ (Cyras and Toni 2016) and ASPIC (Modgil and Prakken 2013) both feature preferences, and deal with them akin to the reduction-based approach used in this thesis.

Regarding choice logics, we want to highlight some applications that have been discussed in the literature, ranging from logic programming (Brewka, Niemelä, and Syrjänen 2004) to alert correlation (Benferhat and Sedki 2010) to database querying (Liétard, Hadjali, and Rocacher 2014) or preference learning (Sedki, Lamy, and Tsopra 2020, 2022).

Besides choice logics, there are other logic-based formalisms that jointly represent hard- and soft-constraints. This includes the lexicographic logic of Charalambidis et al. (2021), in which lists of truth values are used to rank interpretations. Maly and Woltran (2018) propose an alternative semantics for the language of QCL. Logic programs with ordered disjunction (LPODs) (Brewka, Niemelä, and Syrjänen 2004; Charalambidis, Rondogiannis, and Troumpoukis 2021; Delgrande, Schaub, and Tompits 2003) feature ordered disjunction in the head of rules. Another example, and one of the first logic-based formalisms dealing with preferences, is the *preference logic* by von Wright (Liu 2010; von Wright 1963). The focus of this logic, however, lies in reasoning *about* preferences rather than representing them and reasoning *with* them.

As observed in the literature (Brewka, Niemelä, and Truszczynski 2008; Shoham 1987), non-monotonic reasoning is inherently connected to preferences. Indeed, the notion of defeasible inference can be viewed as a form of preferential reasoning: the state of the world that we can defeasibly infer, i.e., the default state, is the preferred state; if new information comes to light that forces us to retract our conclusion and abandon the default state and accept an alternative, less preferred, state. As such, preferred model entailment for choice logics is closely related to the seminal work by Kraus, Lehmann, and Magidor (1990) and entailment in other non-monotonic logics such as circumscription (McCarthy 1980), default-logic (Reiter 1980), or autoepistemic logic (Moore 1985). In contrast to these formalisms, however, choice logics are explicitly designed to represent preferences. Moreover, choice logics are a representative of integrated preferences, whereas, e.g., circumscription, clearly separates hard-constraints (represented as propositional formulas) from soft-constraints (represented via the circumscription-policy).

Conditional preferences have been studied in a variety of settings, and many formalisms have been introduced to this end. Maybe the most prominent example is CP-nets (Boutilier et al. 2004), with conditional preferences being represented as graphs. In logic programming, conditional preferences can occur in the head of rules or as dedicated preference statements (Brewka et al. 2015). In argumentation, conditional preferences have been studied in structured argumentation (Dung, Thang, and Son 2019) and recently also in abstract argumentation (Alfano et al. 2023), although this approach uses the lifting-based approach of handling preferences rather than the reduction-based approach employed in this thesis.

Of course, preferences also play an important role in areas of AI other than KR. We have already mentioned recommender systems (Bobadilla et al. 2013) and computational social choice (Brandt et al. 2016) as examples to this end. For general survey papers on preferences in AI we refer to Bienvenu, Lang, and Wilson (2010), Domshlak et al. (2011), and Pigozzi, Tsoukiàs, and Viappiani (2016).

1.4 Reading Guide

While this thesis can certainly be read from start to finish, this is by no means necessary. We now provide a reading guide for readers that wish to go out of order by specifying the reading requirements for each chapter. Figure 1.2 visualizes this guide.

It is recommended that readers first go through Chapter 1 (the introduction) and at least the first two sections of Chapter 2 (the preliminaries), i.e., Section 2.1 (on propositional logic) and Section 2.2 (on complexity theory).

Before reading Chapter 3 (on conditional preferences in abstract argumentation) it is recommended to read Subsection 2.3.1 (on abstract argumentation) and Subsection 2.3.2 (on preferences in abstract argumentation) in addition to the generally recommended preliminaries (Section 2.1 and Section 2.2).

Before reading Chapter 4 (on preferences in claim-based argumentation) it is recommended to read all of Section 2.3, which, in contrast to the recommended preliminaries for Chapter 3, includes Subsection 2.3.3 (on claim-based argumentation).

Before reading Chapter 5 (on preferred model entailment in choice logics) it is recommended to read Section 2.4 (on choice logics) in addition to the generally required preliminaries (Section 2.1 and Section 2.2).

Finally, before reading Chapter 6 (on the connection between argumentation and choice logics) it is recommended to read at least Subsection 2.3.1 (on abstract argumentation) as well as Chapter 5 (on preferred model entailment in choice logics).

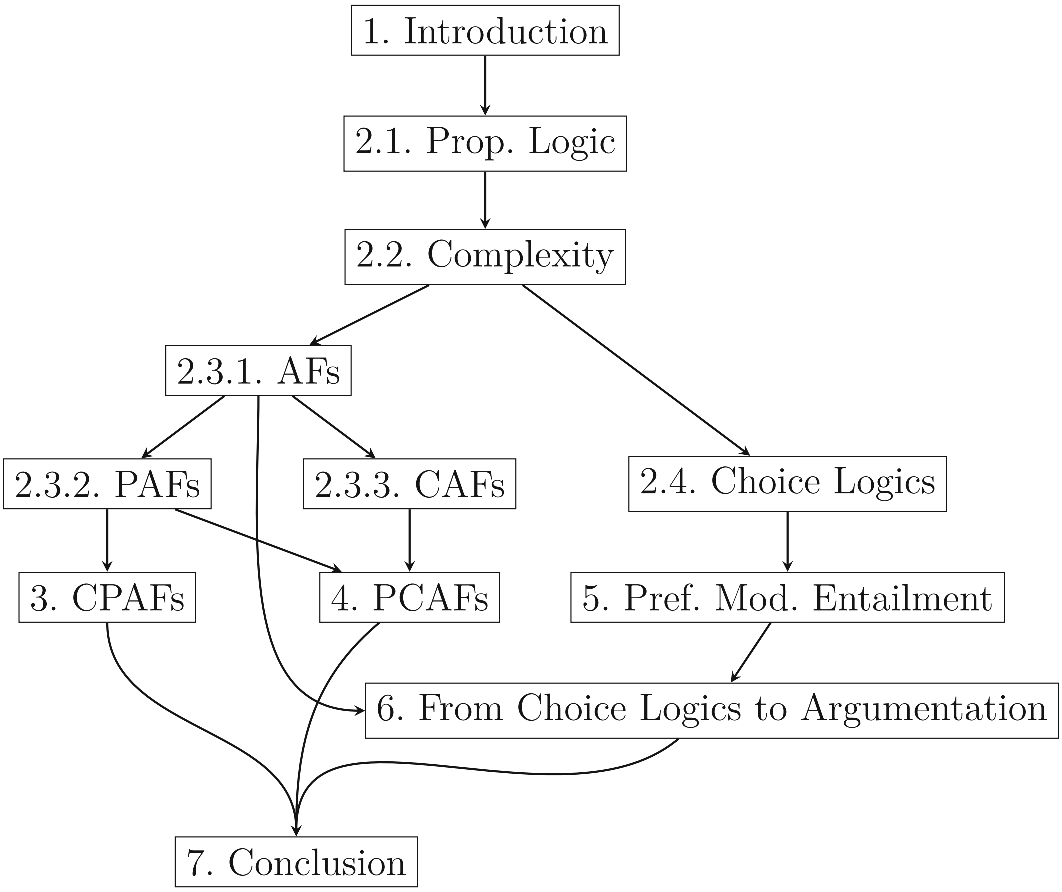

Figure 1.2: Reading guide for this thesis.

Chapter 7 (the conclusion) is best read after reading through all previous chapters.

1.5 Publications

In this section we list the publications used as a basis for this thesis, as well as further publications co-authored by the author.

Chapter 3 of this thesis is based on the following two publications (the second paper is an extended version of the first paper).

- Bernreiter, M., Dvořák, W., and Woltran, S. (2022). „Abstract Argumentation with Conditional Preferences“. In: *Proc. COMMA'22*. Vol. 353. Frontiers in Artificial Intelligence and Applications. IOS Press, pp. 92–103
- Bernreiter, M., Dvořák, W., and Woltran, S. (2023). „Abstract Argumentation with Conditional Preferences“. In: *Argument & Computation* pre-press, pp. 1–29

Chapter 4 is based on the following two publications (the second paper is a revised and shortened version of the first paper). An extended journal publication of these papers is currently under submission.

- Bernreiter, M., Dvořák, W., Rapberger, A., and Woltran, S. (2022a). „The Effect of Preferences in Abstract Argumentation Under a Claim-Centric View“. In: *Proc. NMR'22*. Vol. 3197. CEUR Workshop Proceedings. CEUR-WS.org, pp. 27–38
- Bernreiter, M., Dvořák, W., Rapberger, A., and Woltran, S. (2023). „The Effect of Preferences in Abstract Argumentation under a Claim-Centric View“. In: *Proc. AAAI'23*. AAAI Press, pp. 6253–6261

Chapter 5 is based on the following three publications (the third paper is an extended version of the second paper).

- Bernreiter, M., Maly, J., and Woltran, S. (2022). „Choice logics and Their Computational Properties“. In: *Artif. Intell.* 311, p. 103755
- Bernreiter, M., Lolic, A., Maly, J., and Woltran, S. (2022b). „Sequent Calculi for Choice Logics“. In: *Proc. IJCAR'22*. Springer, pp. 331–349
- Bernreiter, M., Lolic, A., Maly, J., and Woltran, S. (2024a). „Sequent Calculi for Choice Logics“. In: *J. Autom. Reason. (to appear)*

Chapter 6 is based on the following publication.

- Bernreiter, M. and König, M. (2023). „From Qualitative Choice Logic to Abstract Argumentation“. In: *Proc. KR'23*, pp. 737–741

Additionally, the author of this thesis has co-authored the following publications which are, however, not part of the contributions of this thesis.

- Bernreiter, M., Maly, J., and Woltran, S. (2020). „Encoding Choice Logics in ASP“. in: *Proc. ASPOCP@ICLP'20*. Vol. 2678. CEUR Workshop Proceedings. CEUR-WS.org
- Bernreiter, M. (2020). „A General Framework for Choice Logics“. Master's thesis. Vienna, Austria: TU Wien
- Bernreiter, M., Maly, J., and Woltran, S. (2021). „Choice Logics and Their Computational Properties“. In: *Proc. IJCAI'21*. IJCAI Organization, pp. 1794–1800

- Freiman, R. and Bernreiter, M. (2023a). „Truth and Preferences - A Game Approach for Qualitative Choice Logic“. In: *Proc. JELIA'23*. Vol. 14281. LNCS. Springer, pp. 547–560
- Freiman, R. and Bernreiter, M. (2023b). „Validity in Choice Logics - A Game-Theoretic Investigation“. In: *Proc. WoLLIC'23*. Vol. 13923. LNCS. Springer, pp. 211–226
- Bernreiter, M., Maly, J., Nardi, O., and Woltran, S. (2024b). „Combining Voting and Abstract Argumentation to Understand Online Discussions“. In: *Proc. AAMAS'24 (to appear)*

Note that (Bernreiter, Maly, and Woltran 2022), which is used in Chapter 5 of this thesis, is an extended version of (Bernreiter, Maly, and Woltran 2021), which in turn is based on the master thesis (Bernreiter 2020) of the author. The results from (Bernreiter, Maly, and Woltran 2022) that are used in Chapter 5 of this thesis are concerned with preferred model entailment in choice logics, and do not appear in (Bernreiter, Maly, and Woltran 2021) or (Bernreiter 2020). To clearly separate the results of the author's master thesis from the results newly presented in this thesis, the results already established in the master thesis are contained solely in the preliminaries of this thesis (see Section 2.4).

CHAPTER 2

Preliminaries

In this chapter we introduce central notions and results from the literature that are used throughout this thesis. Section 2.1 briefly recalls the basics of propositional logic. In Section 2.2 we define the complexity classes used in this thesis. Section 2.3 covers abstract argumentation, starting with standard Argumentation Frameworks (AFs) (Subsection 2.3.1) before detailing Preference-based AFs (Subsection 2.3.2) and Claim-augmented AFs (Subsection 2.3.3). Finally, in Section 2.4 we define the general choice logic framework and the most prominent choice logics considered in the literature.

Note that basic set theory or graph theory are not covered in this chapter. We assume familiarity with these concepts.

2.1 Propositional Logic

Classical propositional logic is one of the most fundamental knowledge representation languages there is (Russell and Norvig 2021). In this section, we briefly recall the most fundamental notions. For a more thorough overview, we can recommend (Hodel 2013). The syntax of propositional logic is based on (propositional) variables as well as logical connectives. The variables, also called atoms, represent basic statements such as "It is raining" or "I have an umbrella" that can, from a logical perspective, not be subdivided further. The connectives allow us to join together propositions to build logical formulas, also called sentences, such as "It is raining *and* I have an umbrella". In this thesis we use negation (denoted by $\neg$), conjunction (denoted by $\wedge$), and disjunction (denoted by $\vee$). Other connectives such as material implication are not explicitly part of the language as we define it, but can be modeled using the basic connectives $\neg, \wedge, \vee$ since they form a functionally complete set of connectives. Propositional variables will be denoted by lower-case letters such as a, b, c, x, y, z whereas formulas will be denoted by upper-case latin letters such as F, G, H or greek letters such as φ, ψ.

Definition 2.1. *Let $\mathcal{U}$ denote the (countably infinite) set of propositional variables (also called atoms). The set $\mathcal{F}_{\text{PL}}$ of formulas (over $\mathcal{U}$) of propositional logic is defined inductively:*

1. *if $a \in \mathcal{U}$, then $a \in \mathcal{F}_{\text{PL}}$;*
2. *if $F \in \mathcal{F}_{\text{PL}}$, then $(\neg F) \in \mathcal{F}_{\text{PL}}$;*
3. *if $F, G \in \mathcal{F}_{\text{PL}}$, then $(F \wedge G) \in \mathcal{F}_{\text{PL}}$;*
4. *if $F, G \in \mathcal{F}_{\text{PL}}$, then $(F \vee G) \in \mathcal{F}_{\text{PL}}$.*

The semantics of propositional logic assign a truth value to formulas given a specific interpretation (world view). Interpretations set each propositional variable either to true (represented by 1) or false (represented by 0). Formally, an interpretation $\mathcal{I}$ is a set of propositional variables, i.e., $\mathcal{I} \subseteq \mathcal{U}$. If $x \in \mathcal{I}$ for some variable x then x is true under $\mathcal{I}$, and if $x \notin \mathcal{I}$ then x is false under $\mathcal{I}$. The truth value of a non-atomic formula given an interpretation depends on the connectives with which the formula is built, and the truth values of the formula's subformulas.

Definition 2.2. *The truth value of a* PL*-formula under an interpretation $\mathcal{I} \subseteq \mathcal{U}$ is given by the function $val\colon 2^{\mathcal{U}} \times \mathcal{F}_{\text{PL}} \to \{0,1\}$ with*

1. $val(\mathcal{I}, a) = \begin{cases} 1 & \text{if } a \in \mathcal{I} \\ 0 & \text{otherwise} \end{cases}$ *for every $a \in \mathcal{U}$;*
2. $val(\mathcal{I}, \neg F) = \begin{cases} 1 & \text{if } val(\mathcal{I}, F) = 0 \\ 0 & \text{otherwise;} \end{cases}$
3. $val(\mathcal{I}, F \wedge G) = min(val(\mathcal{I}, F), val(\mathcal{I}, G))$;
4. $val(\mathcal{I}, F \vee G) = max(val(\mathcal{I}, F), val(\mathcal{I}, G))$.

We also use the alternative notation $\mathcal{I} \models F$ for $val(\mathcal{I}, F) = 1$ and $\mathcal{I} \not\models F$ for $val(\mathcal{I}, F) = 0$. If $\mathcal{I} \models F$, we say that $\mathcal{I}$ satisfies F, and if $\mathcal{I} \not\models F$, then $\mathcal{I}$ does not satisfy F.

Example 2.3. *Consider the formula*

$$F = (((a \vee c) \wedge (b \vee c)) \wedge \neg(a \wedge b)).$$

There are eight interpretations relevant for F, namely all $\mathcal{I}$ such that $\mathcal{I} \subseteq \{a, b, c\}$. To satisfy F, an interpretation must satisfy $\neg(a \wedge b)$. Thus, we have $\mathcal{I} \not\models F$ for $\mathcal{I} \in \{\{a, b\}, \{a, b, c\}\}$. Moreover, since $(a \vee c)$ and $(b \vee c)$ have to be satisfied, we have $\mathcal{I} \not\models F$ for $\mathcal{I} \in \{\emptyset, \{a\}, \{b\}\}$. For the three remaining interpretations, i.e., $\mathcal{I} \in \{\{a, c\}, \{b, c\}, \{c\}\}$, we have $\mathcal{I} \models F$.

2.2 Complexity Theory

Computational complexity theory aims to classify computational problems with respect to how difficult it is to solve them on a computer. In this section we briefly recall basic notions of complexity theory before defining some complexity classes of the polynomial hierarchy. The main source for this section is (Arora and Barak 2009). Note that we assume familiarity with fundamental concepts of complexity theory such as O-notation, decision problems, reductions between decision problems, and the complexity classes P, NP, and coNP.

Often, Turing machines are used as a computational model in complexity theory. However, due to the strong form of the Church-Turing thesis [p. 26](Arora and Barak 2009), which states that "every physically realizable computation model can be simulated by a Turing machine with polynomial overhead", we can use algorithms written in modern programming languages instead. In fact, informal descriptions of algorithms can also be used, as long as it is clear how they can be implemented.

Recall that P is the class of problems solvable in polynomial time, i.e., $O(n^c)$ time for some constant c with respect to the input size n. NP encompasses the problems solvable in non-deterministic polynomial time, i.e., the problems solvable in polynomial time by a non-deterministic Turing machine. For informal algorithms, this is equivalent to using a guess and check procedure, where a certificate or witness can be guessed and then used in a following deterministic algorithm. The class coNP contains all problems that are the complement[2] of an NP-problem. An example for a problem in P is model checking, where we are given a propositional formula F and an interpretation $\mathcal{I}$ as input, and ask whether $\mathcal{I}$ satisfies F, i.e., whether $\mathcal{I} \models F$. A prototypical NP-problem is SAT, where we are given a propositional formula F and ask whether there is an interpretation $\mathcal{I}$ that satisfies F. Its complement is the coNP-problem UNSAT, where we are given a propositional formula F and ask whether there is *no* interpretation that satisfies F.

A crucial notion in complexity theory is that of hardness and completeness of decision problems for a given complexity class. A decision problem Q is $\mathcal{C}$-hard for a complexity class $\mathcal{C}$ if all problems in $\mathcal{C}$ can be reduced to Q, i.e., any instances of a problem in $\mathcal{C}$ can be transformed into an instance of Q, in polynomial time. Moreover, Q is $\mathcal{C}$-complete if Q is in $\mathcal{C}$ and $\mathcal{C}$-hard. Completeness is a crucial notion that allows us to identify the most difficult to solve problems in a complexity class, and therefore draws the boundaries between classes. Moreover, reducing a problem such as SAT that is known to be hard for a complexity class to another problem allows us to show hardness for this other problem. SAT is known to be NP-complete and UNSAT is known to be coNP-complete. It is famously not known whether the classes P, NP, and coNP are distinct from each other, but it is widely assumed to be the case (Aaronson 2017).

Next, we consider complexity classes that are located on the second level of the polynomial hierarchy, a hierarchy of complexity classes that has P, NP, and coNP as its base. We

[2] A decision problem Q' is the complement of another decision problem Q if for any instance I we have that I is a yes-instance of Q iff I is a no-instance of Q'.

do not define the entire polynomial hierarchy, as this is not needed for this thesis. We simply note that P, NP, and coNP are located on the first level of the hierarchy, and that the polynomial hierarchy contains an infinite number of levels. For further reading, we recommend (Arora and Barak 2009, p. 95). We define the problems on the second level of the polynomial hierarchy via algorithms that have access to a Sat-oracle, which is a function that takes an instance of Sat as input and decides whether this instance is a yes-instance or not in constant time. Note that, with access to such an oracle we can solve any problem in NP (or coNP) in polynomial time, as we can reduce our problem to Sat (or Unsat) in polynomial time and then make a single call to the Sat-oracle to solve our problem. Instead of the term Sat-oracle we will also use the term NP-oracle, as any other NP-complete problem can be used instead of Sat.

Definition 2.4. *A decision problem is*

- *in* Σ_2^{P} *if it can be decided in nondeterministic polynomial time by an algorithm with access to an* NP*-oracle (Arora and Barak 2009);*
- *in* Π_2^{P} *if its complement is in* Σ_2^{P} *(Arora and Barak 2009).*
- *in* Δ_2^{P} *if it can be decided in polynomial time by an algorithm with access to an* NP*-oracle (Krentel 1988);*
- *in* Θ_2^{P}*, also called* $\Delta_2^{\mathsf{P}}[O(\log n)]$*, if it can be decided in polynomial time by an algorithm which is allowed* $O(\log n)$ *number of calls to an* NP*-oracle, where* n *is the size of the input (Wagner 1990);*
- *in* $\Delta_2^{\mathsf{P}}[O(\log^2 n)]$ *if it can be decided in polynomial time by an algorithm which is allowed* $O(\log^2 n)$ *number of calls to an* NP*-oracle, where* n *is the size of the input (Castro and Seara 1992);*
- *in* DP *if it is the intersection of a problem in* NP *and a problem in* coNP *(Papadimitriou and Yannakakis 1982).*

A useful Σ_2^{P}-complete problem is $\mathrm{QBF}_{\exists}^2$, where we are given a Quantified Boolean Formula (QBF) of the form $\Phi = \exists x_1, \ldots, x_n \forall y_1, \ldots, y_m \varphi$ with φ being a propositional formula over variables $x_1, \ldots, x_n, y_1, \ldots, y_m$, and ask whether Φ is valid, meaning that there is an assignment to $x_1, \ldots, x_n$ such that for all assignments to $y_1, \ldots, y_m$ we satisfy φ; For Π_2^{P} we will make use of the analogously defined $\mathrm{QBF}_{\forall}^2$, where we are given a QBF of the form $\Phi \forall x_1, \ldots, x_n \exists y_1, \ldots, y_m \varphi$, see e.g. (Biere et al. 2021). In the case of $\mathrm{QBF}_{\forall}^2$, the propositional formula φ can be assumed to be in conjunctive normal form (CNF), i.e., $\varphi = C_1 \wedge \cdots \wedge C_k$ with $C_i = (z_1 \vee \cdots \vee z_l)$.

For Δ_2^{P} we use the LexMaxSat problem where we are given a propositional formula φ and an ordering $x_1 > \ldots > x_n$ over the variables in φ, and ask if x_n is true in the lexicographically largest model of φ (Creignou, Pichler, and Woltran 2018). The Θ_2^{P}-complete LogLexMaxSat problem and the $\Delta_2^{\mathsf{P}}[O(\log^2 n)]$-complete Log^{2}LexMaxSat

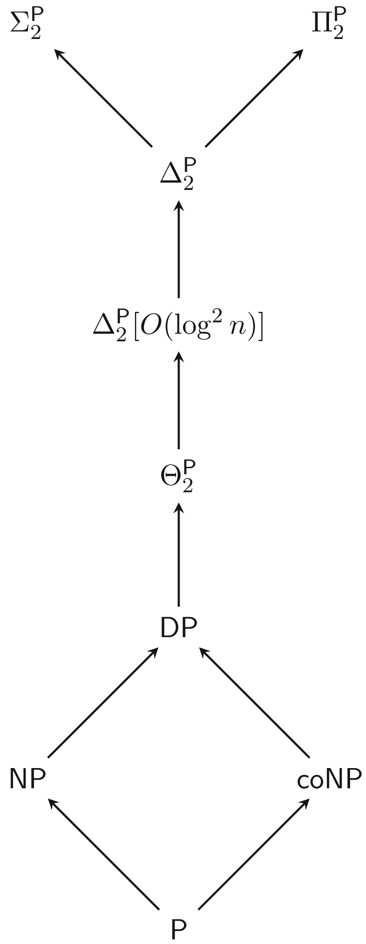

Figure 2.1: The complexity classes used in this thesis. An arrow indicates that the class from which the arrow originates is a subset of the class to which the arrow is pointing to.

problem are defined analogously, but we are given an ordering over only $\log(n)$ (resp. $\log(n)^2$) variables in φ (Creignou, Pichler, and Woltran 2018; Segoufin and ten Cate 2013).

Lastly, we will also make use of the DP-complete SatUnsat problem, where we are given two propositional formulas φ_1 and φ_2 and ask whether φ_1 is satisfiable and φ_2 is unsatisfiable (Papadimitriou and Yannakakis 1982).

Figure 2.1 depicts the complexity classes used in this thesis and how they are related.

2.3 Abstract Argumentation

Abstract argumentation is a simple yet powerful tool for modeling discussions and dealing with conflicting information (Baroni, Caminada, and Giacomin 2018). In Subsection 2.3.1

we recall the basics of abstract argumentation as introduced by Dung (1995). In Subsections 2.3.2 and 2.3.3 we then describe two generalizations of Argumentation Frameworks (AFs) that we will build upon in this thesis, namely Preference-based AFs, which provide the foundation needed to deal with preferences, and Claim-augmented AFs, which allow us to take not only arguments but also their claims/conclusions into account.

2.3.1 Abstract Argumentation Frameworks (AFs)

AFs were first introduced by Dung (1995). Arguments in AFs are abstract in the sense that we are not concerned with the internal structure of the argument themselves. Rather, we are interested in the relationship between arguments, which is modeled via *attacks* between arguments. If there is an attack between two arguments, then the arguments are in conflict and cannot be jointly accepted. Moreover, usually we require that an attacked argument must be *defended* against its attackers in order to be accepted.

Definition 2.5 (AF)**.** *An Argumentation Framework (AF) is a tuple $F = (A, R)$ where A is a finite set of arguments and $R \subseteq A \times A$ is an attack relation between arguments. $S \subseteq A$* attacks *$b \in A$ (in F) if $(a, b) \in R$ for some $a \in S$; $S^+_F = \{b \in A \mid \exists a \in S : (a, b) \in R\}$ denotes the set of arguments attacked by S. $S^\oplus_F = S \cup S^+_F$ is the* range *of S in F. An argument $a \in A$ is* defended *(in F) by S if $b \in S^+_F$ for each b with $(b, a) \in R$.*

Semantics for AFs are defined as functions σ which assign to each AF $F = (A, R)$ a set $\sigma(F) \subseteq 2^A$ of extensions (Baroni, Caminada, and Giacomin 2018). We consider for σ the functions *cf* (conflict-free), *adm* (admissible), *com* (complete), *grd* (grounded), *stb* (stable), *naive* (naive), *prf* (preferred), *sem* (semi-stable), and *stg* (stage).

Definition 2.6 (AF-semantics)**.** *Let $F = (A, R)$ be an AF. For a set $S \subseteq A$ it holds that*

- *$S \in cf(F)$ iff there are no $a, b \in S$, such that $(a, b) \in R$;*
- *$S \in naive(F)$ iff $S \in cf(F)$ and there is no $T \in cf(F)$ with $S \subset T$;*
- *$S \in stg(F)$ iff $S \in cf(F)$ and there is no $T \in cf(F)$ with $S^\oplus_F \subset T^\oplus_F$;*
- *$S \in adm(F)$ iff $S \in cf(F)$ and each $a \in S$ is defended by S in F;*
- *$S \in prf(F)$ iff $S \in adm(F)$ and there is no $T \in adm(F)$ with $S \subset T$;*
- *$S \in sem(F)$ iff $S \in adm(F)$ and there is no $T \in adm(F)$ with $S^\oplus_F \subset T^\oplus_F$;*
- *$S \in stb(F)$ iff $S \in cf(F)$ and each $a \in A \setminus S$ is attacked by S in F;*
- *$S \in com(F)$ iff $S \in adm(F)$ and each $a \in A$ defended by S in F is contained in S;*
- *$S \in grd(F)$ iff $S \in com(F)$ and there is no $T \in com(F)$ with $T \subset S$.*

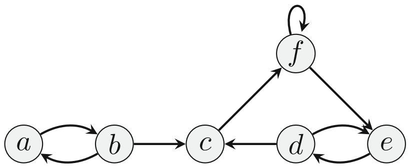

Figure 2.2: Example AF.

Let us provide a small example. AFs will be depicted as directed graphs, where the nodes are arguments and the edges are attacks between arguments.

Example 2.7. *Let $F = (A, R)$ be the AF depicted in Figure 2.2, i.e.,*

$$A = \{a, b, c, d, e, f\},$$
$$R = \{(a, b), (b, a), (b, c), (c, f), (d, c), (d, e), (e, d), (f, e), (f, f)\}.$$

Regarding conflict-free semantics, observe that, e.g., $\{a, b\} \notin cf(F)$ since $(a, b) \in R$. On the other hand, $\{a, c\} \in cf(F)$ since $(a, c) \notin R$, $(c, a) \notin R$, $(a, a) \notin R$, and $(c, c) \notin R$. Moreover, note that $\{f\} \notin cf(F)$ since the argument f is self-attacking, i.e., since $(f, f) \in R$.

$$cf(F) = \{\emptyset, \{a\}, \{b\}, \{c\}, \{d\}, \{e\},$$
$$\{a, c\}, \{a, d\}, \{a, e\}, \{b, d\}, \{b, e\}, \{c, e\}, \{a, c, e\}\}.$$

Since naive extensions are the subset-maximal conflict-free sets, we have

$$naive(F) = \{\{a, d\}, \{b, d\}, \{b, e\}, \{a, c, e\}\}.$$

Note that there is only one conflict-free set S such that $S^{\oplus}_F = A$, namely $S = \{a, c, e\}$. Thus, $\{a, c, e\}^{\oplus}_F \supset T^{\oplus}_F$ for all $T \in cf(F)$ such that $T \neq \{a, c, e\}$. Therefore,

$$stg(F) = stb(F) = \{\{a, c, e\}\}.$$

Regarding admissible semantics we have, e.g., $\{c\} \notin adm(F)$ since c does not defend itself against the attacks from b and d. However, $\{a, c, e\} \in adm(F)$ since a defends c and since c defends e. Overall we have

$$adm(F) = \{\emptyset, \{a\}, \{b\}, \{d\}, \{a, d\}, \{b, d\}, \{a, c, e\}\}.$$

The preferred semantics are the subset-maximal admissible sets, i.e.,

$$prf(F) = \{\{a, d\}, \{b, d\}, \{a, c, e\}\}.$$

Analogously to conflict-free sets, there is only one admissible set S such that $S^{\oplus}_F = A$, namely $S = \{a, c, e\}$. Thus,

$$sem(F) = stb(F) = \{\{a, c, e\}\}.$$

Table 2.1: Complexity of AFs (Dvořák and Dunne 2018).

σ	$Cred_\sigma^{AF}$	$Skept_\sigma^{AF}$	Ver_σ^{AF}
cf	in P	trivial	in P
adm	NP-c	trivial	in P
com	NP-c	P-c	in P
grd	P-c	P-c	P-c
stb	NP-c	coNP-c	in P
$naive$	in P	in P	in P
prf	NP-c	Π_2^{P}-c	coNP-c
sem	Σ_2^{P}-c	Π_2^{P}-c	coNP-c
stg	Σ_2^{P}-c	Π_2^{P}-c	coNP-c

As for complete semantics, we have $com(F) = adm(F)$ since no admissible set defends an argument outside of the set. Thus,

$$com(F) = \{\emptyset, \{a\}, \{b\}, \{d\}, \{a, d\}, \{b, d\}, \{a, c, e\}\}.$$

Lastly, the subset-minimal complete extensions is $\emptyset$, i.e.,

$$grd(F) = \{\emptyset\}.$$

AFs have been studied with regards to many properties such as semantic principles (van der Torre and Vesic 2018) and computational complexity (Dvořák and Dunne 2018) in the literature. Regarding semantic principles, we will consider preference-specific principles (see Subsection 2.3.2) and other properties such as I-maximality (see Definition 2.25 in Subsection 2.3.3). Regarding complexity, which will be a major focus in this thesis, the three central problems are those of credulous acceptance, skeptical acceptance, and verification. See below for their formal definition. Table 2.1 shows their complexity.

Definition 2.8. *Given an AF-semantics σ we define the following decision problems:*

- Credulous Acceptance *($Cred_\sigma^{AF}$): given an AF F and an argument x, is $x \in S$ for some $S \in \sigma(F)$?*
- Skeptical Acceptance *($Skept_\sigma^{AF}$): given an AF F and an argument x, is $x \in S$ in all $S \in \sigma(F)$?*
- Verification *(Ver_σ^{AF}): given an AF F and a set of arguments S, is $S \in \sigma(F)$?*

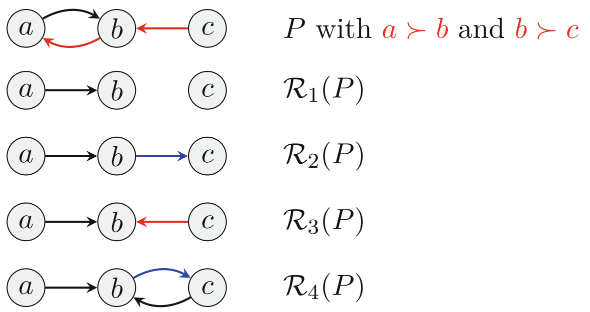

Figure 2.3: The four preference reductions used for PAFs.

2.3.2 Preference-based Argumentation Frameworks (PAFs)

Preference-based AFs constitute one way of introducing preferences, in the sense of variable argument strength, to abstract argumentation (Amgoud and Cayrol 1998, 2002; Kaci et al. 2021). Specifically, they generalize standard Dung-style AFs by introducing preferences between arguments.

Definition 2.9 (PAF)**.** *A Preference-based AF (PAF) is a triple* $P = (A, R, \succ)$ *where* (A, R) *is an AF and* $\succ$ *is an asymmetric preference relation over* A*.*

Notice that preferences in PAFs are not required to be transitive. While transitivity of preferences is often assumed in general and also in argumentation (Amgoud and Vesic 2014; Kaci, van der Torre, and Villata 2018), this is not entirely uncontroversial (Kaci et al. 2021; Schumm 1987). In this thesis, we do not assume transitivity but will consider the effect of transitive orderings when applicable.

If a and b are arguments and $a \succ b$ holds then we say that a is stronger than b (and that b is weaker than a). But what effect should this ordering have? How should this influence, e.g., the set of admissible arguments? One possibility is to remove all attacks from weaker to stronger arguments, and to then determine the set of admissible arguments in the resulting AF. This altering of attacks in a PAF based on its preference-ordering is called a reduction. The literature describes four such reductions for AFs (Kaci et al. 2021), which we now define.

Definition 2.10 (Preference reduction)**.** *Given a PAF* $P = (A, R, \succ)$*, the corresponding AF* $\mathcal{R}_i(P) = (A, R')$ *is constructed via Reduction* i*, where* $i \in \{1, 2, 3, 4\}$*, as follows:*

- $i = 1$*:* $\forall a, b \in A: \ (a, b) \in R' \Leftrightarrow (a, b) \in R, b \not\succ a$
- $i = 2$*:* $\forall a, b \in A: \ (a, b) \in R' \Leftrightarrow ((a, b) \in R, b \not\succ a) \vee ((b, a) \in R, (a, b) \notin R, a \succ b)$
- $i = 3$*:* $\forall a, b \in A: \ (a, b) \in R' \Leftrightarrow ((a, b) \in R, b \not\succ a) \vee ((a, b) \in R, (b, a) \notin R)$
- $i = 4$*:* $\forall a, b \in A: \ (a, b) \in R' \Leftrightarrow ((a, b) \in R, b \not\succ a) \vee ((b, a) \in R, (a, b) \notin R, a \succ b) \vee ((a, b) \in R, (b, a) \notin R)$

Figure 2.3 visualizes the four preference reductions. Intuitively, Reduction 1 removes attacks that contradict the preference ordering while Reduction 2 reverts such attacks. Reduction 3 removes attacks that contradict the preference ordering, but only if the weaker argument is attacked by the stronger argument also. Reduction 4 can be seen as a combination of Reductions 2 and 3: if a weak argument attacks a stronger argument, and there is no reverse attack, add a reverse attack but do not remove the attack from the weak to the strong argument; if a weak argument attacks a stronger argument, but there is a reverse attack, remove the attack from the weaker argument.

The semantics for PAFs are defined in a straightforward way: first, one of the four reductions is applied to the given PAF; then, AF-semantics are applied to the resulting AF.

Definition 2.11 (PAF-semantics)**.** *Let P be a PAF and let $i \in \{1, 2, 3, 4\}$. The preference-based variant of an AF-semantics σ relative to Reduction i is defined as $\sigma^i(P) = \sigma(\mathcal{R}_i(P))$.*

Example 2.12. *Consider the PAF P depicted in Figure 2.3, i.e., $P = (A, R, \succ)$ with*

$$A = \{a, b, c\},$$
$$R = \{(a, b), (b, a), (c, b)\},$$
$$a \succ b \text{ and } b \succ c.$$

The AFs $\mathcal{R}_1(P)$, $\mathcal{R}_2(P)$, $\mathcal{R}_3(P)$, and $\mathcal{R}_4(P)$ resulting from applying the various preference reductions are also depicted in Figure 2.3.

For Reduction 1 we have $cf^1(P) = cf(\mathcal{R}_1(P)) = \{\emptyset, \{a\}, \{b\}, \{c\}, \{a, c\}, \{b, c\}\}$ and thus $naive^1(P) = naive(\mathcal{R}_1(P)) = \{\{a, c\}, \{b, c\}\}$. Moreover, $adm^1(P) = adm(\mathcal{R}_1(P)) = \{\emptyset, \{a\}, \{c\}, \{a, c\}\}$ and therefore $prf^1(P) = prf(\mathcal{R}_1(P)) = \{\{a, c\}\}$.

For Reduction 2, on the other hand, we get $cf^2(P) = cf(\mathcal{R}_2(P)) = \{\emptyset, \{a\}, \{b\}, \{c\}, \{a, c\}\}$, i.e., $\{b, c\} \notin cf^2(P)$. Thus, $naive^2(P) = naive(\mathcal{R}_2(P)) = \{\{b\}, \{a, c\}\}$. Moreover, $adm^2(P) = adm(\mathcal{R}_2(P)) = \{\emptyset, \{a\}, \{a, c\}\}$, i.e., $\{c\} \notin adm^2(P)$. The preferred semantics are the same as under Reduction 1 in this case, i.e., $prf^2(P) = prf(\mathcal{R}_2(P)) = \{\{a, c\}\}$.

For Reductions 3 and 4 the conflict free sets are the same as under Reduction 2 in this case, i.e., $cf^2(P) = cf^3(P) = cf^4(P)$. The admissible sets, on the other hand, are the same as under Reduction 1, i.e., $adm^1(P) = adm^3(P) = adm^4(P)$.

PAFs have the same complexity as standard AFs (see Table 2.1) with respect to the decision problems of Definition 2.8: hardness results follow from the fact that PAFs generalize AFs, and membership results from the fact that the four preference reductions can be carried out in polynomial time.

A principle-based analysis of the four preference reductions was conducted for complete, grounded, preferred, and stable semantics (Kaci et al. 2021; Kaci, van der Torre, and Villata 2018). To this end, ten PAF-properties were laid out and investigated. We now recall them in Definitions 2.13, 2.14, 2.16, and 2.18 according to (Kaci et al. 2021).

Definition 2.13. *Let σ_p^i be a PAF-semantics. Let $\succ, \succ' \subseteq (A \times A)$ such that $\succ \cup \succ'$ is asymmetric.*

- *$P1$ (conflict-freeness): If $(x, y) \in R$ there is no $S \in \sigma_p^i(A, R, \succ)$ such that $\{x, y\} \subseteq S$.*
- *$P2$ (preference selects extensions 1): $\sigma_p^i(A, R, \succ \cup \succ') \subseteq \sigma_p^i(A, R, \succ)$.*
- *$P3$ (preference selects extensions 2): $\sigma_p^i(A, R, \succ) \subseteq \sigma_p^i(A, R, \emptyset)$.*

Intuitively, $P1$ states that if there is an attack between two arguments, then there is no extension containing both of them. $P2$ expresses that adding more preferences to a PAF can exclude extensions, but not introduce them. $P3$ states that this is in particular true if we add preferences to a framework without any preferences, i.e., $P3$ is a special case of $P2$.

Definition 2.14. *Let σ_p^i be a PAF-semantics. Let $\succ, \succ' \subseteq (A \times A)$ such that $\succ \cup \succ'$ is asymmetric.*

- *$P4$ (extension refinement): for all $S' \in \sigma_p^i(A, R, \succ \cup \succ')$ there is $S \in \sigma_p^i(A, R, \succ)$ such that $S \subseteq S'$.*
- *$P5$ (extension growth): $\bigcap(\sigma_p^i(A, R, \succ)) \subseteq \bigcap(\sigma_p^i(A, R, \succ \cup \succ'))$.*
- *$P6$ (number of extensions): $|\sigma_p^i(A, R, \succ \cup \succ')| \leq |\sigma_p^i(A, R, \succ)|$.*

$P4$ states that adding preferences means extensions will be supersets of extensions in the original PAF. $P5$ says that adding preferences will preserve skeptically accepted arguments, and might cause new arguments to be skeptically accepted. $P6$ expresses that the number of extensions will not grow if new preferences are added.

For the next two principles, we need to define the notion of an argument's status.

Definition 2.15. *Let $F = (A, R, \succ)$ be a PAF and $x \in A$. We write*

- *$status(x, F) = sk\text{-}cr$ iff x is skeptically and credulously accepted in F;*
- *$status(x, F) = cr$ iff x is credulously but not skeptically accepted in F;*
- *$status(x, F) = rej$ iff $status(x, F) \notin \{sk\text{-}cr, cr\}$.*

We define the order over theses statuses as follows: $sk\text{-}cr > cr > rej$.

Note that in stable semantics an argument is not always credulously accepted if it is skeptically accepted, since there are AFs without stable extensions. Thus, some argument x might be skeptically accepted with respect to stable semantics, yet we still might have $status(a, F) = rej$.

Definition 2.16. *Let σ_p^i be a PAF-semantics.*

- *$P7$ (status conservation): status$(x, (A, R, \succ \cup \{(x, y)\})) \geq$ status$(x, (A, R, \succ))$.*
- *$P8$ (preference-based immunity): if $(x, x) \notin R$ and $x \succ y$ for all $y \in A \setminus \{x\}$ then status$(x, (A, R, \succ)) \neq$ rej.*

If a semantics satisfies $P7$ then the status of an argument x cannot be lowered by adding a preference $x \succ y$ where x is the preferred (stronger) argument. $P8$ states that if an argument x is not self-attacking and also stronger than all other arguments, then x cannot be rejected.

For principles $P9$ and $P10$ we need the concept of paths between two arguments, by which we mean a path in the underlying undirected graph of a PAF.

Definition 2.17. *Let $F = (A, R, \succ)$ be a PAF. Let $R^- = \{(x, y) \mid (y, x) \in R\}$. There is a path between $x \in A$ and $y \in A$ iff there is a sequence of arguments $z_1, \ldots, z_n \in A$ such that $z_1 = x$, $z_n = y$, and $(z_k, z_{k+1}) \in R \cup R^-$ for all $1 \leq k < n$.*

Definition 2.18. *Let σ_p^i be a PAF-semantics.*

- *$P9$ (path preference influence 1): if there is no path from $x \in A$ to $y \in A$ in $(A, R, \succ)$ then $\sigma_p^i(A, R, \succ) = \sigma_p^i(A, R, \succ \cup \{(x, y)\})$.*
- *$P10$ (path preference influence 2): if $(x, y) \notin R$ and $(y, x) \notin R$ then $\sigma_p^i(A, R, \succ) = \sigma_p^i(A, R, \succ \cup \{(x, y)\})$.*

If $P9$ is satisfied then adding a preference between two arguments x and y that do not occur in the same weakly connected component does not change the extensions of a PAF. $P10$ is similar to $P9$, but only requires that there is no direct connection between arguments x and y.

Table 2.2 shows which semantics satisfy which principle, as investigated in (Kaci et al. 2021; Kaci, van der Torre, and Villata 2018).

2.3.3 Claim-augmented Argumentation Frameworks (CAFs)

CAFs generalize standard AFs by assigning a claim to each argument (Dvořák and Woltran 2020). The notion of enriching arguments with claims/conclusions appears often and under various names in the literature. For instance, Conclusion-based AF (Rocha and Cozman 2022a,b) are equivalent to CAFs as we consider them, while Argument-Conclusion Structures (Baroni, Governatori, and Riveret 2016) are not technically equivalent but strongly related to CAFs.

Table 2.2: Satisfaction of various PAF-principles (Kaci et al. 2021; Kaci, van der Torre, and Villata 2018). C stands for complete, G for grounded, P for preferred, and S for stable. $\times$ indicates that none of those four semantics satisfy this principle.

	$\mathcal{R}_1$	$\mathcal{R}_2$	$\mathcal{R}_3$	$\mathcal{R}_4$
$P1$ (conflict-freeness)	$\times$	$CGPS$	$CGPS$	$CGPS$
$P2$ (preference selects extensions)	$\times$	$\times$	CS	$\times$
$P3$ (preference selects extensions 2)	$\times$	$\times$	CS	$\times$
$P4$ (extension refinement)	$\times$	$\times$	CGS	$\times$
$P5$ (extension growth)	$\times$	$\times$	CG	$\times$
$P6$ (number of extensions)	G	G	$CGPS$	G
$P7$ (status conservation)	$CGPS$	$CGPS$	$CGPS$	$CGPS$
$P8$ (preference-based immunity)	CGP	CGP	$\times$	CPS
$P9$ (path preference influence 1)	$CGPS$	$CGPS$	$CGPS$	$CGPS$
$P10$ (path preference influence 2)	$CGPS$	$CGPS$	$CGPS$	$CGPS$

Definition 2.19 (CAF)**.** *A Claim-augmented AF (CAF) is a triple* $\mathcal{F} = (A, R, cl)$ *where* (A, R) *is an AF and* $cl\colon A \to \mathcal{C}$ *is a function that maps arguments to an infinite domain of claims* $\mathcal{C}$*. The claim-function is extended to sets of arguments via* $cl(S) = \{cl(a) \mid a \in S\}$*. A well-formed CAF (wfCAF) is a CAF* (A, R, cl) *in which all arguments with the same claim attack the same arguments, i.e., for all* $a, b \in A$ *with* $cl(a) = cl(b)$ *we have that* $\{c \mid (a, c) \in R\} = \{c \mid (b, c) \in R\}$*.*

Well-formed CAFs are an important subclass of CAFs that capture a natural behavior common to many structured argumentation formalisms and instantiations (Cyras and Toni 2016; Modgil and Prakken 2013), i.e., that all arguments with the same claim attack the same arguments. Moreover, wfCAFs enjoy advantages when it comes to semantic and computational properties, as we will see below.

There are two types of semantics for CAFs, inherited and hybrid. Inherited semantics apply AF-semantics to the underlying AF of a given CAF, and then collect the claims of arguments contained in an extension.

Definition 2.20 (Inherited semantics)**.** *Let* $\mathcal{F} = (A, R, cl)$ *be a CAF. The inherited CAF-variant of an AF-semantics* σ *is defined as* $\sigma_{inh}(\mathcal{F}) = \{cl(S) \mid S \in \sigma((A, R))\}$*.*

Example 2.21. *Let* $\mathcal{F} = (A, R, cl)$ *be the CAF depicted in Figure 2.4, i.e.,*

$$
\begin{aligned}
&A = \{a, b, c, d, e, f\},\\
&R = \{(a, b), (b, a), (b, c), (c, f), (d, c), (d, e), (e, d), (f, e), (f, f)\},\\
&cl(a) = cl(d) = \alpha,\\
&cl(b) = cl(e) = cl(f) = \beta,\\
&cl(c) = \gamma.
\end{aligned}
$$

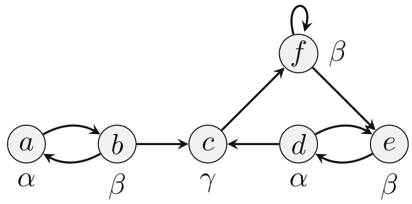

Figure 2.4: Example CAF.

Note that $\mathcal{F}$ is not well-formed, since, e.g., $(a, b) \in R$ but $(d, b) \notin R$ despite $cl(a) = cl(d)$.

The underlying AF (A, R) of $\mathcal{F}$ is the AF we examined in Example 2.7. The extensions of $\mathcal{F}$ on the claim-level can be inferred from the extensions of (A, R) on the argument-level (see Example 2.7). Thus, we have

$$\begin{aligned}
cf_{inh}(\mathcal{F}) &= \{\emptyset, \{\alpha\}, \{\beta\}, \{\gamma\}, \{\alpha,\gamma\}, \{\alpha,\beta\}, \{\beta,\gamma\}, \{\alpha,\beta,\gamma\}\},\\
naive_{inh}(\mathcal{F}) &= \{\{\alpha\}, \{\beta\}, \{\alpha,\beta\}, \{\alpha,\beta,\gamma\}\},\\
adm_{inh}(\mathcal{F}) &= com_{inh}(\mathcal{F}) = \{\{\emptyset, \{\alpha\}, \{\beta\}, \{\alpha,\beta\}, \{\alpha,\beta,\gamma\}\},\\
prf_{inh}(\mathcal{F}) &= \{\{\alpha\}, \{\alpha,\beta\}, \{\alpha,\beta,\gamma\}\},\\
stg_{inh}(\mathcal{F}) &= sem_{inh}(\mathcal{F}) = stb_{inh} = \{\{\alpha,\beta,\gamma\}\},\\
grd(\mathcal{F}) &= \{\emptyset\}.
\end{aligned}$$

Hybrid semantics (Dvořák, Rapberger, and Woltran 2023) employ subset-maximization (such as in preferred semantics) on the claim-level rather than the argument level.

Definition 2.22 (Claim-defeat & claim-range)**.** *A set of arguments $S \subseteq A$* defeats *a claim $\alpha \in cl(A)$ in $\mathcal{F}$ iff S attacks every $a \in A$ with $cl(a) = \alpha$ (in $\mathcal{F}$). $S^*_{\mathcal{F}} = \{\alpha \in cl(A) \mid S$ defeats α in $\mathcal{F}\}$ denotes the set of all claims which are defeated by S in $\mathcal{F}$. The claim-range of a set S of arguments is denoted by $S^{\circledast}_{\mathcal{F}} = cl(S) \cup S^*_{\mathcal{F}}$.*

Definition 2.23 (Hybrid semantics)**.** *Let $\mathcal{F} = (A, R, cl)$ be a CAF with underlying AF $F = (A, R)$. Consider a set of claims $C \subseteq cl(A)$. We call $S \subseteq A$ a σ_{inh}-realization of C in $\mathcal{F}$ iff $S \in \sigma(A, R)$ and $cl(S) = C$.*

- *$C \in prf_{hyb}(\mathcal{F})$ if C is $\subseteq$-maximal in $adm_{inh}(\mathcal{F})$;*
- *$C \in naive_{hyb}(\mathcal{F})$ if C is $\subseteq$-maximal in $cf_{inh}(\mathcal{F})$;*
- *$C \in stb\text{-}adm_{hyb}(\mathcal{F})$ if there is an adm_{inh}-realization S of C which defeats any $\alpha \in cl(A) \setminus C$ (i.e., $S^{\circledast}_{\mathcal{F}} = cl(A)$);*
- *$C \in stb\text{-}cf_{hyb}(\mathcal{F})$ if there is a cf_{inh}-realization S of C which defeats any $\alpha \in cl(A) \setminus C$ (i.e., $S^{\circledast}_{\mathcal{F}} = cl(A)$);*

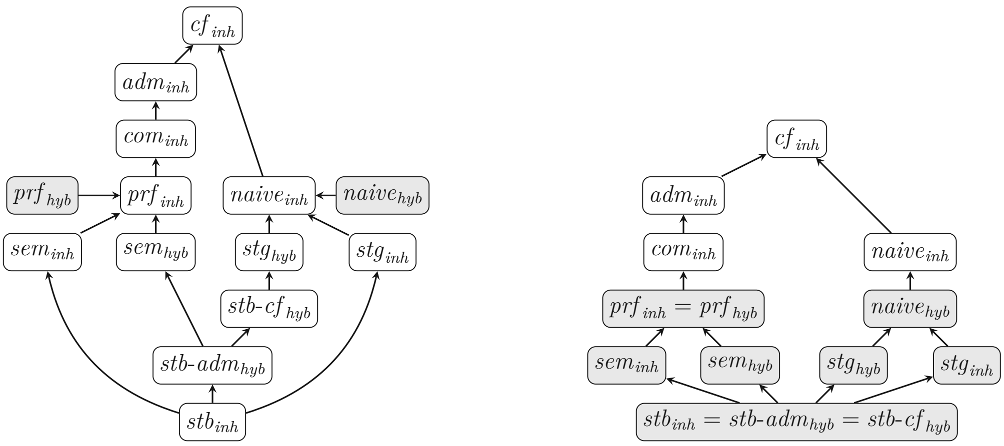

(a) General CAFs. (b) Well-formed CAFs.

Figure 2.5: Relations between semantics on (well-formed) CAFs. If there is an arrow from σ_μ to τ_ν, then $\sigma_\mu(\mathcal{F}) \subseteq \tau_\nu(\mathcal{F})$ for all CAFs $\mathcal{F}$ of the respective CAF-class. Semantics highlighted in gray are I-maximal.

- $C \in sem_{hyb}(\mathcal{F})$ *if there is an* adm_{inh}*-realization* S *of* C *in* $\mathcal{F}$ *such that there is no* $T \in adm(F)$ *with* $S^{\circledast}_{\mathcal{F}} \subset T^{\circledast}_{\mathcal{F}}$*;*
- $C \in stg_{hyb}(\mathcal{F})$ *if there is an* cf_{inh}*-realization* S *of* C *in* $\mathcal{F}$ *such that there is no* $T \in cf(F)$ *with* $S^{\circledast}_{\mathcal{F}} \subset T^{\circledast}_{\mathcal{F}}$*.*

To refer to an arbitrary CAF-semantics we write σ_μ or τ_ν, where $\sigma, \tau \in \{cf, adm, com, naive, stb, stb\text{-}adm, stb\text{-}cf, prf, sem, stg\}$ and $\mu, \nu \in \{inh, hyb\}$.

Example 2.24. *Consider again the CAF* $\mathcal{F} = (A, R, cl)$ *depicted in Figure 2.4. Recall that we already investigated this CAF with regards to inherited semantics inExample 2.21. In contrast to inherited semantics, for hybrid naive and preferred semantics we have*

$$naive_{hyb}(\mathcal{F}) = prf_{hyb}(\mathcal{F}) = \{\{\alpha, \beta, \gamma\}\}.$$

Regarding claim-range, notice that the admissible argument-set $\{a, c, e\}$ *already contains every claim in* $\mathcal{F}$*, i.e.,* $cl(\{a, c, e\}) = cl(A)$*. Thus,* $\{a, c, e\}^{\circledast}_{\mathcal{F}} = cl(A)$*. For the admissible argument set* $\{b, d\}$ *we have* $cl(\{b, d\}) = \{\alpha, \beta\}$ *and* $\{b, d\}^{*}_{\mathcal{F}} = \{\gamma\}$*, i.e.,* $\{b, d\}^{\circledast}_{\mathcal{F}} = cl(A)$*. There is no other admissible argument set* $S \in adm(A, R)$ *such that* $S^{\circledast}_{\mathcal{F}} = cl(A)$*. Thus,*

$$sem_{hyb}(\mathcal{F}) = stb\text{-}adm_{hyb}(\mathcal{F}) = \{\{\alpha, \beta\}, \{\alpha, \beta, \gamma\}\}.$$

For the conflict-free (but not admissible) argument set $\{b, e\}$ *we have* $cl(\{b, e\}) = \{\beta\}$ *and* $\{b, e\}^{*}_{\mathcal{F}} = \{\alpha, \gamma\}$*, i.e.,* $\{b, e\}^{\circledast}_{\mathcal{F}} = cl(A)$*. There is no other conflict-free argument set* $S \in cf((A, R))$ *such that* $S^{\circledast}_{\mathcal{F}} = cl(A)$*. Thus,*

$$stg_{hyb}(\mathcal{F}) = stb\text{-}cf_{hyb}(\mathcal{F}) = \{\{\beta\}, \{\alpha, \beta\}, \{\alpha, \beta, \gamma\}\}.$$

Table 2.3: Complexity of CAFs (Dvořák et al. 2023; Dvořák and Woltran 2020).

σ_μ	$Cred^{\Delta}_{\sigma_\mu}$	$Skept^{\Delta}_{\sigma_\mu}$		$Ver^{\Delta}_{\sigma_\mu}$	
	$\Delta \in \{CAF, wfCAF\}$	$\Delta = CAF$	$\Delta = wfCAF$	$\Delta = CAF$	$\Delta = wfCAF$
cf_{inh}	in P	trivial		NP-c	in P
adm_{inh}	NP-c	trivial		NP-c	in P
com_{inh}	NP-c	P-c		NP-c	in P
stb_{inh} $stb\text{-}adm_{hyb}$ $stb\text{-}cf_{hyb}$	NP-c	coNP-c		NP-c	in P
$naive_{inh}$	in P	coNP-c		NP-c	in P
$naive_{hyb}$		Π_2^{P}-c	coNP-c	DP-c	
prf_{inh}	NP-c	Π_2^{P}-c		Σ_2^{P}-c	coNP-c
prf_{hyb}				DP-c	
sem_{inh} sem_{hyb}	Σ_2^{P}-c	Π_2^{P}-c		Σ_2^{P}-c	coNP-c
stg_{inh} stg_{hyb}	Σ_2^{P}-c	Π_2^{P}-c		Σ_2^{P}-c	coNP-c

The relationship between the various CAF-semantics has been investigated for both general and well-formed CAFs (Dvořák, Rapberger, and Woltran 2023). See Figure 2.5 for a summary of these results. It can be seen that many inherited and hybrid semantics coincide on wfCAFs, but not on general CAFs.

Many argumentation semantics employ argument maximization (e.g. preferred or naive) and therefore deliver incomparable extensions on standard AFs: for all $S, T \in prf(F)$, $S \subseteq T$ implies $S = T$. This property is called I-maximality (Baroni and Giacomin 2007), and is defined analogously for CAFs:

Definition 2.25 (I-maximality)**.** *A CAF-semantics σ_μ is I-maximal for a class $\mathfrak{F}$ of CAFs if, for all CAFs $\mathcal{F} \in \mathfrak{F}$ and all $C, D \in \sigma_\mu(\mathcal{F})$, $C \subseteq D$ implies $C = D$.*

Figure 2.5 shows I-maximality properties of CAFs (Dvořák, Rapberger, and Woltran 2023). For wfCAFs, I-maximality is preserved in all maximization-based semantics except $naive_{inh}$, implying natural behavior analogous to AFs; see, e.g., (van der Torre and Vesic 2018) for a general discussion of such properties.

The computational complexity of CAFs has been investigated as well (Dvořák et al. 2023; Dvořák and Woltran 2020), revealing more differences between general CAFs and wfCAFs. The main decision problems for CAFs are defined analogously to those for AFs (see Definition 2.8), except that we are now interested in the acceptance of claims.

Definition 2.26 (Decision problems for CAFs)**.** *We consider the following decision problems pertaining to a CAF-semantics σ_μ:*

- Credulous Acceptance *($Cred^{CAF}_{\sigma_\mu}$): Given a CAF $\mathcal{F}$ and claim α, is α contained in some $C \in \sigma_\mu(\mathcal{F})$?*
- Skeptical Acceptance *($Skept^{CAF}_{\sigma_\mu}$): Given a CAF $\mathcal{F}$ and claim α, is α contained in each $C \in \sigma_\mu(\mathcal{F})$?*
- Verification *($Ver^{CAF}_{\sigma_\mu}$): Given a CAF $\mathcal{F}$ and a set of claims C, is $C \in \sigma_\mu(\mathcal{F})$?*

We furthermore consider these reasoning problems restricted to wfCAFs and denote them by $Cred^{wfCAF}_{\sigma_\mu}$, $Skept^{wfCAF}_{\sigma_\mu}$, and $Ver^{wfCAF}_{\sigma_\mu}$.

Table 2.3 shows the complexity of these problems. The complexity of the verification problem drops by one level in the polynomial hierarchy when comparing general CAFs to wfCAFs. This is an important advantage of wfCAFs, as a lower complexity in the verification problem allows for a more efficient enumeration of claim-extensions (Dvořák and Woltran 2020).

2.4 Choice Logics

Qualitative Choice Logic (QCL) (Brewka, Benferhat, and Berre 2004) is a formalism for preference representation that extends classical propositional logic by the connective $\vec{\times}$ called ordered disjunction. Intuitively, $A \vec{\times} B$ can be read as "A or B but preferably A". In this way, QCL enables us to express both hard- and soft-constraints, i.e., both truth and preferences, in one unified language. In the master thesis of the author (Bernreiter 2020), a general framework that captures QCL as well as other logics such as Conjunctive Choice Logic (CCL) (Boudjelida and Benferhat 2016) and Lexicographic Choice Logic (LCL) (Bernreiter 2020) was introduced.

In Subsection 2.4.1 we formally define the notion of choice logics in accordance with the choice logic framework introduced in (Bernreiter 2020). In Subsection 2.4.2 we explicitly define QCL, CCL, and LCL. Finally, in Subsection 2.4.3 we recall some crucial properties of choice logic formulas, including their computational complexity, as investigated in the master thesis of the author (Bernreiter 2020).

2.4.1 Syntax and Semantics

Choice logics are an extension of classical propositional logic (see Section 2.1). In addition to the classical connectives ($\neg, \wedge, \vee$) a choice logic can feature one or several additional choice connectives.

Definition 2.27. *Let $\mathcal{U}$ denote the (countably infinite) set of propositional variables (also called atoms). The set of choice connectives $\mathcal{C}_{\mathcal{L}}$ of a choice logic $\mathcal{L}$ is a finite set of symbols such that $\mathcal{C}_{\mathcal{L}} \cap \{\neg, \land, \lor\} = \emptyset$. The set $\mathcal{F}_{\mathcal{L}}$ of formulas of $\mathcal{L}$ is defined inductively as follows:*

1. *if $a \in \mathcal{U}$, then $a \in \mathcal{F}_{\mathcal{L}}$;*
2. *if $F \in \mathcal{F}_{\mathcal{L}}$, then $(\neg F) \in \mathcal{F}_{\mathcal{L}}$;*
3. *if $F, G \in \mathcal{F}_{\mathcal{L}}$, then $(F \circ G) \in \mathcal{F}_{\mathcal{L}}$ for $\circ \in (\{\land, \lor\} \cup \mathcal{C}_{\mathcal{L}})$.*

For instance, in QCL the set of choice connectives is $\mathcal{C}_{\text{QCL}} = \{\vec{\times}\}$. Formulas that do not contain a choice connective are referred to as classical formulas. Note that classical propositional logic is the choice logic containing no choice connectives, i.e., $\mathcal{C}_{\text{PL}} = \emptyset$.

The semantics of a choice logic is given by two functions: satisfaction degree and optionality. The satisfaction degree of a formula given an interpretation is either a natural number or ∞. The lower this degree, the more preferable the interpretation. The optionality of a formula describes the maximal finite satisfaction degree that this formula can be ascribed, and is used to penalize non-satisfaction.

Definition 2.28. *The optionality of a choice connective $\circ \in \mathcal{C}_{\mathcal{L}}$ in a choice logic $\mathcal{L}$ is given by a function $opt^{\circ}_{\mathcal{L}} \colon \mathbb{N}^2 \to \mathbb{N}$ such that $opt^{\circ}_{\mathcal{L}}(k, \ell) \leq (k+1) \cdot (\ell+1)$ for all $k, \ell \in \mathbb{N}$. The optionality of an $\mathcal{L}$-formula is given via $opt_{\mathcal{L}} \colon \mathcal{F}_{\mathcal{L}} \to \mathbb{N}$ with*

1. *$opt_{\mathcal{L}}(a) = 1$, for every $a \in \mathcal{U}$;*
2. *$opt_{\mathcal{L}}(\neg F) = 1$;*
3. *$opt_{\mathcal{L}}(F \land G) = max(opt_{\mathcal{L}}(F), opt_{\mathcal{L}}(G))$;*
4. *$opt_{\mathcal{L}}(F \lor G) = max(opt_{\mathcal{L}}(F), opt_{\mathcal{L}}(G))$;*
5. *$opt_{\mathcal{L}}(F \circ G) = opt^{\circ}_{\mathcal{L}}(opt_{\mathcal{L}}(F), opt_{\mathcal{L}}(G))$ for every $\circ \in \mathcal{C}_{\mathcal{L}}$.*

The optionality of a classical formula is always 1. Moreover, for any choice connective $\circ$, the optionality of $F \circ G$ is bounded such that $opt_{\mathcal{L}}(F \circ G) \leq (opt_{\mathcal{L}}(F)+1) \cdot (opt_{\mathcal{L}}(G)+1)$. The reason for this is that there are $opt_{\mathcal{L}}(F)$ many finite degrees that could be ascribed to F, plus the infinite degree ∞. Likewise for G. Thus, there are at most $(opt_{\mathcal{L}}(F) + 1) \cdot (opt_{\mathcal{L}}(G) + 1)$ possibilities in which the degrees of F and G can be combined.

As in classical propositional logic, interpretations in choice logics are sets of propositional variables. Again, a variable x is true under $\mathcal{I}$ iff $x \in \mathcal{I}$, and false under $\mathcal{I}$ iff $x \notin \mathcal{I}$. Regarding the domain of satisfaction degrees we write $\overline{\mathbb{N}}$ for $(\mathbb{N} \cup \{\infty\})$.

Definition 2.29. *The satisfaction degree of a choice connective* $\circ \in \mathcal{C}_\mathcal{L}$ *in a choice logic* $\mathcal{L}$ *is given by a function* $deg^\circ_\mathcal{L} \colon \mathbb{N}^2 \times \overline{\mathbb{N}}^2 \to \overline{\mathbb{N}}$ *such that* $deg^\circ_\mathcal{L}(k, \ell, m, n) \leq opt^\circ_\mathcal{L}(k, \ell)$ *or* $deg^\circ_\mathcal{L}(k, \ell, m, n) = \infty$ *for all* $k, \ell \in \mathbb{N}$ *and all* $m, n \in \overline{\mathbb{N}}$*. The satisfaction degree of an* $\mathcal{L}$*-formula under an interpretation* $\mathcal{I} \subseteq \mathcal{U}$ *is given via* $deg_\mathcal{L} \colon 2^\mathcal{U} \times \mathcal{F}_\mathcal{L} \to \overline{\mathbb{N}}$ *with*

1. $deg_\mathcal{L}(\mathcal{I}, a) = \begin{cases} 1 & \text{if } a \in \mathcal{I} \\ \infty & \text{otherwise} \end{cases}$ *for every* $a \in \mathcal{U}$*;*

2. $deg_\mathcal{L}(\mathcal{I}, \neg F) = \begin{cases} 1 & \text{if } deg_\mathcal{L}(\mathcal{I}, F) = \infty \\ \infty & \text{otherwise}; \end{cases}$

3. $deg_\mathcal{L}(\mathcal{I}, F \wedge G) = max(deg_\mathcal{L}(\mathcal{I}, F), deg_\mathcal{L}(\mathcal{I}, G))$*;*

4. $deg_\mathcal{L}(\mathcal{I}, F \vee G) = min(deg_\mathcal{L}(\mathcal{I}, F), deg_\mathcal{L}(\mathcal{I}, G))$*;*

5. $deg_\mathcal{L}(\mathcal{I}, F \circ G) = deg^\circ_\mathcal{L}(opt_\mathcal{L}(F), opt_\mathcal{L}(G), deg_\mathcal{L}(\mathcal{I}, F), deg_\mathcal{L}(\mathcal{I}, G))$ *for every* $\circ \in \mathcal{C}_\mathcal{L}$*.*

Note that, by definition, either $deg_\mathcal{L}(\mathcal{I}, F) \leq opt_\mathcal{L}(F)$ or $deg_\mathcal{L}(\mathcal{I}, F) = \infty$ for all $\mathcal{L}$-formulas F. This is as intended, since the optionality of a formula represents its maximal finite satisfaction degree.

We sometimes use the alternative notation $\mathcal{I} \models^\mathcal{L}_m F$ for $deg_\mathcal{L}(\mathcal{I}, F) = m$. If $m < \infty$, we say that $\mathcal{I}$ satisfies F (to a finite degree), and if $m = \infty$, then $\mathcal{I}$ does not satisfy F. If F is a classical formula, then $\mathcal{I} \models^\mathcal{L}_1 F$ iff $\mathcal{I} \models F$ and $\mathcal{I} \models^\mathcal{L}_\infty F$ iff $\mathcal{I} \not\models F$. The symbol $\bot$ is shorthand for the formula $(a \wedge \neg a)$, where a is an arbitrary variable. We have $opt_\mathcal{L}(\bot) = 1$ and $deg_\mathcal{L}(\mathcal{I}, \bot) = \infty$ for any interpretation $\mathcal{I}$ in every choice logic.

The models of a choice logic formula are the interpretations that satisfy the formula, and the preferred models are the models that satisfy the formula to a minimal degree.

Definition 2.30. *Let* $\mathcal{L}$ *be a choice logic,* $\mathcal{I}$ *an interpretation, and* F *an* $\mathcal{L}$*-formula.* $\mathcal{I}$ *is a model of* F*, written as* $\mathcal{I} \in Mod_\mathcal{L}(F)$*, iff* $deg_\mathcal{L}(\mathcal{I}, F) < \infty$*.* $\mathcal{I}$ *is a preferred model of* F*, written as* $\mathcal{I} \in Prf_\mathcal{L}(F)$*, iff* $\mathcal{I} \in Mod_\mathcal{L}(F)$ *and* $deg_\mathcal{L}(\mathcal{I}, F) \leq deg_\mathcal{L}(\mathcal{J}, F)$ *for all other interpretations* $\mathcal{J}$*.*

When specifying the (preferred) models of a formula F, we will often only include interpretations $\mathcal{I}$ that are relevant to F, i.e., $\mathcal{I} \subseteq var(F)$ where $var(F)$ is the set of variables occurring in F. This is purely for succinctness. In general, a formula that has a (preferred) model indeed has an infinite number of them, since $\mathcal{I} \in Mod_{\mathrm{QCL}}(F)$ (resp. $\mathcal{I} \in Prf_{\mathrm{QCL}}(F)$) and $x \notin var(F)$ implies $\mathcal{I} \cup \{x\} \in Mod_{\mathrm{QCL}}(F)$ (resp. $\mathcal{I} \cup \{x\} \in Prf_{\mathrm{QCL}}(F)$).

Table 2.4: The satisfaction degrees resulting from applying the classical connectives $\wedge$, $\vee$ and the choice connectives $\vec{\times}$ (QCL), $\vec{\odot}$ (CCL), $\vec{\diamond}$ (LCL) to atoms.

$\mathcal{I}$	$a \wedge b$	$a \vee b$	$a \vec{\times} b$	$a \vec{\odot} b$	$a \vec{\diamond} b$
$\emptyset$	∞	∞	∞	∞	∞
$\{b\}$	∞	1	2	∞	3
$\{a\}$	∞	1	1	2	2
$\{a, b\}$	1	1	1	1	1

2.4.2 Prominent Choice Logics

So far we introduced choice logics in a quite abstract way. We now introduce three particular instantiations, namely QCL (Brewka, Benferhat, and Berre 2004), the first and most prominent choice logic in the literature, CCL (Boudjelida and Benferhat 2016), which introduces a connective $\vec{\odot}$ called ordered conjunction in place of QCL's ordered disjunction $\vec{\times}$, and LCL (Bernreiter 2020), which replaces ordered disjunction with a lexicographic operator $\vec{\diamond}$.

Definition 2.31. QCL *is the choice logic such that* $\mathcal{C}_{\mathrm{QCL}} = \{\vec{\times}\}$, *and, if* $k = opt_{\mathrm{QCL}}(F)$, $\ell = opt_{\mathrm{QCL}}(G)$, $m = deg_{\mathrm{QCL}}(\mathcal{I}, F)$, *and* $n = deg_{\mathrm{QCL}}(\mathcal{I}, G)$, *then*

$$opt_{\mathrm{QCL}}(F \vec{\times} G) = opt^{\vec{\times}}_{\mathrm{QCL}}(k, \ell) = k + \ell, \text{ and}$$

$$deg_{\mathrm{QCL}}(\mathcal{I}, F \vec{\times} G) = deg^{\vec{\times}}_{\mathrm{QCL}}(k, \ell, m, n) = \begin{cases} m & \text{if } m < \infty; \\ n + k & \text{if } m = \infty, n < \infty; \\ \infty & \text{otherwise.} \end{cases}$$

In the above definition, we can see how optionality is used to penalize non-satisfaction: given a QCL-formula $F \vec{\times} G$ and an interpretation $\mathcal{I}$, if $\mathcal{I}$ satisfies F (to some finite degree), then $deg_{\mathrm{QCL}}(\mathcal{I}, F \vec{\times} G) = deg_{\mathrm{QCL}}(\mathcal{I}, F) \leq opt_{\mathrm{QCL}}(F)$; if $\mathcal{I}$ does not satisfy F, then $deg_{\mathrm{QCL}}(\mathcal{I}, F \vec{\times} G) = opt_{\mathrm{QCL}}(F) + deg_{\mathrm{QCL}}(\mathcal{I}, G) > opt_{\mathrm{QCL}}(F)$. Therefore, interpretations that satisfy F result in a lower degree, i.e., are more preferable, compared to interpretations that do not satisfy F. Table 2.4 shows how ordered disjunction behaves when applied to atoms. The following example highlights how classical conjunction interacts with ordered disjunction.

Example 2.32. *Consider the* QCL*-formula*

$$F = (a \vec{\times} c) \wedge (b \vec{\times} c).$$

Notice that the interpretation $\{a\}$ *satisfies* $(a \vec{\times} c)$ *to a degree of 1 but* $(b \vec{\times} c)$ *to a degree of* ∞. *Thus,* $\{a\} \models^{\mathrm{QCL}}_{\infty} F$. *Analogously for* $\{b\}$. *Moreover,* $\emptyset$ *satisfies both* $(a \vec{\times} c)$ *and* $(b \vec{\times} c)$ *to a degree of* ∞, *i.e.,* $\emptyset \models^{\mathrm{QCL}}_{\infty} F$.

The interpretation $\{a, c\}$ *satisfies* $(a \vec{\times} c)$ *to a degree of 1 but* $(b \vec{\times} c)$ *to a degree of* 2. *Thus,* $\{a, c\} \models_2^{\text{QCL}} F$. *Analogously for* $\{b, c\}$. *Moreover,* $\{c\}$ *satisfies both* $(a \vec{\times} c)$ *and* $(b \vec{\times} c)$ *to a degree of* 2, *i.e.,* $\{c\} \models_2^{\text{QCL}} F$.

Lastly, the interpretations $\{a, b\}$ *and* $\{a, b, c\}$ *satisfy both* $(a \vec{\times} c)$ *and* $(b \vec{\times} c)$ *to a degree of* 1, *i.e.,* $\mathcal{I} \models_1^{\text{QCL}} F$ *for* $\mathcal{I} \in \{\{a, b\}, \{a, b, c\}\}$. *We can conclude that* $\mathit{Mod}_{\text{QCL}}(F) = \{\{c\}, \{a, c\}, \{b, c\}, \{a, b\}, \{a, b, c\}\}$ *while* $\mathit{Prf}_{\text{QCL}}(F) = \{\{a, b\}, \{a, b, c\}\}$.

Suppose now we get the additional information that a *and* b *cannot be jointly satisfied. We encode this in the updated formula*

$$F' = ((a \vec{\times} c) \wedge (b \vec{\times} c)) \wedge \neg(a \wedge b).$$

Now the interpretations $\{a, b\}$ *and* $\{a, b, c\}$ *satisfy* F' *to a degree of* ∞, *while* $\{c\}$, $\{a, c\}$, *and* $\{b, c\}$ *still satisfy* F' *to a degree of 2. Thus, we have that* $\mathit{Prf}_{\text{QCL}}(F') = \{\{c\}, \{a, c\}, \{b, c\}\}$. *This shows that the notion of preferred models in QCL is non-monotonic, since the addition of new information has lead to entirely new preferred models.*

Next, we define CCL. Note that we follow the revised definition of CCL (Bernreiter 2020), which differs from the initial specification[3]. Intuitively, given a CCL-formula $F \vec{\odot} G$ it is best to satisfy both F and G, but also acceptable to satisfy only F. For instance, when buying a new car, one might insist that the car has cruise control (*cruise*), while preferring configurations that additionally feature a lane assistant (*lane*). This can be formalized in CCL as the formula $\mathit{cruise} \vec{\odot} \mathit{lane}$.

Definition 2.33. CCL *is the choice logic such that* $\mathcal{C}_{\text{CCL}} = \{\vec{\odot}\}$, *and, if* $k = \mathit{opt}_{\text{CCL}}(F)$, $\ell = \mathit{opt}_{\text{CCL}}(G)$, $m = \mathit{deg}_{\text{CCL}}(\mathcal{I}, F)$, *and* $n = \mathit{deg}_{\text{CCL}}(\mathcal{I}, G)$, *then*

$$\mathit{opt}_{\text{CCL}}(F \vec{\odot} G) = k + \ell, \text{ and}$$

$$\mathit{deg}_{\text{CCL}}(\mathcal{I}, F \vec{\odot} G) = \begin{cases} n & \text{if } m = 1, n < \infty; \\ m + \ell & \text{if } m < \infty \text{ and } (m > 1 \text{ or } n = \infty); \\ \infty & \text{otherwise.} \end{cases}$$

Example 2.34. *Consider the* CCL*-formula*

$$G = (a \vec{\odot} c) \wedge (b \vec{\odot} c).$$

There are only two interpretations relevant to G *that satisfy* G, *namely* $\mathit{Mod}_{\text{CCL}}(G) = \{\{a, b\}, \{a, b, c\}\}$. *Of these models,* $\{a, b, c\}$ *satisfies* G *to a degree of* 1 *while* $\{a, b\}$ *satisfies* G *to a degree of* 2. *Therefore,* $\mathit{Prf}_{\text{CCL}}(G) = \{\{a, b, c\}\}$.

[3] Under the initial definition of CCL, $a \vec{\odot} b$ is always ascribed a degree of 1 or ∞, i.e., non-classical degrees cannot be obtained (cf. Definition 8 in (Boudjelida and Benferhat 2016)).

The last choice logic we consider, LCL, employs a more fine-grained type of preference: given an LCL-formula $F\vec{\diamond}G$, it is best to satisfy F and G, second-best to satisfy only F, and third-best to satisfy only G.

Definition 2.35. LCL *is the choice logic such that* $\mathcal{C}_{\mathrm{LCL}} = \{\vec{\diamond}\}$, *and, if* $k = opt_{\mathrm{LCL}}(F)$, $\ell = opt_{\mathrm{LCL}}(G)$, $m = deg_{\mathrm{LCL}}(\mathcal{I}, F)$, *and* $n = deg_{\mathrm{LCL}}(\mathcal{I}, G)$, *then*

$$opt_{\mathrm{LCL}}(F\vec{\diamond}G) = (k+1)\cdot(\ell+1) - 1, \text{ and}$$

$$deg_{\mathrm{LCL}}(\mathcal{I}, F\vec{\diamond}G) = \begin{cases} (m-1)\cdot\ell + n & \text{if } m < \infty, n < \infty; \\ k\cdot\ell + m & \text{if } m < \infty, n = \infty; \\ k\cdot\ell + k + n & \text{if } m = \infty, n < \infty; \\ \infty & \text{otherwise.} \end{cases}$$

Example 2.36. *Consider the* LCL-*formula*

$$H = (a\vec{\diamond}c) \wedge (b\vec{\diamond}c).$$

The models relevant to H *are* $Mod_{\mathrm{LCL}}(H) = \{\{c\}, \{a,b\}, \{a,c\}, \{b,c\}, \{a,b,c\}\}$. *Of these models,* $\{a,b,c\}$ *satisfies* H *to a degree of* 1, $\{a,b\}$ *satisfies* H *to a degree of* 2, *and* $\{c\}$, $\{a,c\}$, $\{b,c\}$ *satisfy* H *to a degree of* 3. *Therefore,* $Prf_{\mathrm{LCL}}(H) = \{\{a,b,c\}\}$.

2.4.3 Properties

In this section we present some properties of choice logics that were established in Bernreiter (2020) and Bernreiter, Maly, and Woltran (2021), and that we will need in this thesis.

First, we define the notion of classical counterparts for choice connectives and choice logic formulas.

Definition 2.37. *Let* $\mathcal{L}$ *be a choice logic. The classical counterpart of a choice connective* $\circ \in \mathcal{C}_{\mathcal{L}}$ *is the classical binary connective* $\circledast$ *such that, for all atoms* a *and* b, *we have* $deg_{\mathcal{L}}(\mathcal{I}, a \circ b) < \infty \iff \mathcal{I} \models a \circledast b$. *The classical counterpart of an* $\mathcal{L}$-*formula* F *is denoted as* $cp(F)$ *and is obtained by replacing all occurrences of choice connectives in* F *by their classical counterparts.*

Every choice connective has exactly one classical binary connective as its classical counterpart (Bernreiter 2020, Proposition 22). For example, the classical counterpart of ordered disjunction $\vec{\times}$ is regular disjunction $\vee$, and the classical counterpart of the QCL-formula $F = (a\vec{\times}b) \vee c$ is $cp(F) = (a \vee b) \vee c$. In CCL, the classical counterpart of $\vec{\odot}$ is the first projection, i.e., $cp(A\vec{\odot}B) = A$. In LCL, the classical counterpart of $\vec{\diamond}$ is $\vee$. A natural property of choice logics considered in this thesis is that choice connectives can be replaced by their classical counterpart without affecting satisfiability.

Proposition 2.38. *(Bernreiter 2020, Proposition 23) Let* $\mathcal{L} \in \{\mathrm{QCL}, \mathrm{CCL}, \mathrm{LCL}\}$. *It holds that* $deg_{\mathcal{L}}(\mathcal{I}, F) < \infty$ *iff* $\mathcal{I} \models cp(F)$ *for all interpretations* $\mathcal{I}$ *and all* $\mathcal{L}$-*formulas* F.

Another result we will make use of is that any choice logic can express arbitrary assignments of satisfaction degrees to interpretations, as long as the degrees are obtainable in the following sense:

Definition 2.39. *A degree $m \in \overline{\mathbb{N}}$ is called obtainable in a choice logic $\mathcal{L}$ iff there exists an interpretation $\mathcal{I}$ and an $\mathcal{L}$-formula F such that $deg_{\mathcal{L}}(\mathcal{I}, F) = m$. By $\mathcal{D}_{\mathcal{L}}$ we denote the set of all degrees obtainable in a choice logic $\mathcal{L}$.*

For example, $\mathcal{D}_{\mathrm{PL}} = \{1, \infty\}$ and $\mathcal{D}_{\mathcal{L}} = \overline{\mathbb{N}}$ for $\mathcal{L} \in \{\mathrm{QCL}, \mathrm{CCL}, \mathrm{LCL}\}$. As soon as a degree m is obtainable, any interpretation can be assigned this degree via a suitable formula. This is useful when proving results for choice logics in general, instead of for a specific choice logic.

Proposition 2.40. *(Bernreiter, Maly, and Woltran 2021, Proposition 1) Let $\mathcal{L}$ be a choice logic. Let V be a finite set of propositional variables, and let s be a function $s\colon 2^V \to \mathcal{D}_{\mathcal{L}}$. Then there is an $\mathcal{L}$-formula F such that for every $\mathcal{I} \subseteq V$, $deg_{\mathcal{L}}(\mathcal{I}, F) = s(\mathcal{I})$.*

Lastly, we will make use of complexity results for choice logic formulas established in (Bernreiter 2020). To this end, we require the notion of tractable choice logics.

Definition 2.41. *A choice logic $\mathcal{L}$ is called tractable if the optionality- and degree functions of every choice connective in $\mathcal{L}$ are polynomial-time computable.*

QCL, CCL, and LCL are all tractable in the above sense.

Definition 2.42. *Given a choice logic $\mathcal{L}$ we define the following decision problems:*

- *$\mathcal{L}$-DegreeChecking: given an $\mathcal{L}$-formula F, an interpretation $\mathcal{I}$, and a satisfaction degree $k \in \overline{\mathbb{N}}$, does $deg_{\mathcal{L}}(\mathcal{I}, F) \leq k$ hold?*
- *$\mathcal{L}$-DegreeSat: given an $\mathcal{L}$-formula F and a satisfaction degree k, is there an interpretation $\mathcal{I}$ such that $deg_{\mathcal{L}}(\mathcal{I}, F) \leq k$ holds?*
- *$\mathcal{L}$-PMChecking: given an $\mathcal{L}$-formula F and an interpretation $\mathcal{I}$, does $\mathcal{I} \in Prf_{\mathcal{L}}(F)$ hold?*
- *$\mathcal{L}$-PMContainment: given an $\mathcal{L}$-formula F and a propositional variable x, is there an interpretation $\mathcal{I} \in Prf_{\mathcal{L}}(F)$ such that $x \in \mathcal{I}$?*

The complexity of the above decision problems, as investigated in (Bernreiter 2020), is shown in Table 2.5. Note that $\mathcal{L}$-PMChecking is in P for classical propositional logic, while it is coNP-complete for the non-classical QCL, CCL, and LCL. Moreover, the complexity of $\mathcal{L}$-PMContainment ranges from NP-complete (PL) to Δ_2^{P}-complete (LCL), with QCL and CCL located inbetween (Θ_2^{P}-complete).

Table 2.5: Complexity of choice logic formulas. PL stands for classical propositional logic and "Tract.‘" stands for an arbitrary tractable choice logic.

	Tract.	PL	QCL/CCL	LCL
$\mathcal{L}$-DegreeChecking	in P	in P	in P	in P
$\mathcal{L}$-DegreeSat	NP-c	NP-c	NP-c	NP-c
$\mathcal{L}$-PMChecking	in coNP	in P	coNP-c	coNP-c
$\mathcal{L}$-PMContainment	NP-h/in Δ_2^P	NP-c	Θ_2^P-c	Δ_2^P-c

To further study the complexity of choice logics, as we will do in Chapter 5, we need the notion of a formula's size. As in (Bernreiter 2020) and (Bernreiter, Maly, and Woltran 2021), $|F|$ denotes the total number of variables occurrences in F, e.g. $|(x \wedge x \wedge y)| = 3$. In general, the following holds.

Lemma 2.43. *(Bernreiter, Maly, and Woltran 2021, Lemma 11) Let $\mathcal{L}$ be a choice logic. Then, for every $\mathcal{L}$-formula F it holds that $opt_{\mathcal{L}}(F) < 2^{(|F|^2)}$.*

For QCL and CCL, but not LCL, we have $opt_{\mathcal{L}}(F) \leq |F|$.

CHAPTER 3

Conditional Preferences in Abstract Argumentation

Many situations require the use of conditional preferences, where a choice between two options (e.g. whether to drink tea or coffee) is dependent on other factors (e.g. the time of day). This has lead to the introduction of formalisms explicitly defined to deal with conditional preferences. A prominent example are CP-nets (Boutilier et al. 2004), which use graphs for preference representation. Another example is logic programming, where conditional preferences may occur in the head of rules (Brewka, Niemelä, and Syrjänen 2004; Charalambidis, Rondogiannis, and Troumpoukis 2021; Delgrande, Schaub, and Tompits 2003) or as dedicated preference statements (Brewka et al. 2015).

Despite this, conditional preferences have received only limited attention in the field of argumentation. Dung et al. investigated conditional preferences in the setting of structured argumentation (Dung, Thang, and Son 2019). There, argumentation frameworks (AFs) are built from defeasible knowledge bases containing preference rules of the form $a_1, \ldots, a_n \rightarrow d_0 \succ d_1$, where d_0 and d_1 are defeasible rules. Similarly, there is only one recent paper we are aware of that deals with conditional preferences on the abstract level (Alfano et al. 2023). This is in contrast to unconditional preferences, which are extensively studied both in structured (Modgil and Prakken 2010, 2013, 2018) and abstract (Alfano et al. 2022; Atkinson and Bench-Capon 2021; Bistarelli and Santini 2021; Kaci et al. 2021) argumentation in the literature.

To demonstrate the importance of conditional preferences in common reasoning tasks, we now adapt an example from (Dung, Thang, and Son 2019):

Example 3.1. *Sherlock Holmes is investigating a murder. There are two suspects, Person 1 and Person 2. After analyzing the crime scene, Sherlock is sure:*

- I_1: *Person 1 or Person 2 is the culprit, but not both.*

Moreover, Sherlock adheres to the following rules:

- R_1*: If Person* i *has a motive but Person* j*, with* $j \neq i$*, does not, then this supports the case that Person* i *is the culprit.*
- R_2*: If Person* i *has an alibi but Person* j*, with* $j \neq i$*, does not, then this supports the case that Person* j *is the culprit.*
- R_3*: Alibis have more importance than motives.*

After interrogating the suspects, Sherlock concludes that:

- C_1*: Person 1 has a motive but Person 2 does not.*
- C_2*: Person 1 has an alibi but Person 2 does not.*

If C_1 *is accepted, but* C_2 *is not, then this supports that Person 1 is the culprit. If* C_2 *is accepted then this supports that Person 2 is the culprit, regardless of our stance on* C_1*.*

In this chapter we aim to capture conditional preferences in argumentation on the abstract level rather than the structured level. Doing so will generalize existing formalisms for unconditional preferences in abstract argumentation and may provide a more direct target formalism for structured approaches.

Contributions. We introduce Conditional Preference-based AFs (CPAFs), where each subset of arguments S can be associated with its own preference relation $\succ_S$. Preferences are then resolved via one of four preference reductions (cf. Subsection 2.3.2) which modify the attack relation based on the given preferences. As a consequence, S must be justified in view of its own preferences, i.e., S must be an extension in view of $\succ_S$. We investigate the following topics relevant to CPAFs:

- We show that CPAFs generalize Preference-based AFs (PAFs), and demonstrate that they are capable of dealing with conditional preferences in a general manner.
- We conduct a principle-based analysis of CPAF-semantics and show that especially complete and stable semantics preserve properties that hold on PAFs. This analysis is helpful when aiming to understand the behavior of CPAF-semantics in a general manner, and lets us pinpoint differences to AFs/PAFs formally.
- We analyze the computational complexity of CPAFs in detail, showing that for some semantics (naive, complete, grounded, preferred) the introduction of conditional preferences can cause a rise in complexity compared to AFs. This gives insights into the expressiveness of CPAFs, and differentiates them further from AFs/PAFs.

- Lastly, we compare CPAFs to related formalisms. Specifically, we show that CPAFs can capture other generalizations of AFs such as Value-based AFs (VAFs) (Atkinson and Bench-Capon 2021; Bench-Capon, Doutre, and Dunne 2007) in a straightforward way, and compare CPAFs to Extended AFs (EAFs) (Baroni et al. 2009; Dunne, Modgil, and Bench-Capon 2010; Modgil 2009) in order to highlight similarities and differences. Moreover, we discuss a recently introduced alternative approach to conditional preferences in abstract argumentation (Alfano et al. 2023) and compare it to our CPAFs.

Publications. This chapter is based on the papers (Bernreiter, Dvořák, and Woltran 2022) and (Bernreiter, Dvořák, and Woltran 2023).

Outline. The remainder of this chapter is structured as follows: In Section 3.1 we introduce CPAFs and investigate them with respect to some basic properties. Section 3.2 contains our principle-based analysis, and in Section 3.3 we analyze the computational complexity of CPAFs. We discuss related formalisms in Section 3.4 and conclude in Section 3.5.

Required preliminaries. Before reading this chapter, it is recommended to read Section 2.1 (propositional logic), Section 2.2 (computational complexity), and especially Subsection 2.3.1 (abstract argumentation) and Subsection 2.3.2 (preferences in abstract argumentation).

3.1 Conditional Preference-Based Argumentation Frameworks (CPAFs)

As argued above, our aim is to provide a framework for reasoning with conditional preferences in abstract argumentation. This means that arguments themselves must be capable of expressing preferences, and that those argument-bound preferences are relevant only if the corresponding arguments are themselves accepted. How this is implemented must be considered carefully, as Example 3.1 demonstrates. There, the fact that Person 1 has a motive (let us refer to this as m_1) and the fact that Person 1 has an alibi (a_1) result in opposing preferences. When accepting both m_1 and a_1 it seems natural to combine these opposing preferences, i.e., to cancel them. But this does not allow us to express that alibis are more important than motives, as required in Example 3.1. Therefore, we need to define our formalism in a general way such that the joint acceptance of arguments must not necessarily result in the combination of their associated preferences. We solve this by mapping each subset S of arguments to a separate preference relation $\succ_S$.

Definition 3.2. *A Conditional PAF (CPAF) is a triple* $\mathcal{P} = (A, R, cond)$*, where* (A, R) *is an AF and* $cond\colon 2^A \to 2^{(A\times A)}$ *is a function that maps each set of arguments* $S \subseteq A$ *to an irreflexive and asymmetric binary relation* $\succ_S$ *over* A.

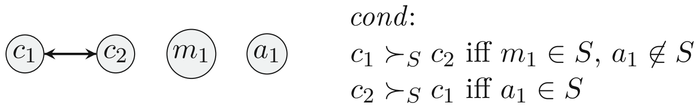

Figure 3.1: The CPAF $\mathcal{P}$ from Example 3.4.

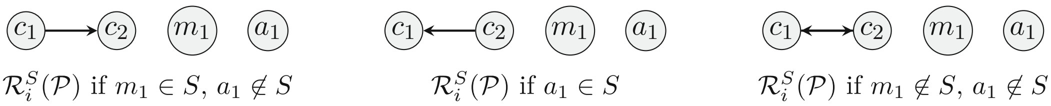

Figure 3.2: The preference-reducts of the CPAF $\mathcal{P}$ from Figure 3.1/Example 3.4.

We set no restriction on how exactly conditional preferences are represented. This is deliberate, as we wish to stay as general as possible. In practice, succinct representations could be achieved, e.g., by expressing the *cond*-function via rules of the form $\varphi \Rightarrow \ x \succ y$ where φ is a propositional formula over the arguments. Indeed, this representation will be used by us in Section 3.3 where we analyze the complexity of CPAFs.

Just as in PAFs, preferences in CPAFs are resolved with the help of the four preference-reductions (cf. Definition 2.10). A set of arguments S is an extension of some CPAF if it is an extension relative to its associated preference relation $cond(S)$.

Definition 3.3. *Let $\mathcal{P} = (A, R, cond)$ be a CPAF and let $S \subseteq A$. The S-reduct of $\mathcal{P}$ with respect to a preference reduction $i \in \{1,2,3,4\}$ is defined as $\mathcal{R}_i^S(\mathcal{P}) = \mathcal{R}_i(A, R, cond(S))$. Given an AF-semantics σ we define the CPAF-semantics σ_{cp}^i as follows: $S \in \sigma_{cp}^i(\mathcal{P})$ iff $S \in \sigma(\mathcal{R}_i^S(\mathcal{P}))$.*

Using CPAFs we can easily formalize our Sherlock Holmes example.

Example 3.4. *We continue Example 3.1 and introduce two arguments c_1 and c_2 expressing that Person 1 (resp. Person 2) is the culprit. Moreover, we introduce m_1 and a_1 to express that Person 1 has a motive (resp. alibi) but Person 2 does not. c_1 and c_2 attack each other while m_1 and a_1 have no incoming or outgoing attacks, but rather express preferences. Formally, we model this via the CPAF $\mathcal{P} = (\{c_1, c_2, m_1, a_1\}, \{(c_1, c_2), (c_2, c_1)\}, cond)$ with cond such that $c_1 \succ_S c_2$ iff $m_1 \in S$ but $a_1 \notin S$, $c_2 \succ_S c_1$ iff $a_1 \in S$, and $cond(S) = \emptyset$ for all other $S \subseteq A$. Figure 3.1 depicts $\mathcal{P}$ and Figure 3.2 shows the S-reducts of $\mathcal{P}$. Note that m_1 and a_1 are unattacked in all S-reducts of $\mathcal{P}$. Therefore, both arguments must be part of any σ_{cp}^i-extension for $\sigma \in \{grd, com, prf, stb\}$ and we can conclude that $\sigma_{cp}^i(\mathcal{P}) = \{\{m_1, a_1, c_2\}\}$.*

Note that, according to Definition 3.3, preferred semantics do not maximize over all admissible sets of a CPAF, but rather over all admissible sets in the given S-reduct. This means that if there is a set S that is admissible in the S-reduct of $\mathcal{P}$, but there is also some $T \supset S$ that is admissible in the S-reduct of $\mathcal{P}$, then S is not preferred in the

S-reduct of $\mathcal{P}$ (and therefore $S \notin \mathit{prf}^i_{cp}(\mathcal{P})$). But this T does not have to be admissible in $\mathcal{P}$, since it might not be admissible in the T-reduct of $\mathcal{P}$. The situation is analogous for naive semantics. The following alternative semantics may be considered more natural:

Definition 3.5. *Let $\mathcal{P} = (A, R, \mathit{cond})$ be a CPAF and let $S \subseteq A$. Then*

- *$S \in \mathit{naive\text{-}glb}^i_{cp}(\mathcal{P})$ iff $S \in \mathit{cf}^i_{cp}(\mathcal{P})$ and there is no T such that $S \subset T$ and $T \in \mathit{cf}^i_{cp}(\mathcal{P})$;*
- *$S \in \mathit{prf\text{-}glb}^i_{cp}(\mathcal{P})$ iff $S \in \mathit{adm}^i_{cp}(\mathcal{P})$ and there is no T such that $S \subset T$ and $T \in \mathit{adm}^i_{cp}(\mathcal{P})$.*

Intuitively, $\mathit{naive\text{-}glb}^i_{cp}$ and $\mathit{prf\text{-}glb}^i_{cp}$ maximize *globally* over all admissible sets of a CPAF, while naive^i_{cp} and prf^i_{cp} maximize *locally* over the admissible sets of the given S-reduct.

Example 3.6. *Let $\mathcal{P}$ be the CPAF from Example 3.4 and recall that $\mathit{prf}^i_{cp}(\mathcal{P}) = \{\{m_1, a_1, c_2\}\}$. Observe that $\{m_1, c_1\}$ is not preferred in the $\{m_1, c_1\}$-reduct of $\mathcal{P}$, but it is a subset-maximal admissible set in $\mathcal{P}$. Thus, $\mathit{prf\text{-}glb}^i_{cp}(\mathcal{P}) = \{\{m_1, a_1, c_2\}, \{m_1, c_1\}\}$.*

The difference between local and global maximization is not only philosophical, but impacts fundamental properties for maximization-based semantics such as I-maximality (Baroni and Giacomin 2007). A semantics σ^i_{cp} is I-maximal if and only if $S \subseteq T$ implies $S = T$ for all CPAFs $\mathcal{P}$ and all $S, T \in \sigma^i_{cp}(\mathcal{P})$.

Proposition 3.7. *$\mathit{prf\text{-}glb}^i_{cp}$ is I-maximal, but prf^i_{cp} is not, where $i \in \{1, 2, 3, 4\}$.*

Proof. I-maximality of $\mathit{prf\text{-}glb}^i_{cp}$ follows from Definition 3.5. Regarding counterexamples for prf^i_{cp} we consider the preference-reductions separately. Reduction 1: consider the CPAF depicted in Figure 3.3a, i.e., $\mathcal{P} = (\{a, b\}, \{(a, b)\}, \mathit{cond}\}$ with cond such that $b \succ_{\{a,b\}} a$. Then $\{a\} \in \mathit{prf}^1_{cp}(\mathcal{P})$ and $\{a, b\} \in \mathit{prf}^1_{cp}(\mathcal{P})$. Reductions 2 and 4: consider the CPAF depicted in Figure 3.3b, i.e., $\mathcal{P}' = (\{a, b, c\}, \{(a, b), (b, c), (c, a)\}, \mathit{cond}\}$ with cond such that $a \succ_{\{a\}} c$. Then $\emptyset \in \mathit{prf}^i_{cp}(\mathcal{P}')$ and $\{a\} \in \mathit{prf}^i_{cp}(\mathcal{P}')$. Reduction 3: consider the CPAF depicted in Figure 3.3c, i.e., $\mathcal{P}'' = (\{a, b, c\}, \{(a, b), (b, a), (b, c), (c, a)\}, \mathit{cond}\}$ with cond such that $a \succ_{\emptyset} b$. Then $\emptyset \in \mathit{prf}^3_{cp}(\mathcal{P}'')$ and $\{b\} \in \mathit{prf}^3_{cp}(\mathcal{P}'')$. □

One may be tempted to deduce from the above proposition that $\mathit{prf\text{-}glb}^i_{cp}$ is more suitable as a default preferred semantics than prf^i_{cp}. However, we will see in Section 3.4.1 that prf^i_{cp} allows us to capture the problems of subjective/objective acceptance in VAFs in a natural way. In our subsequent analysis of CPAFs we consider both local and global subset maximization. Like preferred semantics, naive and stable semantics satisfy I-maximality on AFs. Interestingly, on CPAFs, this depends on the preference-reduction.

Proposition 3.8. *$\mathit{naive\text{-}glb}^i_{cp}$ is I-maximal for $i \in \{1, 2, 3, 4\}$. Moreover, naive^j_{cp} is I-maximal for $j \in \{2, 3, 4\}$ but not for $j = 1$.*

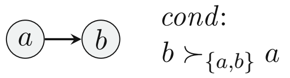

(a) Counterexample for prf^1_{cp}, $naive^1_{cp}$, and stb^1_{cp}.

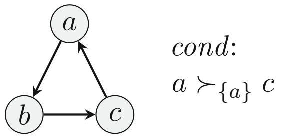

(b) Counterexample for prf^2_{cp} and prf^4_{cp}.

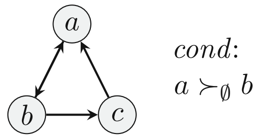

(c) Counterexample for prf^3_{cp}.

Figure 3.3: Counterexamples for I-maximality (cf. Propositions 3.7,3.8,3.9).

Proof. I-maximality of $\textit{naive-glb}^i_{cp}$ follows from Definition 3.5. I-maximality of $naive^j_{cp}$ with $j \in \{2,3,4\}$ follows from the fact that Reductions 2, 3, and 4 do not remove conflicts between arguments, and therefore conflict-free sets are the same across all S-reducts. For $naive^1_{cp}$ we can use the same counter-example as for prf^1_{cp} (cf. Proposition 3.7 and Figure 3.3a). □

Proposition 3.9. *stb^j_{cp} is I-maximal for $j \in \{2,3,4\}$ but not for $j = 1$.*

Proof. For stb^1_{cp} we can use the same counter-example as for prf^1_{cp} (cf. Proposition 3.7 and Figure 3.3a). For stb^j_{cp} with $j \in \{2,3,4\}$ we proceed by contradiction: assume there is a CPAF $\mathcal{P} = (A, R, cond)$ with $S, T \in stb^j_{cp}(\mathcal{P})$ such that $S \subset T$. Then there is $x \in T$ such that $x \notin S$. Since $S \in stb^j_{cp}(\mathcal{P})$ there is $y \in S$ such that $(y, x) \in \mathcal{R}^S_j(\mathcal{P})$. Reductions 2, 3, and 4 do not remove conflicts between arguments, and thus either $(y, x) \in R$ or $(x, y) \in R$. Therefore, $(y, x) \in \mathcal{R}^T_j(\mathcal{P})$ or $(x, y) \in \mathcal{R}^T_j(\mathcal{P})$. But $y \in S$ implies $y \in T$, i.e., $T \notin cf^j_{cp}(\mathcal{P})$. □

A further well-known property of AFs is that if an argument set S is stable in a framework F, then S is also preferred in F (Dung 1995). The same is true for CPAFs, with the exception of preferred semantics utilizing global maximization and Reduction 1.

Proposition 3.10. *If $S \in stb^i_{cp}(\mathcal{P})$ then $S \in prf^i_{cp}(\mathcal{P})$ for $i \in \{1,2,3,4\}$. Moreover, if $S \in stb^j_{cp}(\mathcal{P})$ then $S \in \textit{prf-glb}^j_{cp}(\mathcal{P})$ for $j \in \{2,3,4\}$. However, $S \in stb^1_{cp}(\mathcal{P})$ does not necessarily imply $S \in \textit{prf-glb}^1_{cp}(\mathcal{P})$.*

Proof. Let $\mathcal{P} = (A, R, cond)$ be a CPAF, and let $S \in stb^i_{cp}(\mathcal{P})$, where $i \in \{1,2,3,4\}$. Then $S \in stb(\mathcal{R}^S_i(\mathcal{P}))$. Since $\mathcal{R}^S_i(\mathcal{P})$ is an AF this implies $S \in prf(\mathcal{R}^S_i(\mathcal{P}))$ which means that $S \in prf^i_{cp}(\mathcal{P})$.

Now let $j \in \{2,3,4\}$. If $S \in stb^j_{cp}(\mathcal{P})$ then every argument in $\mathcal{R}^S_j(\mathcal{P})$ is either in S or attacked by it. Towards a contradiction, assume there is $T \in adm^j_{cp}(\mathcal{P})$ such that $T \supset S$.

Then there is some $x \in T \setminus S$. Since S attacks x in $\mathcal{R}_j^S(\mathcal{P})$ there is a conflict between some $y \in S$ and x in the underlying AF (A, R) of $\mathcal{P}$. Note that $y \in T$. But Reductions $2, 3, 4$ cannot remove conflicts between arguments, i.e., $T \notin cf(\mathcal{R}_j^T(\mathcal{P}))$. Contradiction.

For $prf\text{-}glb_{cp}^1$, let $\mathcal{P}$ be the CPAF from Figure 3.3a). Then $\{a\} \in stb_{cp}^1(\mathcal{P})$ but $\{a\} \notin prf\text{-}glb_{cp}^1(\mathcal{P})$. □

A further interesting point is that grounded extensions are not necessarily unique in CPAFs: consider $\mathcal{P} = (\{a, b\}, \{(a, b)\}, cond)$ with $cond$ such that $b \succ_{\{b\}} a$. Then $\{a\} \in grd_{cp}^2(\mathcal{P})$ and $\{b\} \in grd_{cp}^2(\mathcal{P})$. We stress that each grounded extension S is still unique in the S-reduct of the given CPAF and thus unique with respect to its own preferences.

One more crucial difference between PAFs and CPAFs we want to highlight concerns Dung's fundamental lemma (Dung 1995), which says that if a set of arguments S is admissible and x is acceptable w.r.t. S then $S \cup \{x\}$ is admissible.[4] This fundamental lemma is satisfied in AFs (and therefore PAFs) but not in CPAFs.

Proposition 3.11. *Dung's fundamental lemma does not hold for CPAFs.*

Proof. Let $\mathcal{P}$ be the CPAF from Example 3.4/Figure 3.1 with its reducts shown in Figure 3.2. Note that $\{c_2\} \in adm_{cp}^i(\mathcal{P})$ for all $i \in \{1, 2, 3, 4\}$. Moreover, m_1 is acceptable w.r.t. $\{c_2\}$ since m_1 is unattacked in $\mathcal{P}$ and every S-reduct of $\mathcal{P}$. However, $\{c_2, m_1\} \notin adm_{cp}^i(\mathcal{P})$. □

We argue that not satisfying Dung's fundamental lemma is no drawback in the case of CPAFs but rather allows us to deal with conditional preferences in a flexible way. For example, in the proof of Proposition 3.11 the set $\{c_2\}$ is admissible since, when considering only admissibility, we are not forced to include the unattacked m_1, i.e., we do not have to accept that Person 1 has a motive. But if we do accept that Person 1 has a motive, then we can no longer accept c_2. Note that the inclusion of unattacked arguments in CPAFs is handled via more restrictive approaches such as stable or preferred semantics, as usual.

Lastly, by the following proposition we express that every CPAF-semantics considered here generalizes their corresponding PAF-semantics, i.e., that CPAFs generalize PAFs.

Proposition 3.12. *Let $\mathcal{P} = (A, R, cond)$ be a CPAF such that the preference function cond maps every set of arguments to the same binary relation, i.e., there is some $\succ$ such that $cond(S) = \succ$ for all $S \subseteq A$. Let $\sigma \in \{cf, naive, adm, com, grd, prf, stb\}$. Then $\sigma_{cp}^i(\mathcal{P}) = \sigma_p^i(A, R, \succ)$. Furthermore, $naive\text{-}glb_{cp}^i(\mathcal{P}) = naive_p^i(A, R, \succ)$ and $prf\text{-}glb_{cp}^i(\mathcal{P}) = prf_p^i(A, R, \succ)$.*

[4] An argument x is acceptable w.r.t. a set of arguments S iff S defends x against all attackers.

3.2 Principle-Based Analysis

Principles play an important role in argumentation theory, as they allow us to examine the vast amount of semantics defined for AFs in a general way (Baroni and Giacomin 2007; Dvořák et al. 2024; van der Torre and Vesic 2018). In this section, we generalize the principles of Kaci et al. (Kaci et al. 2021) for PAFs (cf. Definitions 2.13, 2.14, 2.16, 2.18 in Subsection 2.3.2) to account for conditional preferences. We then investigate by which semantics these principles are satisfied, and show that there are differences to PAFs.

In the case of PAFs, adding more preferences to a framework $(A, R, \succ)$ means that we now deal with the PAF $(A, R, \succ \cup \succ')$. In the case of CPAFs, if we want $(A, R, cond')$ to have at least the same preferences as $(A, R, cond)$, we must require that $cond(S) \subseteq cond'(S)$ for all $S \subseteq A$. But if we only want to add a single preference $x \succ y$ to a CPAF we add $x \succ_S y$ to a subset S and leave the preferences associated with other subsets unchanged. Given the above considerations, generalizing the PAF-principles to CPAF-principles is quite straightforward. The notions of an argument's status and paths between two arguments in a CPAF are defined analogously to PAFs (cf. Definitions 2.15, 2.17), e.g., $status(x, \mathcal{P}) = cr$ iff x is contained in some but not all extensions of the CPAF $\mathcal{P}$.

Definition 3.13. *Let σ^i_{cp} be a CPAF-semantics. In the following, given a CPAF $(A, R, cond)$, we denote by $cond'$ an arbitrary function such that $cond(S) \subseteq cond'(S)$ for all $S \subseteq A$. Moreover, $cond_{(x,y)}$ is the same as $cond$ except that for some $S \subseteq A$ we have $(x, y) \in cond_{(x,y)}(S)$ but $(x, y), (y, x) \notin cond(S)$. Lastly, $cond_\emptyset(S) = \emptyset$ for all $S \subseteq A$.*

- *P1* (conflict-freeness): If $(x, y) \in R$ there is no $S \in \sigma^i_{cp}(A, R, cond)$ such that $\{x, y\} \subseteq S$.*
- *P2* (preference selects extensions): $\sigma^i_{cp}(A, R, cond') \subseteq \sigma^i_{cp}(A, R, cond)$.*
- *P3* (preference selects extensions 2): $\sigma^i_{cp}(A, R, cond) \subseteq \sigma^i_{cp}(A, R, cond_\emptyset)$.*
- *P4* (extension refinement): for all $S' \in \sigma^i_{cp}(A, R, cond')$ there is $S \in \sigma^i_{cp}(A, R, cond)$ s.t. $S \subseteq S'$.*
- *P5* (extension growth): $\bigcap(\sigma^i_{cp}(A, R, cond)) \subseteq \bigcap(\sigma^i_{cp}(A, R, cond'))$.*
- *P6* (number of extensions): $|\sigma^i_{cp}(A, R, cond')| \leq |\sigma^i_{cp}(A, R, cond)|$.*
- *P7* (status conservation): $status(x, (A, R, cond_{(x,y)})) \geq status(x, (A, R, cond))$.*
- *P8* (preference-based immunity): if $(x, x) \notin R$ and if cond is defined such that for all $S \subseteq A$ and all $y \in A \setminus \{x\}$ we have $x \succ_S y$ then $status(x, (A, R, cond)) \neq rej$.*
- *P9* (path preference influence 1): if there is no path from $x \in A$ to $y \in A$ in $(A, R, cond)$ then $\sigma^i_{cp}(A, R, cond) = \sigma^i_{cp}(A, R, cond_{(x,y)})$.*
- *P10* (path preference influence 2): if $(x, y) \notin R$ and $(y, x) \notin R$ then $\sigma^i_{cp}(A, R, cond) = \sigma^i_{cp}(A, R, cond_{(x,y)})$.*

The following lemma establishes some relationships between the CPAF-principles and is a generalization of known relationships between PAF-principles (Kaci et al. 2021).

Lemma 3.14. *If σ^i_{cp} satisfies $P2^*$ then it also satisfies $P3^*$, $P4^*$, and $P6^*$. If σ^i_{cp} always returns at least one extension, and if it satisfies $P2^*$, then it also satisfies $P5^*$.*

Proof. For $P3^*$, $P4^*$, and $P6^*$ this is easy to see. For $P5^*$ we argue this in detail: let σ^i_{cp} be a semantics that returns at least one extension, but does not satisfy $P5^*$. Thus, there is A, R, $cond$, $cond'$ with $cond(S) \subseteq cond'(S)$ for all $S \subseteq A$, such that $\bigcap(\sigma^i_{cp}(A, R, cond)) \not\subseteq \bigcap(\sigma^i_{cp}(A, R, cond'))$. Then there is $x \in A$ such that $x \in \bigcap(\sigma^i_{cp}(A, R, cond))$ but $x \notin \bigcap(\sigma^i_{cp}(A, R, cond'))$, i.e., there is $E \subseteq A$ such that $x \notin E$ and $E \in \sigma^i_{cp}(A, R, cond')$. Of course, $E \notin \sigma^i_{cp}(A, R, cond)$, otherwise $x \notin \bigcap(\sigma^i_{cp}(A, R, cond))$. But then $\sigma^i_{cp}(A, R, cond') \not\subseteq \sigma^i_{cp}(A, R, cond)$, i.e., σ^i_{cp} does not satisfy $P2^*$. □

Observe that, since CPAFs are a generalization of PAFs (cf. Proposition 3.12), a CPAF-semantics σ^i_{cp} cannot satisfy Pj^* if the corresponding PAF-semantics σ^i_p does not satisfy Pj. Moreover, it is obvious that $P1^*$ is still satisfied under Reductions 2, 3, and 4, as conflicts are not removed by these reductions even if we consider conditional preferences. We can also show that satisfaction of $P2$ carries over from PAFs to CPAFs.

Lemma 3.15. *If σ^i_p satisfies $P2$ then σ^i_{cp} satisfies $P2^*$.*

Proof. By contrapositive, assume σ^i_{cp} does not satisfy $P2^*$. Then there is a CPAF $\mathcal{P} = (A, R, cond)$ and $cond'$ with $cond(S) \subseteq cond'(S)$ for all $S \subseteq A$ such that $\sigma^i_{cp}(A, R, cond') \not\subseteq \sigma^i_{cp}(A, R, cond)$. Thus, there is $E \subseteq A$ such that $E \in \sigma^i_{cp}(A, R, cond')$ but $E \notin \sigma^i_{cp}(A, R, cond)$. Then $E \in \sigma(\mathcal{R}_i(A, R, cond'(E)))$ but $E \notin \sigma(\mathcal{R}_i(A, R, cond(E)))$, i.e., σ^i_p does not satisfy $P2$. □

Lemma 3.15 implies that complete and stable semantics satisfy $P2^*$ on CPAFs under Reduction 3. By Lemma 3.14 these semantics also satisfy $P3^*$, $P4^*$, and $P6^*$. However, we cannot use Lemma 3.14 to show that complete semantics satisfy $P5^*$, since conditional preferences allow for frameworks without complete extensions. Indeed, we can find a counter-example in this case. Counterexamples for the satisfaction of various principles can also be found for grounded semantics, both variants of the preferred semantics, and even stable semantics in the case of $P8^*$.

Lemma 3.16. *The following holds:*

- *grd^i_{cp}, where $i \in \{1, 2, 3, 4\}$, does not satisfy any of $P4^*$, $P5^*$, or $P6^*$;*
- *com^3_{cp} does not satisfy $P5^*$;*
- *prf^3_{cp} and $prf\text{-}glb^3_{cp}$ do not satisfy $P6^*$;*
- *σ^i_{cp}, where for $\sigma \in \{com, grd, prf, stb\}$ and $i \in \{1, 2, 3, 4\}$, does not satisfy $P8^*$.*

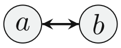

cond:
$a \succ_{\emptyset} b$
$a \succ_{\{a\}} b$

cond':
$a \succ'_{\emptyset} b$
$a \succ'_{\{a\}} b$
$b \succ'_{\{b\}} a$

(a) Counterexample grd^i_{cp} ($P4^*$, $P5^*$, $P6^*$)

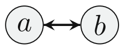

cond:
$a \succ_{\emptyset} b$
$a \succ_{\{b\}} b$

cond':
$a \succ'_{\emptyset} b$
$a \succ'_{\{b\}} b$
$b \succ'_{\{a\}} a$

(b) Counterexample com^3_{cp} ($P5^*$)

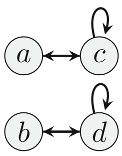

$cond(S) = \emptyset$ for all $S \subseteq A$

cond':
$c \succ'_{\{b\}} a$
$d \succ'_{\{a\}} b$
$c \succ'_{\{a,b\}} a$
$d \succ'_{\{a,b\}} b$

(c) Counterexample prf^3_{cp}, $prf\text{-}glb^3_{cp}$ ($P6^*$)

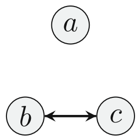

cond:
$a \succ_S b$ for all $S \subseteq A$
$a \succ_S c$ for all $S \subseteq A$
$b \succ_{\{a\}} c$
$b \succ_{\{a,c\}} c$
$c \succ_{\{a,b\}} b$

(d) Counterexample com^i_{cp}, grd^i_{cp}, prf^i_{cp}, stb^i_{cp} ($P8^*$)

Figure 3.4: Counterexamples used in Lemma 3.16.

Proof. We provide counterexamples for all cases.

- For grd^i_{cp} and $P4^*$, $P5^*$, $P6^*$, consider $A = \{a, b\}$, $R = \{(a, b), (b, a)\}$, and $cond/cond'$ as shown in Figure 3.4a. Then $grd^i_{cp}(A, R, cond) = \{\{a\}\}$ while $grd^i_{cp}(A, R, cond') = \{\{a\}, \{b\}\}$.
- For com^3_{cp} and $P5^*$, consider $A = \{a, b\}$, $R = \{(a, b), (b, a)\}$, and $cond/cond'$ as shown in Figure 3.4b. Then $com^3_{cp}(A, R, cond) = \{\{a\}\}$ but $com^3_{cp}(A, R, cond') = \emptyset$.
- For prf^3_{cp}, $prf\text{-}glb^3_{cp}$ and $P6^*$, consider $A = \{a, b, c, d\}$, $R = \{(a, c), (c, a), (b, d), (d, b), (c, c), (d, d)\}$, and $cond/cond'$ as shown in Figure 3.4c. Then $prf^3_{cp}(A, R, cond) = prf\text{-}glb^3_{cp}(A, R, cond) = \{\{a, b\}\}$ while $prf^3_{cp}(A, R, cond') = prf\text{-}glb^3_{cp}(A, R, cond') = \{\{a\}, \{b\}\}$.
- Regarding $P8^*$, consider the CPAF $\mathcal{P} = (A, R, cond)$ shown in Figure 3.4d. Note that $(a, a) \notin R$ and $a \succ_S y$ for all $S \subseteq A$ and all $y \notin \{a\}$. Observe that b is unattacked in $\mathcal{R}^{\{a\}}_i(\mathcal{P})$. Thus, $\{a\} \notin \sigma(\mathcal{R}^{\{a\}}_i(\mathcal{P}))$ for $\sigma \in \{com, grd, prf, stb\}$. Moreover, b is not defended against c in $\mathcal{R}^{\{a,b\}}_i(\mathcal{P})$. Analogously for c in $\mathcal{R}^{\{a,c\}}_i(\mathcal{P})$. Thus, $status(a, \mathcal{P}) = rej$. □

We have now fully investigated the first six principles. It remains to examine principles 7-10, of which we so far only know that $P8^*$ is not satisfied by most semantics. It turns out that $P8^*$ is retained when using preferred semantics with global maximization.

Lemma 3.17. *$prf\text{-}glb^i_{cp}$ satisfies $P8^*$ for $i \in \{1, 2, 4\}$.*

Proof. Let $\mathcal{P} = (A, R, cond)$ be a CPAF containing an argument $x \in A$ such that $(x, x) \notin R$ and for all $S \subseteq A$ and all $y \in A \setminus \{x\}$ we have $x \succ_S y$. Specifically, this means that $x \succ_{\{x\}} y$ for all $y \in A \setminus \{x\}$. Then, by definition of Reduction $i \in \{1, 2, 4\}$, x defends itself against all attacks in $\mathcal{R}_i^{\{x\}}(\mathcal{P})$. Thus, $\{x\} \in adm^i_{cp}(\mathcal{P})$. By this and the definition of *prf-glb*, there is some $E \in prf\text{-}glb^i_{cp}(\mathcal{P})$ such that $x \in E$. □

Now we turn our attention to $P7^*$, where, analogously to $P2^*$ (cf. Lemma 3.15), it turns out that satisfaction carries over from PAFs to CPAFs.

Lemma 3.18. *If σ^i_p satisfies $P7$ then σ^i_{cp} satisfies $P7^*$.*

Proof. By contrapositive, assume σ^i_{cp} does not satisfy $P7^*$. Then there is a CPAF $\mathcal{P} = (A, R, cond)$ such that $status(x, (A, R, cond_{(x,y)})) < status(x, (A, R, cond))$. This means there is some $S \subseteq A \cup \{x\}$ for which $S \in \sigma^i_{cp}(A, R, cond)$ but $S \notin \sigma^i_{cp}(A, R, cond_{(x,y)})$. By the definition of CPAF-semantics this means that $S \in \sigma^i_p(A, R, cond(S))$ but $S \notin \sigma^i_p(A, R, cond(S) \cup \{(x, y)\})$, i.e., σ^i_p does not satisfy $P7$. □

However, unlike in the case of $P2^*$, the above lemma does not constitute an exhaustive investigation of $P7^*$. The reason for this is that $P7$, in contrast to $P2$, is satisfied by preferred semantics on PAFs (cf. Table 2.2). Lemma 3.18 only allows us to conclude that prf^i_{cp} satisfies $P7^*$, but it says nothing about the satisfaction of $prf\text{-}glb^i_{cp}$. We find that $P7^*$ is satisfied also when maximizing admissible sets globally.

Lemma 3.19. *$prf\text{-}glb^i_{cp}$ satisfies $P7^*$.*

Proof. Let $\mathcal{P} = (A, R, cond)$ be an arbitrary CPAF and $x \in A$. Let $\mathcal{P}_{(x,y)} = (A, R, cond_{(x,y)})$ as specified in Definition 3.13. There are three possible cases:

1. $status(x, \mathcal{P}) = rej$. Then $status(x, \mathcal{P}_{(x,y)}) \geq status(x, \mathcal{P})$ trivially holds.
2. $status(x, \mathcal{P}) = cr$. Then there is some $S \in adm^i_{cp}(\mathcal{P})$ with $x \in S$. We distinguish two cases:
   a) $cond(S) = cond_{(x,y)}(S)$. Then clearly $S \in adm^i_{cp}(\mathcal{P}_{(x,y)})$.
   b) $cond(S) \neq cond_{(x,y)}(S)$. Then $cond_{(x,y)}$ is the same as $cond$ except that $(x, y) \in cond_{(x,y)}(S)$ but $(x, y), (y, x) \notin cond(S)$ for some $y \in A \setminus \{x\}$. Adding the preference $x \succ y$ via $cond_{(x,y)}$ does not introduce any new attacks against S in $\mathcal{R}^S_i(\mathcal{P}_{(x,y)})$, no matter which of the preference reductions we consider. Thus, $S \in adm^i_{cp}(\mathcal{P}_{(x,y)})$.

   In both cases we have $S \in adm^i_{cp}(\mathcal{P}_{(x,y)})$ and therefore $T \in prf\text{-}glb^i_{cp}(\mathcal{P}_{(x,y)})$ for some $T \supseteq S$. Thus, $status(x, \mathcal{P}_{(x,y)}) \geq cr$.

Table 3.1: Satisfaction of CPAF-principles. C stands for complete, G for grounded, P for preferred (both prf^i_{cp} and $\mathit{prf\text{-}glb}^i_{cp}$), and S for stable. P_g indicates that $\mathit{prf\text{-}glb}^i_{cp}$ satisfies the principle but prf^i_{cp} does not. If a cell is marked with $\times$ then none of the investigated semantics satisfy this principle.

	$\mathcal{R}_1$	$\mathcal{R}_2$	$\mathcal{R}_3$	$\mathcal{R}_4$
$P1^*$ (conflict-freeness)	$\times$	$CGPS$	$CGPS$	$CGPS$
$P2^*$ (preference selects extensions)	$\times$	$\times$	CS	$\times$
$P3^*$ (preference selects extensions 2)	$\times$	$\times$	CS	$\times$
$P4^*$ (extension refinement)	$\times$	$\times$	CS	$\times$
$P5^*$ (extension growth)	$\times$	$\times$	$\times$	$\times$
$P6^*$ (number of extensions)	$\times$	$\times$	CS	$\times$
$P7^*$ (status conservation)	$CGPS$	$CGPS$	$CGPS$	$CGPS$
$P8^*$ (preference-based immunity)	P_g	P_g	$\times$	P_g
$P9^*$ (path preference influence 1)	$CGPS$	$CGPS$	$CGPS$	$CGPS$
$P10^*$ (path preference influence 2)	$CGPS$	$CGPS$	$CGPS$	$CGPS$

3. $\mathit{status}(x, \mathcal{P}) = \mathit{sk\text{-}cr}$. By the same line of reasoning as in case (2) we have $\mathit{status}(x, \mathcal{P}_{(x,y)}) \geq \mathit{cr}$. Towards a contradiction, assume that $\mathit{status}(x, \mathcal{P}_{(x,y)}) \neq \mathit{sk\text{-}cr}$. Then there is some $S \in \mathit{prf\text{-}glb}^i_{cp}(\mathcal{P}_{(x,y)})$ such that $x \notin S$. Since $\mathit{status}(x, \mathcal{P}) = \mathit{sk\text{-}cr}$ we know that $S \notin \mathit{prf\text{-}glb}^i_{cp}(\mathcal{P})$. One of the following must be the case:

   a) $S \notin \mathit{cf}^i_{cp}(\mathcal{P})$. Since $S \in \mathit{prf\text{-}glb}^i_{cp}(\mathcal{P}_{(x,y)})$ it must be that $S \in \mathit{cf}^i_{cp}(\mathcal{P}_{(x,y)})$. Then it must be that the additional preference in $\mathit{cond}_{(x,y)}$ removes a conflict between two arguments in S. But this preference is $x \succ y$ for some $y \in A \setminus \{x\}$. Thus, $x \in S$. Contradiction.

   b) $S \in \mathit{cf}^i_{cp}(\mathcal{P})$ but $S \notin \mathit{adm}^i_{cp}(\mathcal{P})$. Since $S \in \mathit{prf\text{-}glb}^i_{cp}(\mathcal{P}_{(x,y)})$ it must be that $S \in \mathit{adm}^i_{cp}(\mathcal{P}_{(x,y)})$. However, the additional preference $x \succ y$ added via $\mathit{cond}_{(x,y)}$ at most adds an attack (x, y). Since $x \notin S$ this means that S is still not defended in $\mathcal{R}^S_i(\mathcal{P}_{(x,y)})$. Contradiction.

   c) $S \in \mathit{adm}^i_{cp}(\mathcal{P})$ but there is $T \supset S$ such that $T \in \mathit{prf\text{-}glb}^i_{cp}(\mathcal{P})$. Since $\mathit{status}(x, \mathcal{P}) = \mathit{sk\text{-}cr}$ we know that $x \in T$. Since $T \in \mathit{prf\text{-}glb}^i_{cp}(\mathcal{P})$ we have $T \in \mathit{adm}^i_{cp}(\mathcal{P})$. By the same line of reasoning as in case (2) we can conclude that $T \in \mathit{adm}^i_{cp}(\mathcal{P}_{(x,y)})$ and therefore $S \notin \mathit{prf\text{-}glb}^i_{cp}(\mathcal{P}_{(x,y)})$. Contradiction.

   In all three cases we arrive at a contradiction. Thus, $\mathit{status}(x, \mathcal{P}_{(x,y)}) = \mathit{sk\text{-}cr}$. □

Lastly, we must consider principles 9 and 10. Below, we show that they retain the satisfaction of all principles under all considered semantics.

Lemma 3.20. *σ_{cp}^{i} satisfies $P9^*$ and $P10^*$ for $\sigma \in \{com, grd, prf, prf\text{-}glb, stb\}$.*

Proof. Let $\mathcal{P} = (A, R, cond)$ be an arbitrary CPAF and $x \in A$. Let $\mathcal{P}_{(x,y)} = (A, R, cond_{(x,y)})$ as specified in Definition 3.13. Note that the premise for $P9^*$ (there is no path from $x \in A$ to $y \in A$) implies the premise of $P10^*$ ($(x, y) \notin R$ and $(y, x) \notin R$). If $(x, y) \notin R$ and $(y, x) \notin R$ then the additional preference $x \succ y$ added via $cond_{(x,y)}$ does not delete or add any attacks, regardless of which preference reduction we consider. This means that $\mathcal{R}_i^S(\mathcal{P}) = \mathcal{R}_i^S(\mathcal{P}_{(x,y)})$ for all $S \subseteq A$ and all $i \in \{1, 2, 3, 4\}$ and therefore $\sigma_{cp}^{i}(A, R, cond) = \sigma_{cp}^{i}(A, R, cond_{(x,y)})$ for all $i \in \{1, 2, 3, 4\}$. □

The above results constitute an exhaustive investigation of the ten CPAF-principles for all semantics considered in this chapter. Thus, we can conclude:

Theorem 3.21. *The satisfaction of CPAF-principles depicted in Table 3.1 holds.*

To summarize, complete and stable semantics preserve the satisfaction of PAF-principles in most cases. Grounded semantics no longer satisfies any of the principles 1-6 on CPAFs except $P1^*$ (conflict-freeness) since grounded extensions are not unique on CPAFs, and since there are even CPAFs without a grounded extension (cf. Lemma 3.16). Unlike on PAFs, complete semantics does not satisfy $P5^*$ (extension growth) under Reduction 3. Furthermore, neither variant of preferred semantics satisfies $P6^*$ (number of extensions) under Reduction 3. As for principles 7-10, we note that only $P8^*$ is no longer satisfied by all semantics.

3.3 Complexity

The computational complexity of Dung-style AFs and various generalizations thereof has received considerable attention in the literature (Dvořák and Dunne 2018). Indeed, complexity results give insights into the expressiveness of specific argumentation formalisms and help to find appropriate methods for solving a given problem. Note that AFs and PAFs have the same properties with regards to complexity, i.e., none of the four preference reductions result in a higher complexity when considering unconditional preferences in the setting of Dung-AFs. The situation is not as clear when dealing with conditional preferences. As we have seen in previous sections, CPAFs do not necessarily have unique grounded extensions, there are CPAFs without any complete extensions, and there is more than one way of dealing with subset maximization (recall the *naive*/*naive-glb* and *prf*/*prf-glb* semantics). In this section, we show that these differences between CPAFs and AFs/PAFs have an impact on complexity.

We define $Ver_{\sigma,i}^{CPAF}$, $Cred_{\sigma,i}^{CPAF}$, and $Skept_{\sigma,i}^{CPAF}$ analogously to Ver_{σ}^{AF}, $Cred_{\sigma}^{AF}$, and $Skept_{\sigma}^{AF}$ (cf. Definition 2.8 in Subsection 2.3.1), with the difference that the framework in question is now a CPAF instead of an AF and that we appeal to the σ_{cp}^{i} semantics of Definitions 3.3 and 3.5 rather than the AF-semantics of Definition 2.6:

Definition 3.22. *Given a CPAF-semantics σ^i_{cp} we define the following decision problems:*

- Credulous Acceptance *($Cred^{CPAF}_{\sigma,i}$): given a CPAF $\mathcal{P}$ and an argument x, is $x \in S$ for some $S \in \sigma^i_{cp}(\mathcal{P})$?*
- Skeptical Acceptance *($Skept^{CPAF}_{\sigma,i}$): given a CPAF $\mathcal{P}$ and an argument x, is $x \in S$ in all $S \in \sigma^i_{cp}(\mathcal{P})$?*
- Verification *($Ver^{CPAF}_{\sigma,i}$): given a CPAF $\mathcal{P}$ and a set of arguments S, is $S \in \sigma^i_{cp}(\mathcal{P})$?*

In the interest of generality, we did not impose a specific method to represent conditional preferences in the previous sections. However, when analyzing the computational complexity of CPAFs, it is necessary to decide on more specific representations if tight bounds are to be found. Therefore, we will assume conditional preferences to be expressed succinctly as arbitrary propositional formulas. Note that if preferences would be stored explicitly for each possible set of arguments, the input size of our problems would always be exponentially larger than the underlying AF itself, and thus some decision problems for CPAFs would be in lower complexity classes than their counterparts for AFs.

Specifically, given the set of arguments A in a framework $\mathcal{P} = (A, R, cond)$ we allow a finite number of rules $\varphi \Rightarrow x \succ y$ where $x, y \in A$ and φ is a propositional formula built from atoms in A and the usual connectives $(\neg, \wedge, \vee)$. As for the meaning of these rules, we define that for some $S \subseteq A$ we have $x \succ_S y$ iff there is a rule $\varphi \Rightarrow x \succ y$ such that $S \models \varphi$ and there is no rule $\varphi' \Rightarrow y \succ x$ such that $S \models \varphi'$.[5] Observe that, given $S \subseteq A$, it is possible to compute $\succ_S$ in polynomial time with respect to the size of the given framework $\mathcal{P}$ since $S \models \varphi$ can be decided in polynomial time for each rule $\varphi \Rightarrow x \succ y$.[6]

Our complexity results are summarized in Table 3.2. Note that problems for *naive*/*naive-glb* semantics become harder only under Reduction 1. Intuitively, this is because Reduction 1 can remove conflicts between arguments altogether, unlike Reductions 2-4. Observe that *naive-glb* under Reduction 1 is the only semantics for which the verification problem becomes harder (coNP-complete) compared to AFs (in P). As a result, skeptical acceptance for *naive-glb* is Π^{P}_2-complete, i.e., the complexity rises by two levels in the polynomial hierarchy compared to the case of AFs. For complete semantics, skeptical acceptance is now coNP-complete regardless of which preference reduction is used. With respect to grounded semantics we see an increase in complexity for both credulous acceptance (NP-complete) and skeptical acceptance (coNP-complete). Lastly, for preferred semantics with local maximization, credulous acceptance rises by one level in the polynomial hierarchy compared to AFs.

[5] A set of atoms S can be seen as an interpretation, with x set to true under S iff $x \in S$.

[6] In fact, for our membership results the explicit representation of rules using propositional formulas is not necessary. It suffices to have some representation such that, given $S \subseteq A$, we can determine $\succ_S$ in polynomial time with respect to the size of $\mathcal{P} = (A, R, cond)$. However, for hardness results, a more concrete representation such as via our rules is necessary.

Table 3.2: Complexity of CPAFs with conditional preferences represented via finitely many rules of the form $\varphi \Rightarrow x \succ y$. Underlines indicate a rise in complexity compared to AFs.

σ	$Cred^{CPAF}_{\sigma,1}$ / $Cred^{CPAF}_{\sigma,j\in\{2,3,4\}}$	$Skept^{CPAF}_{\sigma,1}$ / $Skept^{CPAF}_{\sigma,j\in\{2,3,4\}}$	$Ver^{CPAF}_{\sigma,1}$ / $Ver^{CPAF}_{\sigma,j\in\{2,3,4\}}$
cf	in P	trivial	in P
naive	$\underline{\text{NP-c}}$/in P	$\underline{\text{coNP-c}}$/in P	in P
naive-glb	in P	$\underline{\Pi_2^\mathsf{P}\text{-c}}$/in P	$\underline{\text{coNP-c}}$/in P
adm	NP-c	trivial	in P
com	NP-c	$\underline{\text{coNP-c}}$	in P
grd	$\underline{\text{NP-c}}$	$\underline{\text{coNP-c}}$	P-c
stb	NP-c	coNP-c	in P
prf	$\underline{\Sigma_2^\mathsf{P}\text{-c}}$	Π_2^P-c	coNP-c
prf-glb	NP-c	Π_2^P-c	coNP-c

Theorem 3.23. *The complexity results for CPAFs depicted in Table 3.2 hold.*

The remainder of this section is dedicated to proving Theorem 3.23. We first consider the verification problem, for which most semantics have the same complexity as in the case of AFs.

Lemma 3.24. $Ver^{CPAF}_{\sigma,i}$*, where* $i \in \{1,2,3,4\}$*, has the same complexity properties with regards to membership and hardness as* Ver^{AF}_{σ} *for* $\sigma \in \{cf, naive, adm, com, grd, stb, prf\}$*. Moreover,* $Ver^{CPAF}_{prf\text{-}glb,i}$ *is* coNP*-complete for* $i \in \{1,2,3,4\}$*.*

Proof. Hardness follows from the fact that CPAFs are a generalization of AFs. Membership for $\sigma \in \{cf, naive, adm, com, grd, stb, prf\}$: given a CPAF $\mathcal{P}$ and a set of arguments S, we can determine $\succ_S$ and therefore also $\mathcal{R}_i^S(\mathcal{P})$ in polynomial time. It then suffices to check whether $S \in \sigma(\mathcal{R}_i^S(\mathcal{P}))$. Membership for $\sigma = prf\text{-}glb$: let $(\mathcal{P}, S)$ be an arbitrary instance of $Ver^{CPAF}_{prf\text{-}glb,i}$, i.e., $\mathcal{P} = (A, R, cond)$ is a CPAF and $S \subseteq A$ is a set of arguments. First, check in polynomial time whether $S \in adm^i_{cp}(\mathcal{P})$. Then, in coNP-time, check that for all T we have either $T \subseteq S$ or $T \notin adm^i_{cp}(\mathcal{P})$. □

For naive semantics with global maximization (*naive-glb*) we see a rise in complexity, but only when using Reduction 1. The following proof makes use of Reduction 1's ability to delete conflicts between arguments. By SAT we denote the NP-complete satisfiability problem for propositional formulas, and by UNSAT we denote its complementary problem which is coNP-complete.

Lemma 3.25. $Ver^{CPAF}_{naive\text{-}glb,j}$ *is in* P *for* $j \in \{2,3,4\}$*.* $Ver^{CPAF}_{naive\text{-}glb,1}$ *is* coNP*-complete.*

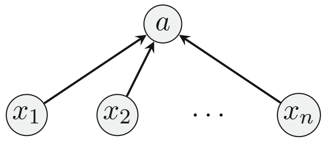

$\varphi \Rightarrow a \succ x_i$ for all $1 \leq i \leq n$

Figure 3.5: Construction used in the proof of Lemma 3.25. Given a formula φ over variables $X = \{x_1, x_2, \dots, x_n\}$, a CPAF $\mathcal{P}$ is constructed such that φ is unsatisfiable iff $\{a\} \in \textit{naive-glb}^1_{cp}(\mathcal{P})$.

Proof. Let $(\mathcal{P}, S)$ be an arbitrary instance of $Ver^{CPAF}_{naive\text{-}glb,i}$, i.e., $\mathcal{P} = (A, R, cond)$ is a CPAF and $S \subseteq A$ is a set of arguments. For Reductions 2, 3, and 4 it suffices to check whether $S \in naive(A, R)$ since these reductions cannot remove or add conflicts.

We now turn our attention to Reduction 1. coNP-membership: first check whether $S \in cf^1_{cp}(\mathcal{P})$. Then, in coNP-time, check that for all T we have either $T \subseteq S$ or $T \notin cf^1_{cp}(\mathcal{P})$.

To show coNP-hardness we provide a reduction from Unsat: let φ be an arbitrary propositional formula over variables X. Let a be a fresh variable, i.e., $a \notin X$. We construct an instance $(\mathcal{P}, \{a\})$ of $Ver^{CPAF}_{naive\text{-}glb,1}$ as follows: $\mathcal{P} = (A, R, cond)$ with $A = X \cup \{a\}$, $R = \{(x, a) \mid x \in X\}$, and $cond$ defined by the rules $\varphi \Rightarrow a \succ x$ for $x \in X$, i.e., $a \succ_S x$ iff $S \models \varphi$. Figure 3.5 depicts the above construction. We now show that φ is unsatisfiable iff $\{a\} \in \textit{naive-glb}^1_{cp}(\mathcal{P})$ (i.e., $(\mathcal{P}, \{a\})$ is a yes-instance of $Ver^{CPAF}_{naive\text{-}glb,1}$).

- Suppose φ is unsatisfiable. This means that for all $x \in X$ and all $S \subseteq A$ we have $a \not\succ_S x$, i.e., for each $x \in X$ the attack (x, a) is present in $\mathcal{R}^S_1(\mathcal{P})$. Thus, there is no conflict-free set containing a other than $\{a\}$ which implies $\{a\} \in \textit{naive-glb}^1_{cp}(\mathcal{P})$.
- Suppose φ is satisfiable. Then there is an interpretation $I \subseteq X$ such that $I \models \varphi$. We assume that $I \neq \emptyset$. This is permissible since we can check in polynomial time whether $\emptyset$ satisfies φ, and if this is the case, return a trivial no-instance of $Ver^{CPAF}_{naive\text{-}glb,1}$. Consider $S = I \cup \{a\}$. Since a does not appear in φ we have $S \models \varphi$ and therefore $a \succ_S x$ for all $x \in X$. Thus, $S \in cf^1_{cp}(\mathcal{P})$ and, since $\{a\} \subset S$, $\{a\} \notin \textit{naive-glb}^1_{cp}(\mathcal{P})$. □

We now consider credulous and skeptical acceptance, starting with semantics based solely on conflict-freeness. Let us first cover the cases in which there is no rise in complexity.

Lemma 3.26. *$Cred^{CPAF}_{\sigma,i}$ is in* P *for $\sigma \in \{cf, naive\text{-}glb\}$, $i \in \{1, 2, 3, 4\}$. $Cred^{CPAF}_{naive,j}$ is in* P *for $j \in \{2, 3, 4\}$.*

Proof. Let $(\mathcal{P}, x)$ be an instance of $Cred^{CPAF}_{\sigma,i}$. For $\sigma = cf$ it suffices to check whether x is self-attacking in the underlying AF of $\mathcal{P}$, since self-attacks are not removed by any of the four reductions. For $\sigma = naive\text{-}glb$ it suffices to test whether $(\mathcal{P}, x)$ is a yes-instance

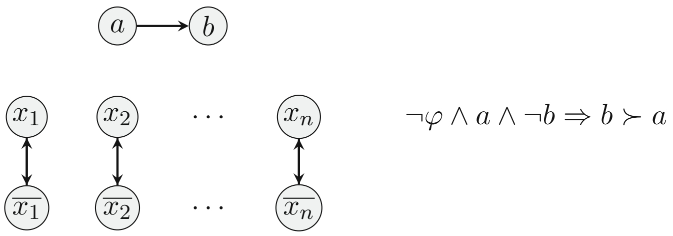

Figure 3.6: Construction used in the proof of Lemma 3.28. Given a formula φ over variables $X = \{x_1, x_2, \ldots, x_n\}$, a CPAF $\mathcal{P}$ is constructed such that φ is satisfiable iff $(\mathcal{P}, a)$ is a yes-instance of $Cred^{CPAF}_{naive,1}$ iff $(\mathcal{P}, b)$ is a no-instance of $Skept^{CPAF}_{naive,1}$.

of $Cred^{CPAF}_{cf,i}$. For $\sigma = naive$ and Reductions 2, 3, and 4 it is enough to check whether x appears in a naive set of the underlying AF, since these reductions cannot remove conflicts. □

Lemma 3.27. *$Skept^{CPAF}_{cf,i}$ is trivial for $i \in \{1, 2, 3, 4\}$. $Skept^{CPAF}_{\sigma,j}$ is in* P *for $\sigma \in \{naive, naive\text{-}glb\}$ and $j \in \{2, 3, 4\}$.*

Proof. Let $(\mathcal{P}, x)$ be an arbitrary instance of $Skept_{naive\text{-}glb,i}$, i.e., $\mathcal{P} = (A, R, cond)$ is a CPAF and $x \in A$ is an argument. Note that $\emptyset$ is always conflict-free in $\mathcal{R}^{\emptyset}_i(\mathcal{P})$, i.e., $(\mathcal{P}, x)$ is trivially a no-instance. For $\sigma \in \{naive, naive\text{-}glb\}$ and Reductions 2, 3, and 4 it is enough to solve the problem on the underlying AF, since these reductions cannot remove conflicts. □

For naive semantics with local maximization the complexity rises by one level in the polynomial hierarchy under Reduction 1.

Lemma 3.28. *$Cred^{CPAF}_{naive,1}$ is* NP*-complete and $Skept^{CPAF}_{naive,1}$ is* coNP*-complete.*

Proof. We will consider the complementary problem of $Skept^{CPAF}_{naive,1}$ and show that it is NP-complete since this allows us to prove both results simultaneously.

NP-Membership: given a CPAF $\mathcal{P} = (A, R, cond)$ and an argument $x \in A$, guess a set $S \subseteq A$ and, in polynomial time, check whether $S \in naive^1_{cp}(\mathcal{P})$ and $x \in S$ (resp. $S \in naive^1_{cp}(\mathcal{P})$ and $x \notin S$).

NP-hardness by reduction from SAT: let φ be an arbitrary propositional formula over a set of variables X. Let a and b be fresh atoms, i.e., $a, b \notin X$. We construct an instance $(\mathcal{P}, a)$ of $Cred^{CPAF}_{naive,1}$ as follows: $\mathcal{P} = (A, R, cond)$ with $A = X \cup \{\overline{x} \mid x \in X\} \cup \{a, b\}$, $R = \{(x, \overline{x}), (\overline{x}, x) \mid x \in X\} \cup \{(a, b)\}$, and $cond$ defined by the rule $\neg\varphi \wedge a \wedge \neg b \Rightarrow b \succ a$, i.e., $b \succ_S a$ iff $S \models \neg\varphi \wedge a \wedge \neg b$. The above construction is visualized in Figure 3.6. We show that φ is satisfiable iff $(\mathcal{P}, a)$ is a yes-instance of $Cred^{CPAF}_{naive,1}$ iff $(\mathcal{P}, b)$ is a no-instance of $Skept^{CPAF}_{naive,1}$:

- Assume φ is satisfiable. Then there is an interpretation $I \subseteq X$ such that $I \models \varphi$. Then also $I \cup \{a\} \models \varphi$ and $I \cup \{a\} \not\models \neg\varphi \wedge a \wedge \neg b$. Let $S = I \cup \{\overline{x} \mid x \notin I\} \cup \{a\}$. Note that $b \not\succ_S a$, which means that a and b are in conflict in $\mathcal{R}_1^S(\mathcal{P})$. Furthermore, for all $x \in X$, we have either $x \in S$ or $\overline{x} \in S$, but not both. Thus, $S \in \mathit{naive}(\mathcal{R}_1^S(\mathcal{P}))$. Note that $a \in S$ but $b \notin S$, i.e., $(\mathcal{P}, a)$ is a yes-instance of $\mathit{Cred}_{\mathit{naive},1}^{\mathit{CPAF}}$ but $(\mathcal{P}, b)$ is a no-instance of $\mathit{Skept}_{\mathit{naive},1}^{\mathit{CPAF}}$.

- Assume φ is unsatisfiable. Consider some $S \subseteq A$. If there is some $x \in X$ such that neither $x \in S$ nor $\overline{x} \in S$, then $S \notin \mathit{naive}(\mathcal{R}_1^S(\mathcal{P}))$. Likewise, if there is some $x \in X$ such that both $x \in S$ and $\overline{x} \in S$ then $S \notin \mathit{cf}(\mathcal{R}_1^S(\mathcal{P}))$ and therefore $S \notin \mathit{naive}(\mathcal{R}_1^S(\mathcal{P}))$. It remains to consider sets S in which for all $x \in X$ we have either $x \in S$ or $\overline{x} \in S$ but not both. Given such a set, we consider four cases:

  1. $a \notin S$ and $b \notin S$. Then $S \notin \mathit{naive}(\mathcal{R}_1^S(\mathcal{P}))$ since $S \cup \{a\} \in \mathit{cf}(\mathcal{R}_1^S(\mathcal{P}))$.
  2. $a \notin S$ and $b \in S$. Then $S \not\models \neg\varphi \wedge a \wedge \neg b$, i.e., $b \not\succ_S a$. This means that the attack (a, b) is present in $\mathcal{R}_1^S(\mathcal{P})$. Note that every argument is either in S or in conflict with S in $\mathcal{R}_1^S(\mathcal{P})$. Moreover, $S \in \mathit{cf}(\mathcal{R}_1^S(\mathcal{P}))$. We can conclude that $S \in \mathit{naive}(\mathcal{R}_1^S(\mathcal{P}))$.
  3. $a \in S$ and $b \notin S$. Then $S \models \neg\varphi \wedge a \wedge \neg b$, i.e., $b \succ_S a$. This means that the attack (a, b) is deleted in $\mathcal{R}_1^S(\mathcal{P})$, which further implies that $S \cup \{b\} \in \mathit{cf}(\mathcal{R}_1^S(\mathcal{P}))$. Thus, $S \notin \mathit{naive}(\mathcal{R}_1^S(\mathcal{P}))$.
  4. $a \in S$ and $b \in S$. Then $S \not\models \neg\varphi \wedge a \wedge \neg b$, i.e., $b \not\succ_S a$. This means that the attack (a, b) is present in $\mathcal{R}_1^S(\mathcal{P})$. Thus, $S \notin \mathit{cf}(\mathcal{R}_1^S(\mathcal{P}))$ and therefore $S \notin \mathit{naive}(\mathcal{R}_1^S(\mathcal{P}))$.

  In conclusion, if $S \in \mathit{naive}_{cp}^1(\mathcal{P})$ then $a \notin S$ and $b \in S$. Thus, $(\mathcal{P}, a)$ is a no-instance of $\mathit{Cred}_{\mathit{naive},1}^{\mathit{CPAF}}$ but $(\mathcal{P}, b)$ is a yes-instance of $\mathit{Skept}_{\mathit{naive},1}^{\mathit{CPAF}}$. □

For naive semantics with global maximization, skeptical acceptance rises by even two levels in the polynomial hierarchy under Reduction 1. The reason for this is the increased complexity of the verification problem in this case (cf. Lemma 3.25). Recall that $\mathrm{QBF}_{\forall}^2$ denotes the Π_2^{P}-complete problem of deciding whether a quantified boolean formula of the form $\forall Y \exists Z \varphi$, where φ is a formula over $Y \cup Z$, is true (cf. Section 2.2).

Lemma 3.29. *$\mathit{Skept}_{\mathit{naive\text{-}glb},1}^{\mathit{CPAF}}$ is Π_2^{P}-complete.*

Proof. Σ_2^{P}-membership for the complementary problem of $\mathit{Skept}_{\mathit{naive\text{-}glb},1}^{\mathit{CPAF}}$: given a CPAF $\mathcal{P} = (A, R, \mathit{cond})$ and an argument x, guess a set $S \subset A$ and check that $x \notin S$ and, in coNP-time, that $S \in \mathit{naive\text{-}glb}_{cp}^1(\mathcal{P})$.

Π_2^{P}-hardness: let $\forall Y \exists Z \varphi$ be an arbitrary instance of $\mathrm{QBF}_{\forall}^2$ over variables $Y = \{y_1, \dots, y_n\}$ and $Z = \{z_1, \dots, z_m\}$. Let $X = Y \cup Z$. Using fresh variables a and z_{m+1} we construct an instance $(\mathcal{P}, a)$ of $\mathit{Skept}_{\mathit{naive\text{-}glb},1}^{\mathit{CPAF}}$ where $\mathcal{P} = (A, R, \mathit{cond})$ with

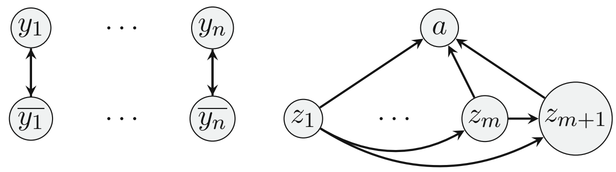

Figure 3.7: Construction used in the proof of Lemma 3.29. Given a quantified boolean formula $\forall Y \exists Z \varphi$ over variables $Y = \{y_1, \dots, y_n\}$ and $Z = \{z_1, \dots, z_m\}$, a CPAF $\mathcal{P}$ is constructed such that $\forall Y \exists Z \varphi$ is true iff $(\mathcal{P}, a)$ is a yes-instance of $\mathit{Skept}^{CPAF}_{naive\text{-}glb,1}$.

- $A = X \cup \{\overline{y} \mid y \in Y\} \cup \{a, z_{m+1}\}$,
- $R = \{(y, \overline{y}), (\overline{y}, y) \mid y \in Y\} \cup \{(z_i, z_j) \mid 1 \leq i < j \leq m+1\} \cup \{(z_i, a) \mid 1 \leq i \leq m+1\}$,
- and *cond* defined by the following rules:
  - $\varphi \wedge z_{m+1} \Rightarrow z_j \succ z_i$ for all $1 \leq i < j \leq m+1$,
  - $\varphi \wedge z_{m+1} \Rightarrow a \succ z_i$ for all $1 \leq i \leq m+1$,
  - $z_i \wedge \bigwedge_{1 \leq j \leq m+1, j \neq i}(\neg z_j) \Rightarrow a \succ z_i$ for all $1 \leq i \leq m$.

  Expressed in natural language, the first two rules remove all conflicts between $z_1, \dots, z_{m+1}, a$ if $\varphi \wedge z_{m+1}$ is satisfied, and the third rule removes the conflict between some $z \in Z$ and a if this z is the only element from Z that is part of the extension, and if z_{m+1} is also not part of the extension.

The above construction is visualized in Figure 3.7. Note that the resulting CPAF is polynomial in the size of φ as we employ $O(m^2)$ rules, each linear in the size of φ. It remains to show that $\forall Y \exists Z \varphi$ is true iff $(\mathcal{P}, a)$ is a yes-instance of $\mathit{Skept}^{CPAF}_{naive\text{-}glb,1}$.

- Assume that $\forall Y \exists Z \varphi$ is true. We want to show that for all $S \in \mathit{naive\text{-}glb}^1_{cp}(\mathcal{P})$ we have $a \in S$. Towards a contradiction assume this is not the case, i.e., there is some $S \in \mathit{naive\text{-}glb}^1_{cp}(\mathcal{P})$ such that $a \notin S$. There are two possibilities:
  1. $S \models \varphi$. Then for $S' = S \cup \{a, z_{m+1}\}$ we also have $S' \models \varphi$ since a and z_{m+1} are fresh variables. Moreover, $S' \in \mathit{cf}^1_{cp}(\mathcal{P})$ since $S' \models \varphi \wedge z_{m+1}$ and thus all conflicts between the arguments $z_1, \dots, z_{m+1}, a$ are removed. But $S \subset S'$, i.e., $S \notin \mathit{naive\text{-}glb}^1_{cp}(\mathcal{P})$. Contradiction.
  2. $S \not\models \varphi$. Then $S \not\models \varphi \wedge z_{m+1}$ and therefore the conflicts between $z_1, \dots, z_{m+1}, a$ are not removed. This means that at most one of $z_1, \dots, z_{m+1}, a$ is in S since

we require $S \in cf^1_{cp}(\mathcal{P})$. Indeed, *exactly* one argument from $z_1, \ldots, z_{m+1}, a$ has to be in S, since if none of them were in S then we could add any of these arguments to S and the resulting set would still be conflict free regardless of preferences. By our assumption, $a \notin S$. Again, we distinguish two cases:

a) $z_i \in S$ with $1 \leq i \leq m$. But then $S \cup \{a\} \in cf^1_{cp}(\mathcal{P})$ because the following rule would apply: $z_i \wedge \bigwedge_{1 \leq j \leq m+1, j \neq i}(\neg z_j) \Rightarrow a \succ z_i$. Thus, $S \notin naive\text{-}glb^1_{cp}(\mathcal{P})$. Contradiction.

b) $z_{m+1} \in S$. Let $I_Y = Y \cap S$. Since $\forall Y \exists Z \varphi$ is true there is some $I_Z \subseteq Z$ such that $I_Y \cup I_Z \models \varphi$. Therefore, $S' \models \varphi$ for $S' = I_Y \cup \{\overline{y} \mid y \notin I_Y\} \cup I_Z \cup \{a, z_{m+1}\}$. Moreover, $S' \in cf^1_{cp}(\mathcal{P})$ since $S' \models \varphi \wedge z_{m+1}$ and thus all conflicts between $z_1, \ldots, z_{m+1}, a$ are removed. Since $S \subset S'$ by construction we have that $S \notin naive\text{-}glb^1_{cp}(\mathcal{P})$. Contradiction.

In all cases we arrive at a contradiction, and we can conclude that $a \in S$ for all $S \in naive\text{-}glb^1_{cp}(\mathcal{P})$.

- Assume that $\forall Y \exists Z \varphi$ is not true. Then there is some $I_Y \subseteq Y$ such that $I_Y \cup I_Z \not\models \varphi$ for all $I_Z \subseteq Z$. Let $S = I_Y \cup \{\overline{y} \mid y \notin I_Y\} \cup \{z_{m+1}\}$. Clearly, $S \in cf^1_{cp}(\mathcal{P})$. Moreover, there can be no $S' \supset S$ such that $S' \in cf^1_{cp}(\mathcal{P})$ since we would need to add at least one argument from $z_1, \ldots, z_m, a$ to S. But these arguments are all in conflict with z_{m+1} unless $S' \models \varphi \wedge z_{m+1}$, which we know to be impossible. □

We now turn our attention to admissibility-based semantics, where, in contrast to semantics based only on conflict-freeness, the choice of preference reduction makes no difference with regards to complexity. Again, let us first consider the cases in which there is no rise in complexity compared to AFs.

Lemma 3.30. *$Cred^{CPAF}_{\sigma,i}$ is* NP*-complete for $\sigma \in \{adm, com, stb, prf\text{-}glb\}$ and $i \in \{1, 2, 3, 4\}$.*

Proof. Hardness follows from hardness for AFs. Regarding membership of $\sigma \in \{adm, com, stb\}$, given a CPAF $\mathcal{P}$ and an argument x we can simply guess a set of arguments S containing x and, by Lemma 3.24, check whether $S \in \sigma^i_{cp}(\mathcal{P})$ in polynomial time. Regarding membership of *prf-glb*, it suffices to test whether $(\mathcal{P}, x)$ is a yes-instance of $Cred^{CPAF}_{adm,i}$. □

Lemma 3.31. *Let $i \in \{1, 2, 3, 4\}$. $Skept^{CPAF}_{\sigma,i}$ is trivial for $\sigma = adm$,* coNP*-complete for $\sigma = stb$, and Π^{P}_2-complete for $\sigma \in \{prf, prf\text{-}glb\}$.*

Proof. Hardness follows from hardness for AFs. Regarding membership, let $\mathcal{P} = (A, R, cond)$ be a CPAF and $x \in A$. Concerning $\sigma = adm$, note that $\emptyset$ is always admissible in $\mathcal{R}^{\emptyset}_i(\mathcal{P})$, i.e., $(\mathcal{P}, x)$ is trivially a no-instance. Regarding $\sigma \in \{stb, prf, prf\text{-}glb\}$ we consider the complementary problem: guess a set $S \subset A$ and check that $x \notin S$

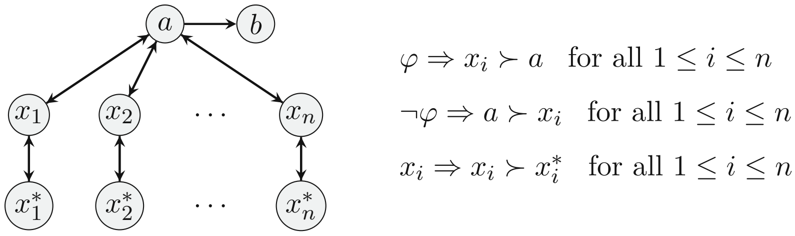

Figure 3.8: Construction used in the proof of Lemma 3.32. Given a formula φ over variables $X = \{x_1, x_2, \dots, x_n\}$, a CPAF $\mathcal{P}$ is constructed such that φ is satisfiable iff $(\mathcal{P}, b)$ is a yes-instance of $Cred^{CPAF}_{grd,i}$ iff $(\mathcal{P}, a)$ is a no-instance of $Skept^{CPAF}_{grd,i}/Skept^{CPAF}_{com,i}$.

and that $S \in \sigma^i_{cp}(\mathcal{P})$ with $\sigma \in \{stb, prf, prf\text{-}glb\}$. Checking $S \in \sigma^i_{cp}(\mathcal{P})$ can be done in polynomial time in the case of $\sigma = stb$ and in coNP-time in the case of $\sigma \in \{prf, prf\text{-}glb\}$ (cf. Lemma 3.24). □

In the case of grounded semantics, both credulous and skeptical acceptance are located one level higher on the polynomial hierarchy compared to AFs. For complete semantics, the same is true for skeptical acceptance.

Lemma 3.32. *Let $i \in \{1, 2, 3, 4\}$. $Cred^{CPAF}_{grd,i}$ is* NP*-complete. $Skept^{CPAF}_{\sigma,i}$ is* coNP*-complete for $\sigma \in \{grd, com\}$.*

Proof. We will consider the complementary problems of $Skept^{CPAF}_{grd,i}$/ $Skept^{CPAF}_{com,i}$ and show that they are NP-complete since this allows us to prove all results simultaneously.

NP-membership: given a CPAF $\mathcal{P} = (A, R, cond)$ and an argument $x \in A$, guess a set $S \subseteq A$ and, in polynomial time, check whether $S \in grd^i_{cp}(\mathcal{P})$ and $x \in S$ (resp. $S \in grd^i_{cp}(\mathcal{P})$ and $x \notin S$ or $S \in com^i_{cp}(\mathcal{P})$ and $x \notin S$).

NP-hardness by reduction from SAT: let φ be an arbitrary propositional formula over variables X. Let a and b be fresh variables, i.e., $a, b \notin X$. We construct an instance $(\mathcal{P}, b)$ of $Cred^{CPAF}_{grd,i}$ as follows: $\mathcal{P} = (A, R, cond)$ with $A = X \cup \{x^* \mid x \in X\} \cup \{a, b\}$, $R = \{(x, a), (a, x), (x, x^*), (x^*, x) \mid x \in X\} \cup \{(a, b)\}$, and $cond$ defined by the rules $\varphi \Rightarrow x \succ a$, $\neg\varphi \Rightarrow a \succ x$, and $x \Rightarrow x \succ x^*$ for $x \in X$, i.e., $x \succ_S a$ iff $S \models \varphi$, $a \succ_S x$ iff $S \models \neg\varphi$, and $x \succ_S x^*$ iff $S \models x$. Figure 3.8 depicts the above construction. In fact, this construction also works for the complementary problem of skeptical acceptance with respect to grounded and complete semantics, except that we will ask for the acceptance of the argument a instead of b. In this spirit, we now show that φ is satisfiable iff $(\mathcal{P}, b)$ is a yes-instance of $Cred^{CPAF}_{grd,i}$ iff $(\mathcal{P}, a)$ is a no-instance of $Skept^{CPAF}_{grd,i}/Skept^{CPAF}_{com,i}$.

- Suppose φ is satisfiable. Then there is an interpretation $I \subseteq X$ such that $I \models \varphi$. We assume that $I \neq \emptyset$. This is permissible since we can check in polynomial time whether $\emptyset$ satisfies φ, and if this is the case, return a trivial yes-instance of $Cred^{CPAF}_{grd,i}$ (or a trivial no-instance of $Skept^{CPAF}_{grd,i}/Skept^{CPAF}_{com,i}$). Consider $S = I \cup \{b\}$.

Then $S \models \varphi$ since $I \models \varphi$ and b does not occur in φ. We then have $x \succ_S a$ and $x \succ x^*$ for all $x \in I$, but $x \not\succ_S x^*$ for $x \in X \setminus I$. Thus, the unattacked arguments in $\mathcal{R}_i^S(\mathcal{P})$ are exactly those in I. Since $I \neq \emptyset$, b is defended in $\mathcal{R}_i^S(\mathcal{P})$ against a by the arguments in I. Thus, S is the minimal complete extension in $\mathcal{R}_i^S(\mathcal{P})$ and therefore also grounded in $\mathcal{R}_i^S(\mathcal{P})$. This implies that $(\mathcal{P}, b)$ is a yes-instance of $\mathit{Cred}_{grd,i}^{CPAF}$ while $(\mathcal{P}, a)$ is a no-instance of both $\mathit{Skept}_{grd,i}^{CPAF}$ and $\mathit{Skept}_{com,i}^{CPAF}$.

- Suppose φ is unsatisfiable. Then, for every $S \subseteq A$ and $x \in X$, we have $a \succ_S x$. Thus, the argument a is unattacked in every $\mathcal{R}_1^S(\mathcal{P})$, i.e., every complete extension (and therefore also every grounded extension) in $\mathcal{P}$ must contain a. Since a and b are in conflict, b is contained in no complete or grounded extension. Thus, $(\mathcal{P}, b)$ is a no-instance of $\mathit{Cred}_{grd,i}^{CPAF}$ while $(\mathcal{P}, a)$ is a yes-instance of both $\mathit{Skept}_{grd,i}^{CPAF}$ and $\mathit{Skept}_{com,i}^{CPAF}$. □

Lastly, we consider credulous acceptance for preferred semantics with local maximization. The following proof is the only one in this section to utilize one of the well-known standard reductions for AFs (Dvořák and Dunne 2018). Only a very limited inclusion of conditional preferences is necessary. Indeed, only a single preference rule, consisting of a very simple propositional formula, is used in the construction.

Lemma 3.33. *$\mathit{Cred}_{prf,i}^{CPAF}$ is Σ_2^{P}-complete for $i \in \{1, 2, 3, 4\}$.*

Proof. Σ_2^{P}-membership: given a CPAF $\mathcal{P} = (A, R, \mathit{cond})$ and an argument x, guess a set $S \subseteq A$ and check that $x \in S$ and, in coNP-time, that $S \in \mathit{prf}_{cp}^i(\mathcal{P})$.

Π_2^{P}-hardness of the complementary problem: let $\forall X \exists Y \varphi$ be an arbitrary instance of $\text{QBF}_{\forall}^2$ in 3-CNF over variables $Y = \{y_1, \ldots, y_n\}$ and $Z = \{z_1, \ldots, z_m\}$ with clauses $C = \{c_1, \ldots, c_k\}$. Let $X = Y \cup Z$. Using fresh variables a and b we construct an instance $(\mathcal{P}, a)$ of *co-*$\mathit{Cred}_{prf,i}^{CPAF}$ where $\mathcal{P} = (A, R, \mathit{cond})$,

- $A = X \cup \{\overline{x} \mid x \in X\} \cup C \cup \{a, b, \top, \bot\}$,
- $R = \{(x, \overline{x}), (\overline{x}, x) \mid x \in X\} \cup$
  $\{(x, c) \mid x \in C\} \cup \{(\overline{x}, c) \mid \neg x \in C\} \cup$
  $\{(c, c), (c, \top) \mid c \in C\} \cup$
  $\{(\bot, z), (\bot, \overline{z}) \mid z \in Z\} \cup$
  $\{(\top, \bot), (\bot, \bot)\} \cup \{(a, b), (b, a)\}$,
- and cond defined by the following rule: $\bigvee_{z \in Z}(z \vee \overline{z}) \Rightarrow b \succ a$.

This construction is exemplified in Figure 3.9. It remains to show that $\forall Y \exists Z \varphi$ is true iff $a \notin S$ for all $S \in \mathit{prf}_{cp}^i(\mathcal{P})$.

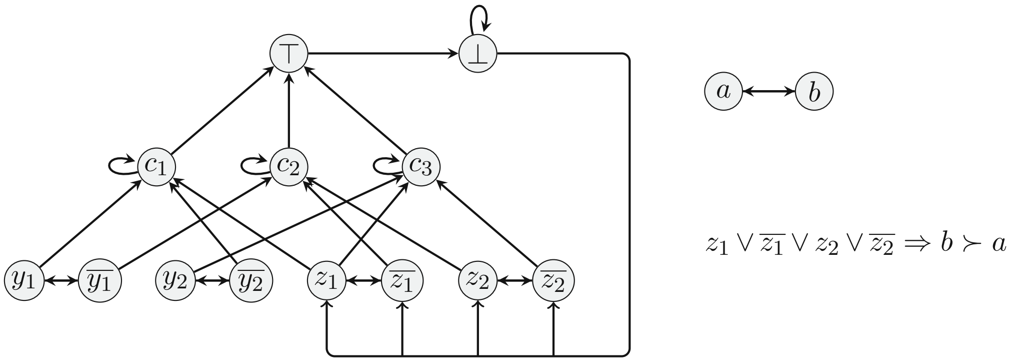

Figure 3.9: Construction used in the proof of Lemma 3.33. Given the quantified boolean formula $\forall y_1 y_2 \exists z_1 z_2\, \varphi$, with φ consisting of clauses $c_1 = (y_1 \vee \neg y_2 \vee z_1)$, $c_2 = (\neg y_1 \vee \neg z_1 \vee z_2)$, and $c_3 = (y_2 \vee z_1 \vee \neg z_2)$, a CPAF $\mathcal{P}$ is constructed such that $\forall y_1 y_2 \exists z_1 z_2\, \varphi$ is true iff $(\mathcal{P}, a)$ is a no-instance of $\mathit{Cred}^{\mathit{CPAF}}_{\mathit{prf},i}$.

- Assume that $\forall Y \exists Z \varphi$ is true. Towards a contradiction, assume that there is some $S \in \mathit{prf}^i_{cp}(\mathcal{P})$ such that $a \in S$. Then it must be that $b \not\succ_S a$, otherwise a is undefended in $\mathcal{R}^S_i(\mathcal{P})$. Thus, for all $z \in Z$ we have $z \notin S$ and $\overline{z} \notin S$. Let $I_Y = S \cap Y$. Since $\forall Y \exists Z \varphi$ is true there is some $I_Z \subseteq Z$ such that $I_Y \cup I_Z \models \varphi$. Let $S' = I_Y \cup \{\overline{y} \mid y \in Y, y \notin I_Y\} \cup I_Z \cup \{\overline{z} \mid z \in Z, z \notin I_Z\} \cup \{\top, a\}$. Clearly, $S \subset S'$. Moreover, S' is admissible in $\mathcal{R}^S_i(\mathcal{P})$: since $I_Y \cup I_Z \models \varphi$ all clause-arguments $c \in C$ are attacked by arguments in S', and therefore $\top$ is defended by S'. This further implies that all arguments $z, \overline{z}$ are defended by S' against $\bot$. We can conclude that S is not preferred in $\mathcal{R}^S_i(\mathcal{P})$. Contradiction.
- Assume that $\forall Y \exists Z \varphi$ is not true. Then there is $I_Y \subseteq Y$ such that $I_Y \cup I_Z \not\models \varphi$ for all $I_Z \subseteq Z$. Let $S = I_Y \cup \{\overline{y} \mid y \in Y, y \notin I_Y\} \cup \{a\}$. Note that $b \not\succ_S a$ and that all arguments $y, \overline{y}$ defend themselves. Thus, S is admissible in $\mathcal{R}^S_i(\mathcal{P})$. Towards a contradiction, assume there is $S' \supset S$ such that S' is admissible in $\mathcal{R}^S_i(\mathcal{P})$. This means one of the following must be the case:
  - $\top \in S'$. Then $\top$ needs to be defended by S' against the clause arguments $c \in C$. But this means that $I = (S' \cap Y) \cup (S' \cap Z)$ satisfies all clauses in φ, i.e., $I \models \varphi$. Note that S contains exactly one of $y, \overline{y}$ for every $y \in Y$. Thus, $S' \cap Y = S \cap Y = I_Y$. The fact that $I_Y \cup (S' \cap Z) \models \varphi$ contradicts $I_Y \cup I_Z \not\models \varphi$ for all $I_Z \subseteq Z$.
  - $z \in S'$ for some $z \in Z$. Then z needs to be defended by S' against $\bot$. This is only possible if $\top \in S'$, which we already have shown to not be the case. Contradiction.
  - $\overline{z} \in S'$ for some $z \in Z$. Analogous to the case that $z \in S'$.

Since we arrive at a contradiction in all cases, $S \in \mathit{prf}^i_{cp}(\mathcal{P})$. Moreover, note that $a \in S$. □

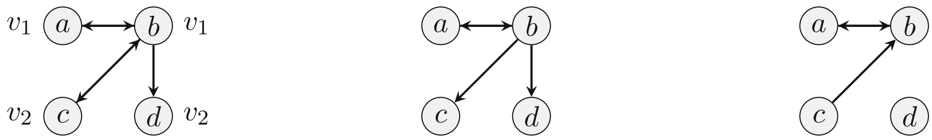

(a) Example VAF (b) AF for p_1 (c) AF for p_2

Figure 3.10: Example VAF with two audiences p_1 ($v_1 \succ v_2$) and p_2 ($v_2 \succ v_1$).

3.4 Related Formalisms

We now investigate the connection between CPAFs and related formalisms. First, we show that Value-based Argumentation Frameworks (VAFs) (Atkinson and Bench-Capon 2021; Bench-Capon, Doutre, and Dunne 2007) can be captured by CPAFs in a straightforward way. Secondly, we consider Extended Argumentation Frameworks (EAFs) (Baroni et al. 2009; Dunne, Modgil, and Bench-Capon 2010; Modgil 2009) and highlight similarities and differences to CPAFs. Lastly, we compare our CPAFs with a recently introduced alternative approach to conditional preferences in abstract argumentation (Alfano et al. 2023).

3.4.1 Capturing Value-Based Argumentation

VAFs, similarly to CPAFs, are capable of dealing with multiple preference relations. But, in contrast to CPAFs, these preferences are not over individual arguments but over values associated with arguments. Which values are preferred depends on the audience. A set of arguments may then be accepted in view of one audience, but not in view of another.

More formally, a VAF is a quintuple (A, R, V, val, P) such that (A, R) is an AF, V is a set of values, $val\colon A \to V$ is a mapping from arguments to values, and P is a finite set of audiences. Each audience $p \in P$ is associated with a preference relation $\succ_p$ over values, and $F_p = (A, R, V, val, \succ_p)$ is called an audience-specific VAF (AVAF). The extensions of VAFs are determined for each audience separately. Specifically, an argument x successfully attacks y in F_p iff $(x, y) \in R$ and $val(y) \not\succ_p val(x)$. Conflict-freeness and admissibility are then defined over these successful attacks. In essence, this boils down to using Reduction 1 on F_p, i.e., deleting attacks that contradict the preference ordering.

Figure 3.10a shows a VAF with two values v_1 and v_2. Let us say there are two audiences in this VAF, p_1 with the preference $v_1 \succ_{p_1} v_2$ and p_2 with $v_2 \succ_{p_2} v_1$. The AFs associated with p_1 and p_2, i.e., the AFs containing only the successful attacks in the AVAFs of p_1 and p_2, are depicted in Figures 3.10b and 3.10c.

The reasoning tasks typically associated with VAFs are those of subjective and objective acceptance. Let $F = (A, R, V, val, P)$ be a VAF and $x \in A$. Then x is subjectively accepted in F iff there is $p \in P$ such that x is in a preferred extension of the AVAF $(A, R, V, val, \succ_p)$. Similarly, x is objectively accepted in F iff for all $p \in P$ we have that x is in all preferred extensions of the AVAF $(A, R, V, val, \succ_p)$.

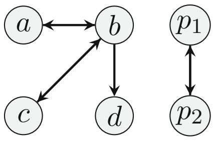

cond (let $x \in \{a, b\}, y \in \{c, d\}$):
$x \succ_S y$ iff $S \cap P = \{p_1\}$
$y \succ_S x$ iff $S \cap P = \{p_2\}$

Figure 3.11: CPAF obtained by translating the VAF of Figure 3.10 according to Definition 3.34.

We now provide a translation where an arbitrary VAF $F = (A, R, V, val, P)$ is transformed into a CPAF $Tr(F) = (A', R', cond)$ such that the subjectively and objectively accepted arguments in F correspond to the credulously and skeptically preferred arguments in $Tr(F)$ respectively.

Definition 3.34. *Let $F = (A, R, V, val, P)$ be a VAF. Then $Tr(F) = (A', R', cond)$ is the CPAF such that*

- $A' = A \cup P$,
- $R' = R \cup \{(p, p'), (p', p) \mid p, p' \in P, p \neq p'\}$,
- *for every $S \subseteq A'$, $a \succ_S b$ iff there is $p \in P$ with $S \cap P = \{p\}$ and $val(a) \succ_p val(b)$.*

Intuitively, each audience in the initial VAF is added as its own argument in our CPAF. The attacks of the VAF are preserved and symmetric attacks are added between all audience-arguments. Lastly, the preferences in our CPAF correspond to the preferences of each audience and are controlled by the newly introduced audience-arguments. Figure 3.11 shows the CPAF that results if the above translation is applied to the VAF of Figure 3.10.

Observe that the successful attacks in some AVAF $F_p = (A, R, V, val, \succ_p)$ are also attacks in $\mathcal{R}_1^{S\cup\{p\}}(Tr(F))$, where $S \subseteq A$, and vice versa. This means that the admissible sets in the initial VAF F stand in direct relationship to the admissible sets in our constructed CPAF, as expressed be the following lemma.

Lemma 3.35. *Let $F = (A, R, V, val, P)$ be a VAF, $S \subseteq A$, and $p \in P$. Then S is admissible in the AVAF $F_p = (A, R, V, val, \succ_p)$ iff $S \cup \{p\} \in adm^1_{cp}(Tr(F))$.*

Furthermore, note that all audience-arguments in $Tr(F)$ attack each other, i.e., an admissible set in $Tr(F)$ contains at most one audience-argument. In fact, each audience-argument defends itself, and thus every preferred extension in $Tr(F)$ must contain exactly one audience-argument $p \in P$ if we appeal to the prf^1_{cp}-semantics. This allows us to conclude that the direct correspondence between admissible sets observed in Lemma 3.35 carries over to preferred extensions.

Proposition 3.36. *Given a VAF $F = (A, R, V, val, P)$, $x \in A$ is subjectively (resp. objectively) accepted in F iff x is credulously (resp. skeptically) preferred in $Tr(F)$ w.r.t. Reduction 1.*

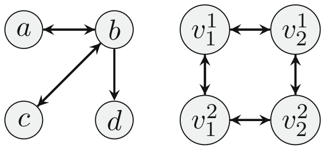

cond (let $x \in \{a, b\}, y \in \{c, d\}$):
$x \succ_S y$ iff $v_1^1 \in S, v_2^2 \in S$
$y \succ_S x$ iff $v_2^1 \in S, v_1^2 \in S$

Figure 3.12: CPAF obtained by translating the VAF of Figure 3.10 according to Definition 3.37.

It must be pointed out that the translation provided in Definition 3.34 was designed for VAFs in which each audience is given explicitly. However, VAFs can also be defined with preferences given implicitly as the set of all possible audiences (Atkinson and Bench-Capon 2021), where each audience corresponds to a linear ordering over all values. In this case, the translation of Definition 3.34 is not polynomial as the number of audience arguments would be factorial in the number of values. We now provide an alternative translation that can handle this implicit definition of audiences and where we only need $|V|^2$ additional arguments.

Definition 3.37. *Let $F = (A, R, V, val, P)$ be a VAF, with P implicitly given as the set of all possible linear orderings over V. Then $Tr_2(F) = (A', R', cond)$ is the CPAF such that*

- $A' = A \cup \{v^k \mid v \in V, 1 \leq k \leq |V|\}$,
- $R' = R \cup \{(v^k, w^k), (w^k, v^k) \mid v \in V, w \in V, 1 \leq k \leq |V|, v \neq w\} \cup \{(v^k, v^l) \mid v \in V, k \neq l\}$,
- *cond is defined as follows: for every $a, b \in A$ such that $val(a) \neq val(b)$ and every $1 \leq k < |V|$ we introduce the rule $val(a)^k \wedge (\bigvee_{l=k+1}^{|V|} val(b)^l) \Rightarrow a \succ b$.*

Figure 3.12 shows the CPAF that results if the above translation is applied to the VAF of Figure 3.10. As with our first translation (cf. Definition 3.34), there is a direct semantic correspondence between the initial VAF and the constructed CPAF. The idea is the following: along with the arguments and attacks of the original VAF, we introduce arguments $v^1, \dots, v^{|V|}$ for each value $v \in V$. If v^k is accepted, this means that v is considered the k-th best value. Since v^k attacks all other value-arguments w^k with $w \neq v$ we know that no other value is simultaneously ascribed the k-th best position. Moreover, v^k attacks all v^l with $l \neq k$, i.e., v is only ascribed the k-th best position and no other. Then, we prefer an argument a to another argument b if the value of a is preferred (appears at an earlier position in the linear ordering) than b. In this way, each extension corresponds to a linear ordering over all values, i.e., each extension corresponds to an audience. This further implies that each S-reduct of the constructed CPAF has exactly the same attacks between the arguments of the initial VAFs as the AVAF corresponding to the value-ordering encoded in S. This gives us a result analogous to Proposition 3.36.

Proposition 3.38. *Given a VAF $F = (A, R, V, val, P)$ where P is implicitly given as the set of all possible linear orderings over V, $x \in A$ is subjectively (resp. objectively) accepted in F iff x is credulously (resp. skeptically) preferred in $Tr_2(F)$ w.r.t. Reduction 1.*

Our translations highlight the versatility of our formalism. On the one hand, conditional preferences can be tied to dedicated arguments (in this case the audience-arguments). On the other hand, these dedicated arguments themselves may be part of the argumentation process. Note that we used CPAFs with Reduction 1 since preferences in VAFs are usually handled by deleting attacks. However, our approach also allows for the use of other preference-reductions in VAF-settings.

Moreover, note that the problem of subjective acceptance in VAFs is NP-complete (Bench-Capon, Doutre, and Dunne 2007), even if the set of all audiences is represented implicitly. In contrast, we have shown that credulous acceptance in CPAFs is Σ_2^{P}-complete (cf. Table 3.2). Thus, assuming that the polynomial hierarchy does not collapse, finding a polynomial translation from CPAFs to VAFs analogous to our Proposition 3.36 (resp. Proposition 3.38) is not possible when considering credulously/subjectively accepted arguments.

3.4.2 Relationship to Extended Argumentation Frameworks

EAFs allow arguments to express preferences between other arguments by permitting attacks themselves to be attacked. While EAFs are related to our CPAFs conceptually, we will see that there are crucial differences in how exactly preferences are handled.

Formally, an EAF is a triple (A, R, D) such that (A, R) is an AF, $D \subseteq A \times R$, and if $(a, (b, c)), (a', (c, b)) \in D$ then $(a, a'), (a', a) \in R$. The definition of admissibility in EAFs is quite involved and requires so-called reinstatement sets. Essentially, a set of arguments S is admissible in an EAF if all arguments $x \in S$ are defended from other arguments $y \in A \setminus S$, and if all attacks (z, y) used for defending x are in turn defended from attacks on attacks $(w, (z, y))$ and thus reinstated. It is possible that a chain of such reinstatements is required which is formalized with the aforementioned reinstatement sets. Formally defining these concepts is not necessary for our purposes, but the corresponding definitions can be found in (Modgil 2009). Observe that the notion of attacks on attacks in EAFs is similar to Reduction 1 in the sense that attacks between arguments can be unsuccessful, but they are never reversed. Therefore, we will compare EAFs to CPAFs with Reduction 1.

Recall our Sherlock Holmes example from above (Example 3.1) that we modeled as a CPAF (Example 3.4). Let us first consider a slimmed-down variation without an argument stating that Person 1 has an alibi. We can model this as an EAF with three arguments c_1 (Person 1 is the culprit), c_2 (Person 2 is the culprit), and m_1 (Person 1 has a motive) in which m_1 attacks the attack from c_2 to c_1. The corresponding EAF is depicted in Figure 3.13b. Compare this to the formalization via a CPAF in Figure 3.13a. Note that $\{c_1\}$ is admissible in the EAF but $\{c_2\}$ is not since (c_2, c_1) is used to defend

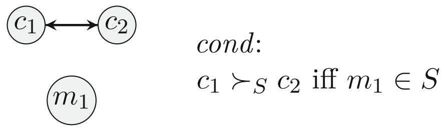

(a) Simple CPAF

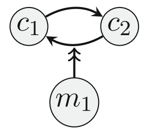

(b) Simple EAF

Figure 3.13: A simplified version of the Sherlock Holmes example modeled via a CPAF and an EAF.

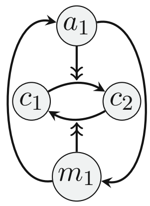

(a) Conflicting preferences

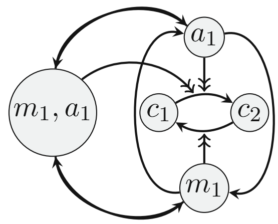

(b) Combination of preferences

Figure 3.14: The full Sherlock Holmes example modeled by two different EAFs.

against (c_1, c_2) but not reinstated against $(m_1, (c_2, c_1))$. In the CPAF, $\{c_2\}$ is admissible (but not stable).

This simple example highlights a fundamental difference in how preferences are viewed in the two formalisms. In CPAFs, preferences are relevant exactly if the argument that expresses them (e.g. m_1) are part of the set under inspection. In EAFs, preference are relevant even if the argument that expresses them is not accepted. Modgil (2009) states that admissibility for EAFs was defined in this way because it was deemed important to satisfy Dung's fundamental lemma (Dung 1995), which says that if S is admissible and x is acceptable w.r.t. S then $S \cup \{x\}$ is admissible. This fundamental lemma is not satisfied in our CPAFs (see Proposition 3.11). However, in our opinion, this is no drawback but rather a necessary property of formalisms that can deal with conditional preferences in a flexible way. For example, in Figure 3.13a it is clear that $\{c_2\}$ should be admissible since, when considering only admissibility, we are not forced to include the unattacked m_1, i.e., we do not have to accept that Person 1 has a motive. The inclusion of unattacked arguments in CPAFs is handled via more restrictive approaches such as stable or preferred semantics, as usual.

Another difference between CPAFs and EAFs becomes clear when considering the entire Sherlock Holmes example. Recall our formalization for CPAFs (cf. Figure 3.1). In order to express our preference in case Person 1 has an alibi we extend our EAF from Figure 3.13b by adding an attack from a_1 to the attack (c_1, c_2), as shown in Figure 3.14a[7]. Note that a_1 and m_1 must attack each other in this EAF by definition since they express

[7]The EAFs of Figure 3.13b and Figure 3.14a are also used as examples in (Modgil 2009).

conflicting preferences. But this formalization is unsatisfactory since it should be possible for Person 1 to have both a motive and an alibi. The fact that the preference of one argument may change in view of another argument must be modeled indirectly in EAFs. For example, we can introduce an additional argument to express that Person 1 has both a motive and an alibi. This is depicted in Figure 3.14b. Thus, we can see that CPAFs allow for more flexibility when combining preferences associated with several arguments.

There are also some differences between CPAFs and EAFs when it comes to preferred semantics. For instance, stable extensions in EAFs are not necessarily preferred extensions (Dunne, Modgil, and Bench-Capon 2010). In CPAFs, every stable extension is also preferred, except if we use global maximization and Reduction 1 (cf. Proposition 3.10). Moreover, credulous acceptance under preferred semantics is in NP for EAFs (Dunne, Modgil, and Bench-Capon 2010), but Σ_2^{P}-complete for CPAFs when using local maximization (cf. Table 3.2).

To summarize, CPAFs are designed to express conditional preferences in abstract argumentation, whereas preferences in EAFs are unconditional in the sense that they may always influence the argumentation process, even if the argument associated with the preference is not accepted. Moreover, since CPAFs can make use of all four preference reductions, they allow for more flexibility in how preferences are handled compared to EAFs, where unsuccessful attacks are always deleted. However, the two formalisms are similar in that arguments are capable of reasoning about the argumentation process itself, i.e., they constitute a form of metalevel argumentation (Modgil and Bench-Capon 2011). Lastly, note that a generalization of EAFs has been suggested in which attacks on attacks can themselves be attacked, i.e., unlike in EAFs and CPAFs, preferences can be directly attacked themselves (Baroni et al. 2011).

3.4.3 Lifting Preferences over Arguments to Sets of Arguments

In our CPAFs, we deal with preferences by using preference reductions which modify the attack relation (see Definition 2.10). There exist other approaches to preferences in argumentation, where preference orderings over arguments are lifted to sets of arguments (Alfano et al. 2022; Amgoud and Vesic 2014; Brewka, Truszczynski, and Woltran 2010; Kaci, van der Torre, and Villata 2018), and the most preferred extensions are then selected according to this new preference ordering.

Recently in (Alfano et al. 2023), conditional preferences in abstract argumentation have been investigated using the aforementioned preference liftings. We refer to the CPAFs introduced in that work as *lifting-based* CPAFs. Similarly to our reduction-based CPAFs, a lifting-based CPAF is given as (A, R, Γ) where (A, R) is an AF and Γ is a set of conditional preference rules of the form $a_1 \succ a_2 \leftarrow b_1 \wedge \cdots \wedge b_m \wedge \neg c_1 \wedge \cdots \wedge \neg c_n$ built from arguments $a_1, a_2, b_1, \ldots, b_m, c_1, \ldots, c_n$. The conditional preferences over arguments given by Γ are lifted to preferences over sets of arguments according to one of three criteria (democratic, elitist, KTV), and then the 'best' extensions are selected according to this lifted preference ordering.

Note that lifting-based CPAFs, in contrast to our reduction-based CPAFs, satisfy principle $P2^*$ (cf. Definition 3.13) by design, since the 'best' extensions selected in a lifting-based CPAF (A, R, Γ) are always extensions of (A, R). We note that, for complete and stable semantics, Reduction 3 satisfies $P2^*$ as well and thus selects extensions in the style of preference liftings (cf. Table 3.1).

The conditional preference rules Γ of a lifting-based CPAF (A, R, Γ) are usually assumed to be well-formed[8], which ensures that arguments a_1, a_2 occurring in the head of a rule $a_1 \succ a_2 \leftarrow b_1 \wedge \cdots \wedge b_m \wedge \neg c_1 \wedge \cdots \wedge \neg c_n$ do not occur in the body of the same rule. This is to prevent counterintuitive results, as explained in (Alfano et al. 2023) via the following example: given (A, R, Γ) with extensions $\{\{a, b\}, \{a, c\}\}$ and Γ given by $c \succ b \leftarrow b$ and $c \succ b \leftarrow c$, one would expect the only 'best' extension to be $\{a, c\}$. However, under semantics of lifting-based CPAFs, both $\{a, b\}$ and $\{a, c\}$ are 'best'. This problem does not occur with the well-formed $\Gamma' = \{c \succ b \leftarrow\}$. In our reduction-based CPAFs we have no analogous assumption of well-formedness. Despite this, the counter-intuitive behavior observed above does not necessarily occur in our reduction-based CPAFs. For example, consider (A, R, cond) with $A = \{a, b, c\}$, $R = \{(b, c), (c, b)\}$, and cond given by the rules $b \Rightarrow c \succ b$ and $c \Rightarrow c \succ b$. Then $\mathit{prf}((A, R)) = \{\{a, b\}, \{a, c\}\}$. However, under all four preference reductions, the attack (b, c) is deleted as soon as b or c is in the extension under inspection. Thus, $\mathit{prf}^i_{cp}((A, R, \mathit{cond})) = \{\{a, c\}\}$.

Another difference between lifting-based CPAFs and our reduction-based CPAFs lies in their computational complexity, which is higher for lifting-based CPAFs in most cases. For example, verification for stable semantics is coNP-complete in lifting-based CPAFs (Alfano et al. 2023) but remains in P in reduction-based CPAFs (see Table 3.2). As a result, credulous and skeptical acceptance for stable semantics are Σ_2^{P}-complete and Π_2^{P}-complete respectively in lifting-based CPAFs, while they remain NP-complete and coNP-complete respectively in reduction-based CPAFs. Some problems, such as credulous and skeptical acceptance of preferred semantics under elitist and KTV criteria, may even lie on the third level of the polynomial hierarchy for lifting-based CPAFs (tight bounds for the complexity of these problems have not been established yet). We observe that the increased complexity of lifting-based CPAFs is in many cases not due to the introduction of conditional preferences, but rather due to the preference-liftings themselves, as the complexity of lifting-based PAFs (featuring only unconditional preferences) is already considerably higher than that of standard Dung-style AFs (Alfano et al. 2022).

Moreover, there are lifting-based CPAFs where not every (best) stable extension is also a (best) preferred extension (Alfano et al. 2023, Example 2). In contrast, every stable extension in a reduction-based CPAF is also a preferred extension, except when considering Reduction 1 and preferred semantics with global maximization (cf. Proposition 3.10).

Lastly, we want to emphasize that there is a conceptual difference between the reduction-based and lifting-based approaches to resolving preferences in argumentation: when using preference reductions, $x \succ y$ expresses that x is stronger than y; when using preference

[8]Note that this notion of well-formedness is unrelated to the well-formed CAFs used in Chapter 4.

liftings, $x \succ y$ expresses that we prefer outcomes containing x rather than y. Which of the two approaches should be chosen depends on the task at hand. In fact, it is also possible to combine the reduction-based and lifting-based approaches to preferences in a single formalism. This has already been suggested by Amgoud and Vesic (2014) in the form of their so-called Rich PAFs.

3.5 Conclusion

In this chapter, we introduced Conditional Preference-based AFs (CPAFs) which generalize PAFs and allow to flexibly handle conditional preferences in abstract argumentation.

We conducted a principle-based analysis for CPAFs and showed that complete and stable semantics satisfy the same principles as on PAFs in most cases while grounded semantics no longer satisfies many of the principles. We further investigated the computational complexity of CPAFs and showed that this complexity can be influenced by the chosen preference reduction (in case of naive semantics) or by how maximization is handled (in case of naive and preferred semantics). Our results also show that the satisfaction of I-maximality can depend on how maximization is dealt with (in case of preferred semantics) and on which preference-reduction is chosen (in case of stable semantics).

Moreover, we compared CPAFs to related formalisms. On the one hand, we showed that CPAFs can be used to capture VAFs via a straightforward translation. On the other hand, we demonstrated that CPAFs exhibit significant differences to EAFs in terms of how preferences are handled. We also discussed a recently introduced alternative approach to conditional preferences in abstract argumentation, where preferences over arguments are lifted to preferences over sets of arguments.

Regarding future work, in addition to the ten principles from the literature studied in this chapter, new principles that allow us to further examine conditional preferences in argumentation may be formalized. Another avenue for future work is to explore possible restrictions to the structure of CPAFs, be it to the underlying AF or the conditional preferences themselves, and how such restrictions impact the formal properties of CPAFs. Moreover, the relationship between CPAFs and existing approaches in structured argumentation (Dung, Thang, and Son 2019) shall be investigated. Related to this point, it may also be interesting to see whether conditional preferences can be adapted to other formalisms such as bipolar argumentation frameworks (Amgoud et al. 2008), in which both attack and support relations are present. As for preference representation, it could be investigated how existing formalisms designed to handle conditional preferences such as CP-nets (Boutilier et al. 2004) or various forms of logic programming (Brewka et al. 2015; Brewka, Niemelä, and Syrjänen 2004; Charalambidis, Rondogiannis, and Troumpoukis 2021; Delgrande, Schaub, and Tompits 2003) relate to CPAFs. Lastly, given the fact that CPAF semantics behave differently from AF semantics in many respects, it would be interesting to examine how argument justification can be explained in CPAFs with methods such as discussion games (Caminada 2018).

CHAPTER 4

Preferences in Claim-based Argumentation

Claim-centric argumentation, i.e., the evaluation of commonly acceptable statements while disregarding their particular justifications, has received considerable attention in the literature (Baroni and Riveret 2019; Dvořák and Woltran 2020; Horty 2002; Rocha and Cozman 2022a). A simple yet powerful generalization of AFs that allows for claim-based evaluation are Claim-augmented AFs (CAFs) (Dvořák and Woltran 2020), where each argument is assigned a claim (sometimes also referred to as conclusion). CAFs serve as an ideal target formalism for structured argumentation formalisms which utilize abstract argumentation semantics whilst also considering the claims of the arguments in the evaluation (Dung, Kowalski, and Toni 2009; Modgil and Prakken 2018) and therefore help to bridge the gap between abstract and structured argumentation.

Although the acceptance of claims is closely related to argument acceptance, there are crucial differences as observed in (Dvořák and Woltran 2020; Modgil and Prakken 2018; Prakken and Vreeswijk 2002) stemming from the fact that a single claim can be associated with several different arguments. As a consequence, several properties of AF semantics cannot be taken for granted when considered in terms of the arguments' claims. For instance, I-maximality, which gives insights into the expressiveness of semantics (Dunne et al. 2015) and skeptical argument justification (Baroni and Giacomin 2007) is not satisfied by most CAF semantics (Dvořák, Rapberger, and Woltran 2023). Furthermore, the introduction of claims causes a rise in the computational complexity of some standard decision problems in argumentation (Dvořák et al. 2023). Luckily, these drawbacks can be alleviated by taking fundamental properties of the attack relation into account: the basic observation that attacks typically depend on the claim of the attacking arguments gives rise to the central class of *well-formed CAFs* (wfCAFs). This class satisfies that all arguments with the same claim attack the same arguments; thus modeling a very natural behavior of arguments that is common to structured argumentation formalisms and instantiations (Dung, Kowalski, and Toni 2009; Modgil and Prakken 2018). Well-formed CAFs have the main advantage that most of the semantics behave 'as expected', e.g., they retain I-maximality, and their computational complexity is located at the same level of the polynomial hierarchy as for AFs (Dvořák et al. 2023; Dvořák, Rapberger, and Woltran 2023).

Unfortunately, it turns out that well-formedness cannot be assumed if one deals with preferences in argumentation, as arguments with the same claim are not necessarily equally plausible. The following example demonstrates this.

Example 4.1. *Consider two arguments a, a' with claim α, and another argument b having claim β. Moreover, both a and a' attack b, while b attacks a. Furthermore assume that we are given the additional information that b is preferred over a' (for example, if assumptions in the support of b are stronger than assumptions made by a'). A common method to integrate such information on argument rankings is to* delete *attacks from arguments that attack preferred arguments. In this case, we delete the attack from a' to b.*

Both frameworks are depicted below: $\mathcal{F}$ represents the original situation while $\mathcal{F}'$ is the CAF resulting from deleting the unsuccessful attack from a' on the argument b.

Note that $\mathcal{F}$ is well-formed since all arguments with the same claims attack the same arguments. The unique acceptable argument-set w.r.t. stable semantics (cf. Definition 2.6) is $\{a, a'\}$ which translates to $\{\alpha\}$ on the claim-level.

The CAF $\mathcal{F}'$, on the other hand, is no longer well-formed since a' does not attack b. In F', the argument-sets $\{a, a'\}$ and $\{a', b\}$ are both acceptable w.r.t. to stable semantics. In terms of claims this translates to $\{\alpha\}$ and $\{\alpha, \beta\}$, which shows that I-maximality is violated on the claim-level.

Although well-formedness cannot be guaranteed in view of preferences, this does not imply arbitrary behavior of the resulting CAF: on the one hand, preferences conform to a certain type of ordering (e.g., asymmetric, transitive) over the set of arguments; on the other hand, it is evident that the deletion, reversion, and other types of attack manipulation impose restrictions on the structure of the resulting CAF. Combining both aspects, we obtain that, assuming well-formedness of the initial framework, it is unlikely that preference incorporation results in arbitrary behavior. The key motivation of this chapter is to identify and exploit structural properties of preferential argumentation in the scope of claim acceptance. The aforementioned restrictions suggest beneficial impact on both the computational complexity and on desired semantical properties such as I-maximality.

Contributions. In this chapter, we introduce Preference-based CAFs (PCAFs) in order to study the effect of preferences on wfCAFs. We make use of the four commonly used preference reductions from the literature (see Subsection 2.3.2) and study their effects under a claim-centric view. In particular, we investigate the following points:

- For each of the four reductions, we characterize the possible structure of CAFs that are obtained by applying the reduction to a wfCAF and a preference relation. This results in four novel CAF classes, each of which constitutes a proper extension of wfCAFs not retaining full expressiveness of general CAFs. We investigate the relationship between these classes.
- We study semantic properties of the novel CAF classes. Our results highlight a significant advantage of a particular reduction (namely Reduction 3) when it comes to admissibility-based semantics: under this modification, subset-maximization (as used in preferred semantics for example) on the argument-level coincides with subset-maximization on the claim-level. Moreover, this modification preserves I-maximality. The other reductions fail to preserve these properties in most cases; moreover, for the conflict-free-based naive and stage semantics, I-maximality cannot be guaranteed for any of the four reductions.
- We investigate the complexity of reasoning for CAFs with preferences. We show that for three of the four reductions (namely Reductions 2, 3, and 4), the verification problem drops by one level in the polynomial hierarchy for all except complete semantics and is thus not harder than for wfCAFs (which in turn has the same complexity as the corresponding AF problems). Complete semantics remain hard for all but one preference reduction. Moreover, it turns out that verification for the reduction which deletes attacks from weaker arguments (i.e., Reduction 1) remains as hard as for general CAFs.

Our results constitute a systematic study of the structural and computational effect of preferences on claim acceptance. Since we use CAFs as our base formalism, our investigations extend to large classes of formalisms that can be represented as CAFs, just like results on AFs yield insights for formalisms that can be captured by AFs.

Publications. This chapter is based on the paper (Bernreiter et al. 2023). The following contributions are new in this chapter: in addition to inherited CAF-semantics, we now also consider hybrid CAF-semantics (see Definition 2.23) and investigate them with respect to their semantic properties (in Section 4.3) and their computational complexity (in Section 4.4). Moreover, this version contains full proofs for our results, as well as additional figures and explanations.

Outline. This chapter is organized as follows. In Section 4.1, we introduce preference-based CAFs (PCAFs) which combine PAFs with wfCAFs. We characterize the novel CAF classes based on the preference reductions in Section 4.2, study the I-maximality of the semantics in Section 4.3, and their computational complexity in Section 4.4. We conclude in Section 4.5.

Required preliminaries. Before reading this chapter, it is recommended to read Section 2.1 (propositional logic), Section 2.2 (computational complexity), and especially Section 2.3 (formal argumentation).

4.1 Preference-based Claim-augmented AFs (PCAFs)

As discussed above and in Subsection 2.3.3, wfCAFs are a natural subclass of CAFs with advantageous semantic and computational properties. However, when resolving preferences among arguments, the resulting CAFs are typically no longer well-formed (cf. Example 4.1). In order to study preferences under a claim-centric view we introduce Preference-based CAFs. These frameworks enrich the notion of wfCAFs with the concept of preferences in terms of argument strength. Our main goals are then to understand the effect of resolved preferences on the structure of the underlying wfCAF, and to determine whether the advantages of wfCAFs are maintained. Given this motivation, it is reasonable to consider the impact of preferences on wfCAFs only.

Definition 4.2 (PCAF)**.** *A Preference-based Claim-augmented AF (PCAF) is a quadruple* $\mathcal{P} = (A, R, cl, \succ)$ *where* (A, R, cl) *is a wfCAF and* $(A, R, \succ)$ *is a PAF.*

Given a PCAF $\mathcal{P} = (A, R, cl, \succ)$ we sometimes write $a \in \mathcal{P}$ for $a \in A$ and $(a, b) \in \mathcal{P}$ for $(a, b) \in R$. Analogous notation will be used for CAFs and AFs.

Preferences in PCAFs are resolved via one of the four preference reductions, analogously to how they are resolved in PAFs (cf. Definition 2.10). Observe that all four reductions are polynomial time computable with respect to the input PCAF.

Definition 4.3 (Preference reductions applied to PCAFs)**.** *Let* $\mathcal{P} = (A, R, cl, \succ)$ *be a PCAF. The corresponding CAF* $\mathcal{R}_i(\mathcal{P}) = (A, R', cl)$ *is obtained by applying Reduction* i*, where* $i \in \{1, 2, 3, 4\}$*, to the underlying PAF* $P = (A, R, \succ)$ *of* $\mathcal{P}$*, i.e.,* $(A, R') = \mathcal{R}_i(P)$*.*

The semantics of PCAFs work by first resolving preferences between arguments, and then applying CAF-semantics to the resulting CAF.

Definition 4.4 (PCAF-semantics)**.** *Let* $\mathcal{P}$ *be a PCAF and let* $i \in \{1, 2, 3, 4\}$*. The PCAF-variant of a CAF-semantics* σ_μ *relative to Reduction* i *is defined as* $\sigma^i_\mu(\mathcal{P}) = \sigma_\mu(\mathcal{R}_i(\mathcal{P}))$*.*

Note that many structured argumentation formalisms use preference reductions. For instance, ABA+ (Cyras and Toni 2016) employs attack reversal similar to Reduction 2 while some instances of ASPIC (Modgil and Prakken 2013) delete attacks from weaker arguments in the spirit of Reduction 1.

Example 4.5. *Let $\mathcal{P} = (A, R, cl, \succ)$ be the PCAF with arguments $A = \{a, a', b\}$, attacks $R = \{(a, b), (a', b), (b, a)\}$, claims $cl(a) = cl(a') = \alpha$ and $cl(b) = \beta$, and the preference $b \succ a'$. The underlying CAF (A, R, cl) of $\mathcal{P}$ was examined in Example 2.21.*

Note that $\mathcal{R}_1(\mathcal{P}) = (A, R', cl)$ with $R' = \{(a, b), (b, a)\}$, which is the same CAF as $\mathcal{F}'$ in Example 4.1. It can be verified that, e.g., $adm^1_{inh}(\mathcal{P}) = adm_{inh}(\mathcal{R}_1(\mathcal{P})) = \{\{\emptyset, \{\alpha\}, \{\beta\}, \{\alpha, \beta\}\}$ and $stb^1_{inh}(\mathcal{P}) = \{\{\alpha\}, \{\alpha, \beta\}\}$.

As in PAFs (cf. Example 2.12) the choice of reduction can influence the extensions of a PCAF. For example, $\mathcal{R}_2(\mathcal{P}) = (A, R'', cl)$ with $R'' = \{(a, b), (b, a), (b, a)\}$, $adm^2_{inh}(\mathcal{P}) = \{\emptyset, \{\alpha\}, \{\beta\}\}$, and $stb^2_{inh}(\mathcal{P}) = \{\{\alpha\}, \{\beta\}\}$.

Remark. *In this chapter we require the underlying CAF of a PCAF to be well-formed. The reason for this is that we are interested in whether the benefits of well-formed CAFs are preserved when preferences have to be taken into account. Even from a technical perspective, admitting PCAFs with a non-well-formed underlying CAF is not very interesting with respect to the questions addressed in this chapter. Indeed, any CAF could be obtained from such general PCAFs, regardless of which preference reduction we are using, by simply specifying the desired CAF and an empty preference relation. Thus, such general PCAFs have the same properties regarding I-maximality and complexity as general CAFs.*

4.2 Syntactic Characterization & Expressiveness

Our first step towards understanding the effect of preferences on wfCAFs is to examine the impact of resolving preferences on the *structure* of the underlying CAF. To this end, we consider four new CAF classes which are obtained from applying the reductions of Definition 2.10 to PCAFs.

Definition 4.6 (CAF-classes)**.** $\mathcal{R}_i$-**CAF** *denotes the set of CAFs that can be obtained by applying Reduction i to PCAFs, i.e.,* $\mathcal{R}_i$-**CAF** $= \{\mathcal{R}_i(\mathcal{P}) \mid \mathcal{P}$ *is a PCAF*$\}$.

It is easy to see that $\mathcal{R}_i$-**CAF**, where $i \in \{1, 2, 3, 4\}$, contains all wfCAFs (we can simply specify the desired wfCAF and an empty preference relation). Moreover, not all CAFs are contained in $\mathcal{R}_i$-**CAF**, i.e., the four new classes are located in-between wfCAFs and general CAFs:

Proposition 4.7. *Let* **CAF** *be the set of all CAFs and* **wfCAF** *the set of all wfCAFs. For all $i \in \{1, 2, 3, 4\}$ it holds that* **wfCAF** $\subset \mathcal{R}_i$-**CAF** $\subset$ **CAF**.

Proof. Let $i \in \{1, 2, 3, 4\}$. **wfCAF** $\subseteq \mathcal{R}_i$-**CAF** follows from the fact that any $(A, R, cl) \in$ **wfCAF** can be obtained via Reduction i from the PCAF $(A, R, cl, \emptyset)$.

wfCAF $\subset \mathcal{R}_i$-**CAF**: consider the PCAF $\mathcal{P} = (\{a, b\}, \{(a, a), (a, b), (b, a), (b, b)\}, cl, \succ)$ with $cl(a) = cl(b)$, and $b \succ a$. For all $i \in \{1, 2, 3, 4\}$ we have $\mathcal{R}_i(\mathcal{P}) = (\{a, b\}, \{(a, a),$

$(b, a), (b, b)\}, cl)$, i.e., the resulting CAF $\mathcal{R}_i(\mathcal{P})$ is not well-formed since b is attacked by itself but not by a, even though $cl(a) = cl(b)$.

$\mathcal{R}_i$-**CAF** $\subset$ **CAF**: Towards a contradiction, assume there is a PCAF $\mathcal{P} = (A, R, cl, \succ)$ such that $\mathcal{R}_i(\mathcal{P}) = (A, R', cl)$ with $(a, b), (b, a) \in R'$ but $(a, a), (b, b) \notin R'$ for some $a, b \in A$ with $cl(a) = cl(b)$. This means that either $(a, b) \in R$ or $(b, a) \in R$, since none of four reductions can introduce the attacks (a, b) and (b, a) at the same time. By symmetry, we only look at the case that $(a, b) \in R$. Then, since (A, R, cl) is well-formed and since $cl(a) = cl(b)$, $(b, b) \in R$. But $\succ$ is non-reflexive, i.e., (b, b) is not removed by Reduction i and therefore $(b, b) \in R'$. Contradiction. □

Furthermore, the new classes are all distinct from each other, i.e., we are indeed dealing with *four* new CAF classes. Specifically, $\mathcal{R}_1$-**CAF**, $\mathcal{R}_2$-**CAF**, and $\mathcal{R}_4$-**CAF** are incomparable while $\mathcal{R}_3$-**CAF** is strictly contained in the other three classes. This reflects the fact that Reduction 3 is the most conservative of the four preference reductions, removing attacks from weak to strong arguments only when there is a counter-attack from the strong argument.

Proposition 4.8. *For all $i \in \{1, 2, 4\}$ and all $j \in \{1, 2, 3, 4\}$ such that $i \neq j$ it holds that $\mathcal{R}_i$-**CAF** $\not\subseteq \mathcal{R}_j$-**CAF** and $\mathcal{R}_3$-**CAF** $\subset \mathcal{R}_i$-**CAF**.*

Proof. We show the various statements separately.

- $\mathcal{R}_1$-**CAF** $\not\subseteq \mathcal{R}_j$-**CAF** with $j \in \{2, 3, 4\}$: let $\mathcal{F}$ be the CAF shown in Figure 4.1a. $\mathcal{F}$ is in $\mathcal{R}_1$-**CAF** as it can be obtained by applying Reduction 1 to the PCAF $(A, R, cl, \succ)$ with $R = \{(a, b), (b, b)\}$ and $b \succ a$. Towards a contradiction, assume there is a PCAF $\mathcal{P}$ such that $\mathcal{R}_j(\mathcal{P}) = \mathcal{F}$. Since self-attacks cannot be removed by any of the four reductions, $(b, b) \in \mathcal{P}$. Since the underlying CAF of $\mathcal{P}$ must be well-formed, also $(a, b) \in \mathcal{P}$. But then, by the definition of Reduction j, either $(a, b) \in \mathcal{R}_j(\mathcal{P})$ or $(b, a) \in \mathcal{R}_j(\mathcal{P})$. Contradiction.

- $\mathcal{R}_2$-**CAF** $\not\subseteq \mathcal{R}_j$-**CAF** with $j \in \{1, 3, 4\}$: let $\mathcal{F}$ be the CAF shown in Figure 4.1b. $\mathcal{F}$ is in $\mathcal{R}_2$-**CAF** as it can be obtained by applying Reduction 2 to the PCAF $(A, R, cl, \succ)$ with $R = \{(a, b), (b, b)\}$ and $b \succ a$. Towards a contradiction, assume there is a PCAF $\mathcal{P}$ such that $\mathcal{R}_j(\mathcal{P}) = \mathcal{F}$. Then $(b, b) \in \mathcal{P}$ and therefore also $(a, b) \in \mathcal{P}$. But $(b, a) \notin \mathcal{P}$, since $(a, a) \notin \mathcal{F}$ and therefore also $(a, a) \notin \mathcal{P}$. But Reductions 1 and 3 cannot introduce (b, a) in this case, while Reduction 4 cannot introduce (b, a) without retaining (a, b).

- $\mathcal{R}_4$-**CAF** $\not\subseteq \mathcal{R}_j$-**CAF** with $j \in \{1, 2, 3\}$: let $\mathcal{F}$ be the CAF shown in Figure 4.1c. $\mathcal{F}$ is in $\mathcal{R}_4$-**CAF** as it can be obtained by applying Reduction 4 to the PCAF $(A, R, cl, \succ)$ with $R = \{(a, b), (b, b)\}$ and $b \succ a$. Towards a contradiction, assume there is a PCAF $\mathcal{P}$ such that $\mathcal{R}_j(\mathcal{P}) = \mathcal{F}$. Then $(b, b) \in \mathcal{P}$ and therefore also $(a, b) \in \mathcal{P}$. But $(b, a) \notin \mathcal{P}$, since $(a, a) \notin \mathcal{P}$. But Reduction 1, 2 and 3 cannot introduce (b, a), at least not without deleting (a, b).

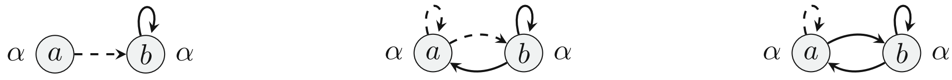

(a) only in $\mathcal{R}_1$-**CAF** (b) only in $\mathcal{R}_2$-**CAF** (c) only in $\mathcal{R}_4$-**CAF**

Figure 4.1: CAFs that are contained only in $\mathcal{R}_1$-**CAF**, $\mathcal{R}_2$-**CAF**, and $\mathcal{R}_4$-**CAF** respectively. Dashed arrows are attacks that are missing for the CAF to be well-formed.

- $\mathcal{R}_3$-**CAF** $\subset \mathcal{R}_j$-**CAF** with $j \in \{1,2,4\}$: let $\mathcal{F}$ be any CAF in $\mathcal{R}_3$-**CAF**. Then there is a PCAF $\mathcal{P} = (A, R', cl, \succ)$ such that $\mathcal{R}_3(\mathcal{P}) = \mathcal{F}$. If $(a,b) \in \mathcal{P}$ and $(a,b) \in \mathcal{F}$ we can assume that $b \not\succ a$ without loss of generality. If $(a,b) \in \mathcal{P}$ but $(a,b) \notin \mathcal{F}$, then, by definition of Reduction 3, $(b,a) \in \mathcal{P}$ and $b \succ a$. In this case, Reduction j functions in the same way as Reduction 3 (cf. Definition 2.10 and Figure 2.3), i.e., $\mathcal{R}_j(\mathcal{P}) = \mathcal{F}$. This proves $\mathcal{R}_3$-**CAF** $\subseteq \mathcal{R}_j$-**CAF**. $\mathcal{R}_3$-**CAF** $\subset \mathcal{R}_j$-**CAF** follows from $\mathcal{R}_j$-**CAF** $\not\subseteq \mathcal{R}_3$-**CAF**. □

We now know that applying preferences to wfCAFs results in four distinct CAF-classes that lie in-between wfCAFs and general CAFs. It is still unclear, however, how to determine whether some CAF belongs to one of these classes or not. Especially for $\mathcal{R}_2$-**CAF** and $\mathcal{R}_4$-**CAF** this is not straightforward, since Reductions 2 and 4 not only remove but also introduce attacks and therefore allow for several possibilities by which a particular CAF can be obtained. We tackle this problem by characterizing the new classes via the so-called wf-problematic part of a CAF.

Definition 4.9 (wf-problematic part)**.** *A pair of arguments (a,b) is wf-problematic in a CAF $\mathcal{F} = (A, R, cl)$ iff $a, b \in A$, $(a,b) \notin R$, and there is $a' \in A$ with $cl(a') = cl(a)$ and $(a', b) \in R$. The set $wfp(\mathcal{F}) = \{(a,b) \mid (a,b)$ is wf-problematic in $\mathcal{F}\}$ is called the wf-problematic part of $\mathcal{F}$.*

Intuitively, the wf-problematic part of a CAF $\mathcal{F}$ consists of those attacks that are missing for $\mathcal{F}$ to be well-formed (cf. Figure 4.1). Indeed, $\mathcal{F}$ is a wfCAF if and only if $wfp(\mathcal{F}) = \emptyset$.

The four new classes can be characterized as follows:

Proposition 4.10. *Let $\mathcal{F} = (A, R, cl)$ be a CAF. Then*

- *$\mathcal{F} \in \mathcal{R}_1$-**CAF** iff $(a,b) \in wfp(\mathcal{F})$ implies $(b,a) \notin wfp(\mathcal{F})$;*
- *$\mathcal{F} \in \mathcal{R}_2$-**CAF** iff there are no arguments a, a', b, b' in $\mathcal{F}$ with $cl(a) = cl(a')$ and $cl(b) = cl(b')$ such that $(a,b) \in wfp(\mathcal{F})$, $(b,a) \notin R$, $(a',b) \in R$, and either $(b,a') \in R$ or $((a',b') \notin R$ and $(b',a') \notin R)$;*
- *$\mathcal{F} \in \mathcal{R}_3$-**CAF** iff $(a,b) \in wfp(\mathcal{F})$ implies $(b,a) \in R$;*
- *$\mathcal{F} \in \mathcal{R}_4$-**CAF** iff there are no arguments a, a', b, b' in $\mathcal{F}$ with $cl(a) = cl(a')$ and $cl(b) = cl(b')$ such that $(a,b) \in wfp(\mathcal{F})$, $(b,a) \notin R$, $(a',b) \in R$, and either $(b,a') \notin R$ or $((a',b') \notin R$ and $(b',a') \notin R)$.*

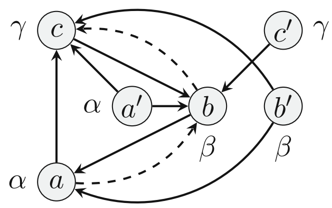

Figure 4.2: A CAF which shows that $\mathcal{R}_3$-$\mathbf{CAF}_{tr} \not\subseteq \mathcal{R}_j$-$\mathbf{CAF}_{tr}$ for $j \in \{1, 2, 4\}$. Dashed arrows are edges in the wf-problematic part.

Proof. Here we consider $\mathcal{R}_1$-**CAF**. The remaining cases can be found in the appendix (Lemma A.1 for $\mathcal{R}_2$-**CAF**, Lemma A.2 for $\mathcal{R}_3$-**CAF**, and Lemma A.3 for $\mathcal{R}_4$-**CAF**).

"$\Longrightarrow$": By contrapositive. Suppose there is $(a, b) \in \mathit{wfp}(\mathcal{F})$ such that $(b, a) \in \mathit{wfp}(\mathcal{F})$. Towards a contradiction, assume $\mathcal{F} \in \mathcal{R}_1$-**CAF**. Then there is a PCAF $\mathcal{P} = (A, R', cl, \succ)$ such that $\mathcal{R}_1(\mathcal{P}) = \mathcal{F}$. Since Reduction 1 can only delete but not introduce attacks, and since the underlying CAF of $\mathcal{P}$ must be well-formed, $(a, b) \in R'$ and $(b, a) \in R'$. However, then also $(b \succ a)$ and $(a \succ b)$ which means that $\mathcal{P}$ is not asymmetric. Contradiction.

"$\Longleftarrow$": Suppose that $(a, b) \in \mathit{wfp}(\mathcal{F})$ implies $(b, a) \notin \mathit{wfp}(\mathcal{F})$. Then $\mathcal{R}_1(\mathcal{P}) = \mathcal{F}$ for the PCAF $\mathcal{P} = (A, R', cl, \succ)$ with $R' = R \cup \{(a, b) \mid (a, b) \in \mathit{wfp}(\mathcal{F})\}$ as well as $a \succ b$ iff $(b, a) \in R' \setminus R$. The underlying CAF of $\mathcal{P}$ is well-formed since $\mathit{wfp}((A, R', cl)) = \emptyset$. Furthermore, $\succ$ is asymmetric since $(a, b) \in \mathit{wfp}(\mathcal{F})$ implies $(b, a) \notin \mathit{wfp}(\mathcal{F})$ and by construction of $\mathcal{P}$. □

The above characterizations give us some insights into the effect of the various reductions on wfCAFs. Indeed, the similarity between the characterizations of $\mathcal{R}_1$-**CAF** and $\mathcal{R}_3$-**CAF**, resp. $\mathcal{R}_2$-**CAF** and $\mathcal{R}_4$-**CAF**, can intuitively be explained by the fact that Reductions 1 and 3 only remove attacks, while Reductions 2 and 4 can also introduce attacks. Proposition 4.10 allows us to decide in polynomial time whether a given CAF $\mathcal{F}$ can be obtained by applying one of the four preference reductions to a PCAF. Moreover, in the proof of Proposition 4.10 we see how, given $\mathcal{F} \in \mathcal{R}_i$-**CAF**, we can construct a PCAF $\mathcal{P}$ such that $\mathcal{R}_i(\mathcal{P}) = \mathcal{F}$ in polynomial time.

But what happens if we restrict ourselves to transitive preferences? Analogously to $\mathcal{R}_i$-**CAF** (cf. Definition 4.6), by $\mathcal{R}_i$-$\mathbf{CAF}_{tr}$ we denote the set of CAFs obtained by applying Reduction i to PCAFs with a transitive preference relation. It is clear that $\mathcal{R}_i$-$\mathbf{CAF}_{tr} \subseteq \mathcal{R}_i$-**CAF** for all $i \in \{1, 2, 3, 4\}$. Moreover, in the proof of Proposition 4.7 we actually made use of transitive preferences, i.e., **wfCAF** $\subset \mathcal{R}_i$-$\mathbf{CAF}_{tr}$ for all $i \in \{1, 2, 3, 4\}$. Interestingly, however, the relationships between the classes $\mathcal{R}_i$-$\mathbf{CAF}_{tr}$ is different to that between $\mathcal{R}_i$-**CAF** (Proposition 4.8). Specifically, $\mathcal{R}_3$-$\mathbf{CAF}_{tr}$ is not contained in the other classes. The reason for this is that, in certain PCAFs $\mathcal{P}$, transitivity can force $a_1 \succ a_n$ via $a_1 \succ a_2 \succ \ldots \succ a_n$ such that $(a_n, a_1) \in \mathcal{P}$ but $(a_1, a_n) \notin \mathcal{P}$. In this case, only Reduction 3 leaves the attacks between a_1 and a_n unchanged.

Proposition 4.11. *For all $i, j \in \{1,2,3,4\}$ with $i \neq j$ it holds that $\mathcal{R}_i$-$\mathbf{CAF}_{tr} \not\subseteq \mathcal{R}_j$-$\mathbf{CAF}_{tr}$. Moreover, for all $i \in \{1,2,4\}$ and all $j \in \{1,2,3,4\}$ such that $i \neq j$ it holds that $\mathcal{R}_i$-$\mathbf{CAF}_{tr} \not\subseteq \mathcal{R}_j$-$\mathbf{CAF}$.*

Proof. Note that the preference relations of the PCAFs used in the proof of Proposition 4.8 are transitive. We therefore have $\mathcal{R}_i$-$\mathbf{CAF}_{tr} \not\subseteq \mathcal{R}_j$-$\mathbf{CAF}$, which also means $\mathcal{R}_i$-$\mathbf{CAF}_{tr} \not\subseteq \mathcal{R}_j$-$\mathbf{CAF}_{tr}$, for every $i \in \{1,2,4\}$ and $j \in \{1,2,3,4\}$ such that $i \neq j$.

It remains to show $\mathcal{R}_3$-$\mathbf{CAF}_{tr} \not\subseteq \mathcal{R}_j$-$\mathbf{CAF}_{tr}$ for $j \in \{1,2,4\}$. Let $\mathcal{F}$ be the CAF shown in Figure 4.2. $\mathcal{F}$ is in $\mathcal{R}_3$-$\mathbf{CAF}_{tr}$: to see this, let $\mathcal{P}$ be the PCAF with the same arguments and attacks as $\mathcal{F}$, and additionally attacks (a,b) and (b,c); Moreover, let $c \succ b$, $b \succ a$, and $c \succ a$; the attack (a,c) is not deleted by Reduction 3 if there is no attack (c,a); Thus, $\mathcal{R}_3(\mathcal{P}) = \mathcal{F}$. We show that $\mathcal{F} \notin \mathcal{R}_j$-$\mathbf{CAF}_{tr}$ for $j \in \{1,2,4\}$.

- $\mathcal{F}$ is not in $\mathcal{R}_1$-$\mathbf{CAF}_{tr}$ since a PCAF that reduces to $\mathcal{F}$ would need to have $c \succ b$, $b \succ a$, and therefore also $c \succ a$. But Reduction 1 would delete the attack (a,c).
- Towards a contradiction, assume there is a PCAF $\mathcal{P}$ such that $\mathcal{R}_2(\mathcal{P}) = \mathcal{F}$. First, we show that $(a,b) \in \mathcal{P}$, $(b,a) \in \mathcal{P}$, $(b,c) \in \mathcal{P}$, and $(c,b) \in \mathcal{P}$.
  - Assume $(a,b) \notin \mathcal{P}$. Then two things most hold. Firstly, it must be that $(b,a) \in \mathcal{P}$, otherwise $(b,a) \notin \mathcal{F}$. Secondly, $(a',b) \notin \mathcal{P}$, otherwise the underlying CAF of $\mathcal{P}$ would not be well-formed. This means that (a',b) must have been introduced into $\mathcal{F}$ by applying Reduction 2, i.e., by reversing (b,a'). Therefore, $(b,a') \in \mathcal{P}$. But then also $(b',a') \in \mathcal{P}$, otherwise the underlying CAF of $\mathcal{P}$ is not well-formed. But then, by the definition of Reduction 2, either $(b',a') \in \mathcal{F}$ or $(a',b') \in \mathcal{F}$, which is not the case. Contradiction.
  - Assume $(b,a) \notin \mathcal{P}$. Then, since the underlying CAF of $\mathcal{P}$ must be well-formed, $(b',a) \notin \mathcal{P}$. This means $(a,b') \in \mathcal{P}$, otherwise we cannot obtain $\mathcal{F}$ from $\mathcal{P}$ via Reduction 2. This means that $(a',b') \in \mathcal{P}$, which is not possible since neither $(a',b') \in \mathcal{F}$ nor $(b',a') \in \mathcal{P}$.
  - Assume $(b,c) \notin \mathcal{P}$. Then two things most hold. Firstly, it must be that $(c,b) \in \mathcal{P}$, otherwise $(c,b) \notin \mathcal{F}$. Secondly, $(b',c) \notin \mathcal{P}$, otherwise the underlying CAF of $\mathcal{P}$ would not be well-formed. This means that (b',c) must have been introduced into $\mathcal{F}$ by applying Reduction 2, i.e., by reversing (c,b'). Therefore, $(c,b') \in \mathcal{P}$. But then also $(c',b') \in \mathcal{P}$, otherwise the underlying CAF of $\mathcal{P}$ is not well-formed. But then, by the definition of Reduction 2, either $(c',b') \in \mathcal{F}$ or $(b',c') \in \mathcal{F}$, which is not the case. Contradiction.
  - Assume $(c,b) \notin \mathcal{P}$. Then, since the underlying CAF of $\mathcal{P}$ must be well-formed, $(c',b) \notin \mathcal{P}$. This means $(b,c') \in \mathcal{P}$, otherwise we cannot obtain $\mathcal{F}$ from $\mathcal{P}$ via Reduction 2. This means that $(b',c') \in \mathcal{P}$, which is not possible since neither $(b',c') \in \mathcal{F}$ nor $(c',b') \in \mathcal{P}$.

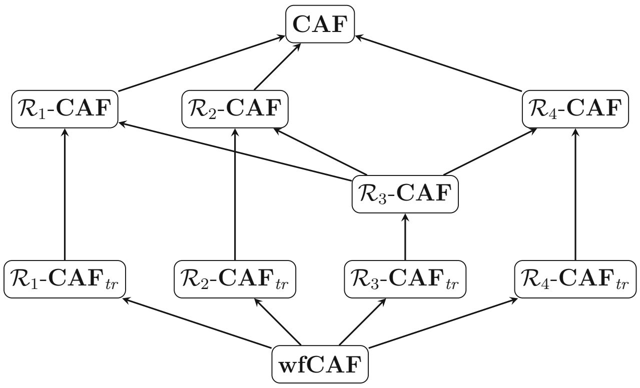

Figure 4.3: Relations between the various CAF-classes. An arrow indicates that a class is a strict subset of the other, e.g. $\mathcal{R}_3$-**CAF** $\subset$ $\mathcal{R}_4$-**CAF**.

Since $(a, b) \in \mathcal{P}$, $(b, a) \in \mathcal{P}$, $(b, c) \in \mathcal{P}$, and $(c, b) \in \mathcal{P}$, the only way to obtain $\mathcal{F}$ form $\mathcal{P}$ via Reduction 2 is to set $c \succ b$ and $b \succ a$. But then $c \succ a$ which means that $(a, c) \notin \mathcal{F}$. Contradiction, i.e., $\mathcal{F} \notin \mathcal{R}_2$-**CAF**$_{tr}$.

- Now assume there is a PCAF $\mathcal{P}'$ such that $\mathcal{R}_4(\mathcal{P}') = \mathcal{F}$. It must be that $(a, b) \in \mathcal{P}'$ since we cannot obtain $(a', b) \in \mathcal{R}_4(\mathcal{P}')$ and $(b, a') \notin \mathcal{R}_4(\mathcal{P}')$ without $(a', b) \in \mathcal{P}'$. Analogously, it must be that $(b, c) \in \mathcal{P}'$. Then in order to have $\mathcal{R}_4(\mathcal{P}') = \mathcal{F}$ we need to set $c \succ b$ and $b \succ a$. But then $c \succ a$ which means that it cannot be that $(a, c) \in \mathcal{R}_4(\mathcal{P}')$ and $(c, a) \notin \mathcal{R}_4(\mathcal{P}')$. Contradiction, i.e., $\mathcal{F} \notin \mathcal{R}_4$-**CAF**$_{tr}$. □

The above result also implies that $\mathcal{R}_i$-**CAF**$_{tr}$ $\subset$ $\mathcal{R}_i$-**CAF** for $i \in \{1, 2, 4\}$ since we have $\mathcal{R}_3$-**CAF**$_{tr}$ $\subseteq$ $\mathcal{R}_3$-**CAF** $\subset$ $\mathcal{R}_i$-**CAF** (cf. Proposition 4.8) and $\mathcal{R}_3$-**CAF**$_{tr}$ $\not\subseteq$ $\mathcal{R}_i$-**CAF**$_{tr}$ (cf. Proposition 4.11), which implies $\mathcal{R}_i$-**CAF**$_{tr}$ $\neq$ $\mathcal{R}_i$-**CAF**. It is also easy to see that $\mathcal{R}_3$-**CAF**$_{tr}$ $\subset$ $\mathcal{R}_3$-**CAF**: take the CAF from Figure 4.2 and add the additional attack (c, a). The resulting CAF is in $\mathcal{R}_3$-**CAF** since we do not need to set the preference $c \succ a$, whereas it is not in $\mathcal{R}_3$-**CAF**$_{tr}$ since $c \succ a$ is enforced by $c \succ b \succ a$. Figure 4.3 summarizes the relationship between the CAF-classes.

We will not characterize all four classes $\mathcal{R}_i$-**CAF**$_{tr}$ for transitive preferences. Indeed, while each $\mathcal{R}_i$-**CAF** and $\mathcal{R}_i$-**CAF**$_{tr}$ are distinct syntactically, we will show that their semantic properties (cf. Section 4.3) and their computational complexity (cf. Section 4.4) are the same. However, we will characterize $\mathcal{R}_1$-**CAF**$_{tr}$ as this will prove useful when analyzing the computational complexity of PCAFs using Reduction 1. Note that $wfp(\mathcal{F})$ can be seen as a directed graph, with an edge between vertices a and b whenever $(a, b) \in wfp(\mathcal{F})$. Thus, we may use notions such as paths and cycles in the wf-problematic part of a CAF.

Proposition 4.12. *$\mathcal{F} \in \mathcal{R}_1$-**CAF**$_{tr}$ for a CAF $\mathcal{F}$ iff (1) $wfp(\mathcal{F})$ is acyclic and (2) $(a, b) \in \mathcal{F}$ implies that there is no path from a to b in $wfp(\mathcal{F})$.*

Proof. Let $\mathcal{F} = (A, R, cl)$.

- Suppose $wfp(\mathcal{F})$ is acyclic and there is no $(a, b) \in \mathcal{F}$ with a path from a to b in $wfp(\mathcal{F})$. Construct the PCAF $\mathcal{P} = (A, R', cl, \succ)$ with $R' = R \cup \{(a, b) \mid (a, b) \in wfp(\mathcal{F})\}$ and $b \succ a$ iff there is a path from a to b in $wfp(\mathcal{F})$. (A, R', cl) is well-formed by construction. $\succ$ is transitive because if there is a path from a to b and from b to c, then there is also a path from a to c. $\succ$ is asymmetric because otherwise there would be a path from a to b and from b to a, which again would mean that there is a cycle. It remains to show that $\mathcal{R}_1(\mathcal{P}) = \mathcal{F}$. Let (a, b) be any attack in $\mathcal{P}$. We distinguish two cases:
  - $(a, b) \in \mathcal{F}$. Then, since there is no path from a to b in $wfp(\mathcal{F})$, $b \not\succ a$. Therefore, $(a, b) \in \mathcal{R}_1(\mathcal{P})$.
  - $(a, b) \notin \mathcal{F}$. Then, by construction, $(a, b) \in wfp(\mathcal{F})$ and therefore $b \succ a$. Thus, (a, b) is removed from $\mathcal{P}$ by Reduction 1, i.e., $(a, b) \notin \mathcal{R}_1(\mathcal{P})$.

  Note also that, by construction of $\mathcal{P}$, there can be no $(a, b) \in \mathcal{F}$ such that $(a, b) \notin \mathcal{P}$.

- Suppose $wfp(\mathcal{F})$ is cyclic. Then there are arguments $x_1, \ldots, x_n \in \mathcal{F}$ such that $x_1 = x_n$ and $(x_i, x_{i+1}) \in wfp(\mathcal{F})$ for all $1 \leq i < n$. Towards a contradiction, assume there is a PCAF $\mathcal{P} = (A, R', cl, \succ)$ such that $\mathcal{R}_1(\mathcal{P}) = \mathcal{F}$. Then $(x_i, x_{i+1}) \in \mathcal{P}$ for all $1 \leq i < n$, otherwise (A, R', cl) would not be well-formed. In order to have $\mathcal{R}_1(\mathcal{P}) = \mathcal{F}$ we must have $x_{i+1} \succ x_i$ for all $1 \leq i < n$. But then, by transitivity and since $x_1 = x_n$ we obtain $x_1 \succ x_1$, which is in contradiction to $\succ$ being asymmetric.

  On the other hand, suppose there is an attack $(a, b) \in \mathcal{F}$ with a path from a to b in $wfp(\mathcal{F})$. Let us denote this path as $x_1, \ldots, x_n$ with $x_1 = a$ and $x_n = b$. By the same argument as above, if there were a PCAF $\mathcal{P} = (A, R', cl, \succ)$ such that $\mathcal{R}_1(\mathcal{P}) = \mathcal{F}$, then $x_n \succ x_1$, i.e., $b \succ a$. But then $(a, b) \notin \mathcal{R}_1(\mathcal{P})$. Contradiction. □

In summary, we have shown that the four new CAF-classes that result from applying preferences to wfCAFs lie strictly inbetween wfCAFs and general CAFs (see Proposition 4.7) and that they are distinct from each other (see Propositions 4.8 and 4.11). Figure 4.3 summarizes the relationship between the CAF-classes. Furthermore, we characterize the four classes (see Proposition 4.10), which allows us to take any CAF, and, in polynomial time, decide whether this CAF belongs to one of the four classes.

From a high-level point of view, these characterization results yield insights into the expressiveness of argumentation formalisms that allow for preferences. Propositions 4.10 and 4.12 show which situations can be captured by formalisms which (i) construct attacks based on the claim of the attacking argument (i.e., formalisms with well-formed attack relation) and (ii) incorporate asymmetric or transitive preference relations on arguments using one of the four reductions.

4.3 Semantic Properties

There are key differences between wfCAFs and general CAFs with respect to semantic properties. It has been shown (Dvořák, Rapberger, and Woltran 2023) that inherited and hybrid variants of stable and preferred semantics coincide on wfCAFs but not on general CAFs (cf. Figure 2.5). This simplifies the choice of semantics when working with wfCAFs. Moreover, wfCAFs, unlike general CAFs, preserve I-maximality under most maximization-based semantics (cf. Figure 2.5). This leads to more intuitive behavior of these semantics when considering extensions on the claim-level. As we have seen in Section 4.2, resolving preferences on wfCAFs results in four new CAF-classes that, from a syntactic perspective, lie inbetween wfCAFs and general CAFs. We now investigate whether these new CAF-classes retain the benefits of wfCAFs when it comes to semantic properties. We summarize and discuss our results at the end of this section (cf. Theorem 4.26 and Figure 4.6).

Firstly, we observe that the basic relations between semantics carry over from general CAFs, i.e., if we have $\sigma_\mu(\mathcal{F}) \subseteq \tau_\nu(\mathcal{F})$ for two CAF-semantics σ_μ, τ_ν and all CAFs $\mathcal{F}$, then we also have also $\sigma^i_\mu(\mathcal{P}) \subseteq \tau^i_\nu(\mathcal{P})$ for all PCAFs $\mathcal{P}$ and Reduction $i \in \{1,2,3,4\}$. Likewise, if we have $\sigma_\mu(\mathcal{F}) \not\subseteq \tau_\nu(\mathcal{F})$, then we also have $\sigma^i_\mu(\mathcal{P}) \not\subseteq \tau^i_\nu(\mathcal{P})$.

Secondly, we note that Reductions 2, 3, and 4 cannot entirely remove conflicts between arguments, and that therefore the resolution of preferences has no impact on conflict-free extensions (both on the argument- and claim-level) under these preference reductions.

Lemma 4.13. *Let $\mathcal{P} = (A, R, cl, \succ)$ be a PCAF and let $\mathcal{R}_i(\mathcal{P}) = (A, R', cl)$ with $i \in \{2,3,4\}$. Then $cf((A,R)) = cf((A,R'))$ and $cf_{inh}((A,R,cl)) = cf_{inh}((A,R',cl))$.*

Proof. Let $\mathcal{P} = (A, R, cl, \succ)$ be a PCAF and let $\mathcal{R}_i(\mathcal{P}) = (A, R', cl)$ for $i \in \{2,3,4\}$. By definition of Reduction i, if $(a,b) \in R$ then either $(a,b) \in R'$ or $(b,a) \in R'$. Conversely, if $(a,b) \in R'$, then it must be that either $(a,b) \in R$ or $(b,a) \in R$. Thus, for any $S \subseteq A$ we have $S \in cf((A,R))$ iff $S \in cf((A,R'))$. This further implies that for any $C \subseteq cl(A)$ we have $C \in cf_{inh}((A,R,cl))$ iff $C \in cf_{inh}((A,R',cl))$. □

The fact that Reductions 2–4 do not remove conflicts, and the well-formedness of a PCAF's underlying CAF, allow us to show that inherited stable semantics and hybrid admissibility-based stable semantics coincide under Reductions 2–4. Under Reduction 1 the two semantics do not coincide.

Proposition 4.14. *$stb^i_{inh}(\mathcal{P}) = stb\text{-}adm^i_{hyb}(\mathcal{P})$, where $i \in \{2,3,4\}$, holds for every PCAF $\mathcal{P}$.*

Proof. $stb_{inh}(\mathcal{F}) \subseteq stb\text{-}adm_{hyb}(\mathcal{F})$ holds for all CAFs. We must show that $stb\text{-}adm_{hyb}(\mathcal{F}) \subseteq stb_{inh}(\mathcal{F})$ for all $\mathcal{F} \in \mathcal{R}_i$-**CAF**, where $i \in \{2,3,4\}$. Let $\mathcal{F} = (A,R,cl) \in \mathcal{R}_i$-**CAF**, and let $\mathcal{P} = (A, R', cl, \succ)$ be a PCAF such that $\mathcal{R}_i(\mathcal{P}) = \mathcal{F}$. Moreover, let $C \in stb\text{-}adm_{hyb}(\mathcal{F})$. Then there is an argument-set $S \subseteq A$ such that $S \in adm(\mathcal{F})$, $cl(S) = C$, and $C \cup S^*_\mathcal{F} = cl(A)$. Let $S' = S \cup \{x \in A \setminus S \mid (x,y) \notin R, (y,x) \notin R \text{ for all } y \in S\}$, i.e., S' is obtained by adding all arguments to S that are not in conflict with S. We show that $C \in stb_{inh}(\mathcal{F})$ by showing the following:

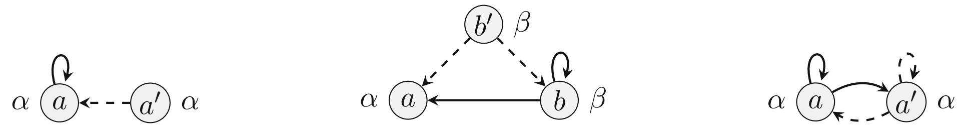

(a) $\mathcal{R}_1(\mathcal{P})$ from Proposition 4.15.
(b) $\mathcal{R}_1(\mathcal{P})$ from Proposition 4.17.
(c) $\mathcal{R}_2(\mathcal{P})$ from Proposition 4.18.

Figure 4.4: CAFs used to show that some variants of stable semantics do not coincide under Reductions 1 and 2. Dashed arrows are attacks in the wf-problematic part of the CAF.

1. $cl(S') = C$: clearly, $cl(S) \subseteq cl(S')$. Now consider any $x \in S' \setminus S$. Then S does not defeat $cl(x)$, since there is no conflict between S and x. Since $cl(S) \cup S^*_{\mathcal{F}} = cl(A)$ there must be $x' \in S$ with $cl(x') = cl(x)$. Thus, $cl(S) \supseteq cl(S')$.

2. $S' \in cf(A, R)$: since $S \in adm(A, R)$, there is no conflict between any two arguments in S. Moreover, by construction, there is no conflict between arguments in S and arguments in $S' \setminus S$. It remains to show there is no conflict between any two arguments in $S' \setminus S$. Towards a contradiction, assume there are $x', y' \in S' \setminus S$ such that $(x', y') \in R$. Since there is no conflict between S and x' (resp. y'), S does not defeat $cl(x')$ (resp. $cl(y')$). Since $cl(S) \cup S^*_{\mathcal{F}} = cl(A)$, there must be $x, y \in S$ with $cl(x) = cl(x')$ and $cl(y) = cl(y')$. Since Reductions 2,3,4 cannot remove conflicts, we have $(x', y') \in R'$ or $(y', x') \in R'$ in the original PCAF $\mathcal{P}$. By the well-formedness of $\mathcal{P}$, we have $(x, y') \in R'$ or $(y, x') \in R'$. Since Reductions 2,3,4 cannot remove conflicts, if $(x, y') \in R'$ then $(x, y') \in R$ or $(y', x) \in R$, and if $(y, x') \in R'$ then $(y, x') \in R$ or $(x', y) \in R$. But then either $x' \notin S'$ or $y' \notin S'$. Contradiction.

3. for all $z \in A \setminus S'$ there is $x \in S'$ such that $(x, z) \in R$: let $z \in A \setminus S'$. Then there must be $x \in S$ such that either $(x, z) \in R$ or $(z, x) \in R$, otherwise we would have $z \in S'$. If $(z, x) \in R$ but $(x, z) \notin R$, there must be $y \in S$ such that $(y, z) \in R$, otherwise we would have $S \notin adm(\mathcal{F})$. □

Proposition 4.15. *There is a PCAF $\mathcal{P}$ such that $stb^1_{inh}(\mathcal{P}) \neq stb\text{-}adm^1_{hyb}(\mathcal{P})$.*

Proof. Let $\mathcal{P} = (A, R, cl, \succ)$ with $A = \{a, a'\}$, $R = \{(a, a), (a', a)\}$, $cl(a) = cl(a') = \alpha$, and $a \succ a'$. Figure 4.4a depicts $\mathcal{R}_1(\mathcal{P}) = (A, R', cl)$, i.e., $R' = \{(a, a)\}$. Note that $stb_{inh}(\mathcal{R}_1(\mathcal{P})) = \emptyset$ while $stb\text{-}adm_{hyb}(\mathcal{R}_1(\mathcal{P})) = stb\text{-}cf_{hyb}(\mathcal{R}_1(\mathcal{P})) = \{\{\alpha\}\}$. □

Similarly, we can show that both variants (conflict-free and admissibility-based) of stable semantics coincide under Reductions 3 and 4, but not under Reductions 1 and 2.

Proposition 4.16. *$stb\text{-}adm^i_{hyb}(\mathcal{P}) = stb\text{-}cf^i_{hyb}(\mathcal{P})$, where $i \in \{3, 4\}$, holds for every PCAF $\mathcal{P}$.*

Proof. $stb\text{-}adm_{hyb}(\mathcal{F}) \subseteq stb\text{-}cf_{hyb}(\mathcal{F})$ holds for all CAFs. We show that $stb\text{-}cf_{hyb}(\mathcal{F}) \subseteq stb\text{-}adm_{hyb}(\mathcal{F})$ for $\mathcal{F} \in \mathcal{R}_i\text{-}\mathbf{CAF}$, where $i \in \{3,4\}$. Let $\mathcal{F} = (A, R, cl) \in \mathcal{R}_i\text{-}\mathbf{CAF}$, and let $\mathcal{P} = (A, R', cl, \succ)$ be a PCAF such that $\mathcal{R}_i(\mathcal{P}) = \mathcal{F}$. Moreover, let $C \in stb\text{-}cf_{hyb}(\mathcal{F})$. Then there is an argument-set $S \subseteq A$ such that $S \in cf(A, R)$, $cl(S) = C$, and $C \cup S^*_{\mathcal{F}} = cl(A)$. We show that $C \in stb\text{-}adm_{hyb}(\mathcal{F})$ by showing that $S \in adm(\mathcal{F})$:

Consider any $x \in S$ and $y \in A \setminus S$ such that $(y, x) \in R$ but $(x, y) \notin R$. Under Reductions 3 and 4 a non-symmetric attack (y, x) in $\mathcal{R}_3(\mathcal{P})$ means that (y, x) was also present in the original PCAF $\mathcal{P}$, i.e., $(y, x) \in R'$. Towards a contradiction, assume that S does not defeat $cl(y)$ in $\mathcal{F}$. Since $cl(S) \cup S^*_{\mathcal{F}} = cl(A)$, this means that there is $y' \in S$ with $cl(y') = cl(y)$. By the well-formedness of $\mathcal{P}$ this further implies $(y', x) \in R'$. But Reductions 3 and 4 cannot remove conflicts, i.e., either $(y', x) \in R$ or $(x, y') \in R$. Thus, $S \notin cf(A, R)$. Contradiction. Therefore, S defeats $cl(y)$ in $\mathcal{F}$, i.e., there is $z \in S$ such that $(z, y) \in R$. We can conclude that $S \in adm(\mathcal{F})$. □

Proposition 4.17. *There is a PCAF $\mathcal{P}$ such that $stb\text{-}adm^1_{hyb}(\mathcal{P}) \neq stb\text{-}cf^1_{hyb}(\mathcal{P})$.*

Proof. Let $\mathcal{P} = (A, R, cl, \succ)$ with $A = \{a, b, b'\}$, $R = \{(b, a), (b, b), (b', a), (b', b)\}$, $cl(a) = \alpha$, $cl(b) = cl(b') = \beta$, and $a \succ b'$, $b \succ b'$. The attacks (b', a) and (b', b) are deleted in $\mathcal{R}_1(\mathcal{P})$, see Figure 4.4b. Moreover, $stb\text{-}adm_{hyb}(\mathcal{R}_1(\mathcal{P})) = \emptyset$ but $stb\text{-}cf_{hyb}(\mathcal{R}_1(\mathcal{P})) = \{\{\alpha, \beta\}\}$. □

Proposition 4.18. *There is a PCAF $\mathcal{P}$ such that $stb\text{-}adm^2_{hyb}(\mathcal{P}) \neq stb\text{-}cf^2_{hyb}(\mathcal{P})$.*

Proof. Consider the PCAF $\mathcal{P} = (A, R, cl, \succ)$ with $A = \{a, a'\}$, $R = \{(a, a), (a', a)\}$, $cl(a) = cl(a') = \alpha$, and $a \succ a'$. Then $\mathcal{R}_2(\mathcal{P}) = (A, R', cl)$ with $R' = \{(a, a), (a, a')\}$, see Figure 4.4c. Note that $stb_{inh}(\mathcal{R}_2(\mathcal{P})) = stb\text{-}adm_{hyb}(\mathcal{R}_2(\mathcal{P})) = \emptyset$ while $stb\text{-}cf_{hyb}(\mathcal{R}_2(\mathcal{P})) = \{\{\alpha\}\}$. □

Before investigating whether inherited and hybrid preferred semantics coincide, we examine the I-maximality property. The following is analogous to Definition 2.25.

Definition 4.19 (I-maximality for PCAFs)**.** *σ^i_μ is I-maximal for PCAFs if, for all PCAFs $\mathcal{P}$ and all $C, D \in \sigma^i_\mu(\mathcal{P})$, $C \subseteq D$ implies $C = D$.*

From known properties of wfCAFs (cf. Figure 2.5) it follows directly that $naive^i_{inh}$, where $i \in \{1, 2, 3, 4\}$, is not I-maximal for PCAFs. Likewise, from the properties of general CAFs we know that $naive^i_{hyb}$ and prf^i_{hyb} are I-maximal for all $i \in \{1, 2, 3, 4\}$. It remains to investigate I-maximality of prf^i_{inh} and all inherited and hybrid variants of stable, semi-stable, and stage semantics.

As it turns out, Reduction 3 manages to preserve I-maximality in all cases except for inherited and hybrid stage semantics.

Proposition 4.20. *prf^3_{inh}, sem^3_{inh}, sem^3_{hyb}, stb^3_{inh}, $stb\text{-}adm^3_{hyb}$, and $stb\text{-}cf^3_{hyb}$ are I-maximal for PCAFs.*

Proof. We show this for prf^3_{inh}. The other results follow from $sem^3_{inh}(\mathcal{P}) \subseteq prf^3_{inh}(\mathcal{P})$ (by properties of general CAFs), $sem^3_{hyb}(\mathcal{P}) \subseteq prf^3_{inh}(\mathcal{P})$ (by properties of general CAFs), and $stb^3_{inh}(\mathcal{P}) = stb\text{-}adm^3_{hyb}(\mathcal{P}) = stb\text{-}cf^3_{hyb}(\mathcal{P}) \subseteq prf^3_{inh}(\mathcal{P})$ (by Propositions 4.14 and 4.16 as well as properties of general CAFs). Towards a contradiction, assume there is a PCAF $\mathcal{P} = (A, R, cl, \succ)$ such that $C \subset D$ for some $C, D \in prf^3_{inh}(\mathcal{P})$. Then there must be $S \subseteq A$ such that $S \in prf(\mathcal{R}_3(\mathcal{P}))$ and $cl(S) = C$, as well as $T \subseteq A$ with $T \in prf(\mathcal{R}_3(\mathcal{P}))$ and $cl(T) = D$. Observe that $S \not\subseteq T$, otherwise $S \notin prf(\mathcal{R}_3(\mathcal{P}))$. Thus, there is $x \in S$ (with $cl(x) \in C$) such that $x \notin T$. However, $cl(x) \in D$ since $C \subset D$, i.e., there is some $x' \in T$ such that $cl(x') = cl(x)$. There are two possibilities for why x is not in T:

1. $T \cup \{x\} \notin cf(\mathcal{R}_3(\mathcal{P}))$. By Lemma 4.13, $T \cup \{x\} \notin cf((A, R, cl))$. Therefore, there is some $y \in T$ such that $y \notin S$ and either $(x, y) \in \mathcal{P}$ or $(y, x) \in \mathcal{P}$. Actually, it cannot be that $(x, y) \in \mathcal{P}$, otherwise, by the well-formedness of (A, R, cl), we would have $(x', y) \in \mathcal{P}$ which, also by Lemma 4.13, would mean that $T \notin cf(\mathcal{R}_3(\mathcal{P}))$. Thus, $(y, x) \in \mathcal{P}$. Since $(x, y) \notin \mathcal{P}$, and by the definition of Reduction 3, $(y, x) \in \mathcal{R}_3(\mathcal{P})$. S must defend x from y in $\mathcal{R}_3(\mathcal{P})$, i.e., there is some $z \in S$ such that $(z, y) \in \mathcal{R}_3(\mathcal{P})$. Therefore, also $(z, y) \in \mathcal{P}$. Since we have that $C \subset D$ there is some $z' \in T$ such that $cl(z') = cl(z)$. $(z', y) \in \mathcal{P}$ by the well-formedness of (A, R, cl). But then, by Lemma 4.13, $T \notin cf(\mathcal{R}_3(\mathcal{P}))$. Contradiction.

2. x is not defended by T. Then there is some $y \in A$ such that $(y, x) \in \mathcal{R}_3(\mathcal{P})$ and such that y is not attacked by any argument in T. But S must defend x against y in $\mathcal{R}_3(\mathcal{P})$, i.e., there is $z \in S$ such that $(z, y) \in \mathcal{R}_3(\mathcal{P})$. Then also $(z, y) \in \mathcal{P}$. Since $C \subset D$ there is some $z' \in T$ such that $cl(z') = cl(z)$. $(z', y) \in \mathcal{P}$ by the well-formedness of (A, R, cl). It cannot be that $(z', y) \in \mathcal{R}_3(\mathcal{P})$, i.e., $y \succ z'$. But then, by the definition of Reduction 3, we must have $(y, z') \in \mathcal{P}$ and also $(y, z') \in \mathcal{R}_3(\mathcal{P})$, which means that T is attacked by y but not defended against it, i.e., $T \notin adm(\mathcal{R}_3(\mathcal{P}))$. Contradiction. □

For negative results, it suffices to show that I-maximality is not preserved for transitive preference orderings to obtain results for the more general case.

Proposition 4.21. *stg^3_{inh} and stg^3_{hyb} are not I-maximal for PCAFs, even when considering only transitive preferences.*

Proof. Let $\mathcal{F} = (A, R, cl)$ be the CAF shown in Figure 4.5a, and let $F = (A, R)$ be its underlying AF. Clearly, $\mathcal{F} \in \mathcal{R}_3\text{-}\mathbf{CAF}_{tr}$.

We can see that $cf(F) = \{\emptyset, \{a\}, \{a'\}, \{b\}, \{c\}, \{a, a'\}, \{a', c\}\}$ and thus $naive(F) = \{\{a, a'\}, \{a', c\}, \{b\}\}$.

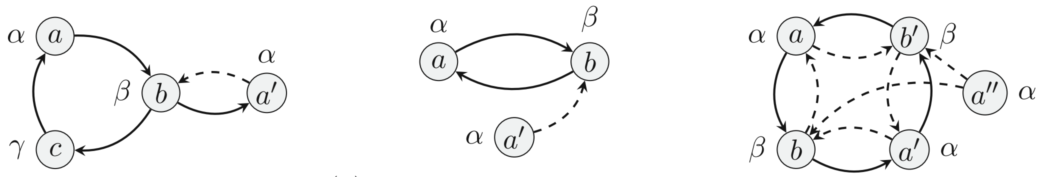

(a) in $\mathcal{R}_3$-$\mathbf{CAF}_{tr}$ (b) in $\mathcal{R}_1$-$\mathbf{CAF}_{tr}$ and $\mathcal{R}_4$-$\mathbf{CAF}_{tr}$ (c) in $\mathcal{R}_2$-$\mathbf{CAF}_{tr}$

Figure 4.5: CAFs used as counter examples for I-maximality of some semantics. Dashed arrows are edges in the respective wf-problematic part.

Regarding stg^3_{inh}, we have $\{a,a'\}^{\oplus}_F = \{a,a',b\}$, $\{a',c\}^{\oplus}_F = \{a,a',c\}$, and $\{b\}^{\oplus}_F = \{a',b,c\}$. The three ranges are incomparable, i.e., $stg(F) = naive(F)$ and therefore $stg_{inh}(\mathcal{F}) = \{\{\alpha\},\{\alpha,\gamma\},\{\beta\}\}$.

Regarding stg^3_{hyb}, $\{a,a'\}$ defeats $\{\beta\}$ while $\{b\}$ defeats $\{\gamma\}$. Thus, $\{a,a'\}^{\circledast}_{\mathcal{F}} = \{\alpha,\beta\}$, $\{a',c\}^{\circledast}_{\mathcal{F}} = \{\alpha,\gamma\}$, and $\{b\}^{\circledast}_{\mathcal{F}} = \{\beta,\gamma\}$. The three claim-ranges are incomparable, and we have $stg_{hyb}(\mathcal{F}) = \{\{\alpha\},\{\alpha,\gamma\},\{\beta\}\}$. □

Reductions 1, 2, and 4 lose I-maximality for *all* semantics that are I-maximal on wfCAFs but not on general CAFs.

Proposition 4.22. *For $i \in \{1,2,4\}$, the following semantics are not I-maximal for PCAFs, even when considering only transitive preferences: stb^i_{inh}, $stb\text{-}adm^i_{hyb}$, $stb\text{-}cf^i_{hyb}$, sem^i_{inh}, sem^i_{hyb}, prf^i_{inh}, stg^i_{inh}, stg^i_{hyb}.*

Proof. We show this for stb^i_{inh}. For all other σ^i_μ this follows from $stb^i_{inh}(\mathcal{P}) \subseteq \sigma^i_\mu(\mathcal{P})$ (which holds by the properties of general CAFs).

For $i \in \{1,4\}$, let $\mathcal{F}$ be the CAF shown in Figure 4.5b. $\mathcal{F} \in \mathcal{R}_1$-$\mathbf{CAF}_{tr}$ by Proposition 4.12. $\mathcal{F} \in \mathcal{R}_4$-$\mathbf{CAF}_{tr}$ since $\mathcal{R}_4(\mathcal{P}) = \mathcal{F}$ for $\mathcal{P} = (A,R,cl,\succ)$ with $A = \{a,a',b\}$, $R = \{(b,a)\}$, $cl(a) = cl(a') = \alpha$, $cl(b) = \beta$, and $a \succ b$. As required, the underlying CAF of $\mathcal{P}$ is well-formed. It can be verified that $stb(\mathcal{F}) = \{\{a,a'\},\{a',b\}\}$ and thus $stb_{inh}(\mathcal{F}) = \{\{\alpha\},\{\alpha,\beta\}\}$.

For $i = 2$, let $\mathcal{F}'$ be the CAF of Figure 4.5c. $\mathcal{F}' \in \mathcal{R}_2$-$\mathbf{CAF}_{tr}$ since $\mathcal{R}_2(\mathcal{P}') = \mathcal{F}'$ for the PCAF $\mathcal{P}' = (A',R',cl',\succ)$ with $R' = \{(b,a),(b,a'),(b',a),(b',a')\}$, $a \succ b$, and $a' \succ b'$. As required, the underlying CAF of $\mathcal{P}'$ is well-formed. It can be verified that $stb(\mathcal{F}') = \{\{a,a',a''\},\{a'',b,b'\}\}$ and thus $stb_{inh}(\mathcal{F}') = \{\{\alpha\},\{\alpha,\beta\}\}$. □

We can now use the fact that inherited preferred semantics are I-maximal under Reduction 3 to show that inherited and hybrid preferred semantics coincide under Reduction 3.

Proposition 4.23. $prf^3_{inh}(\mathcal{P}) = prf^3_{hyb}(\mathcal{P})$ *for every PCAF* $\mathcal{P}$.

Proof. $prf_{hyb}(\mathcal{F}) \subseteq prf_{inh}(\mathcal{F})$ holds for all CAFs. We must show $prf_{inh}(\mathcal{F}) \subseteq prf_{hyb}(\mathcal{F})$ for all $\mathcal{F} \in \mathcal{R}_3$-**CAF**. Towards a contradiction, assume there is $\mathcal{F} = (A, R, cl) \in \mathcal{R}_3$-**CAF** such that $C \in prf_{inh}(\mathcal{F})$ but $C \notin prf_{hyb}(\mathcal{F})$ for some $C \subseteq cl(A)$. Then $C \in adm_{inh}(\mathcal{F})$. Since $C \notin prf_{hyb}(\mathcal{F})$, there must be $D \in prf_{hyb}(\mathcal{F})$ such that $D \supset C$. Since $prf_{hyb}(\mathcal{F}) \subseteq prf_{inh}(\mathcal{F})$ we have $D \in prf_{inh}(\mathcal{F})$. But then we have $C, D \in prf_{inh}(\mathcal{F})$ and $D \supset C$. This means that prf_{inh} is not I-maximal for CAFs in $\mathcal{R}_3$-**CAF**, which contradicts Proposition 4.20. □

Our results regarding I-maximality also allow us to infer negative results regarding the relationship between semantics: if σ_μ^i is I-maximal while τ_ν^i is not, then there must be a PCAF $\mathcal{P}$ such that $\sigma_\mu^i(\mathcal{P}) \not\subseteq \tau_\nu^i(\mathcal{P})$. Thus, we can conclude:

Proposition 4.24. *For every $i \in \{1, 2, 4\}$ there is:*

- *a PCAF $\mathcal{P}$ such that $prf_{inh}^i(\mathcal{P}) \not\subseteq prf_{hyb}^i(\mathcal{P})$;*
- *a PCAF $\mathcal{P}$ such that $sem_{inh}^i(\mathcal{P}) \not\subseteq prf_{hyb}^i(\mathcal{P})$;*
- *a PCAF $\mathcal{P}$ such that $sem_{hyb}(\mathcal{P}) \not\subseteq prf_{hyb}(\mathcal{P})$.*

Proposition 4.25. *For every $i \in \{1, 2, 3, 4\}$ there is:*

- *a PCAF $\mathcal{P}$ such that $stg_{inh}^i(\mathcal{P}) \not\subseteq naive_{hyb}^i(\mathcal{P})$;*
- *a PCAF $\mathcal{P}$ such that $stg_{hyb}^i(\mathcal{P}) \not\subseteq naive_{hyb}^i(\mathcal{P})$.*

We have now determined the relationship between PCAF-semantics and their properties with respect to I-maximality. In summary:

Theorem 4.26. *The results depicted in Figure 4.6 hold, even when considering only PCAFs with transitive preferences.*

Reduction 3 preserves the properties of wfCAFs for semantics that are based on admissibility (stable, semi-stable, preferred) but not for semantics that are based on conflict-freeness (stage, naive). Reductions 1, 2, and 4 on the other hand lose the I-maximality properties of wfCAFs in all cases (except for those semantics that are I-maximal on general CAFs already). Under Reduction 4 all variants of stable semantics coincide, while under Reduction 2 the inherited and admissibility-based hybrid stable semantics coincide. Reduction 1 preserves none of the investigated semantic properties of wfCAFs.

Intuitively, these results can be explained by the fact that Reduction 3 is the most conservative of the reductions, not adding new attacks and preserving conflict-freeness (i.e., given a PCAF $\mathcal{P}$, a set of arguments E is conflict-free in the underlying CAF of $\mathcal{P}$ iff E is conflict-free in $\mathcal{R}_3(\mathcal{P})$). Reductions 2 and 4 preserve conflict-freeness too, but they

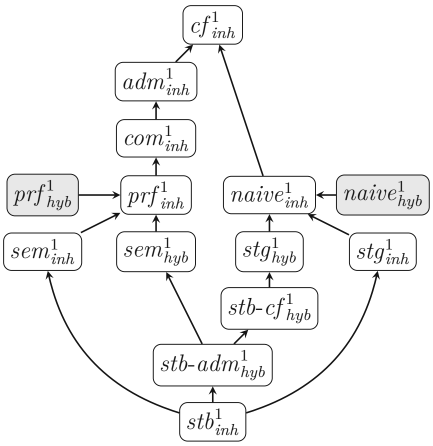

(a) Reduction 1

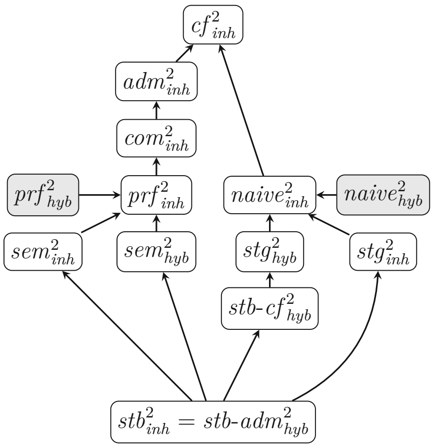

(b) Reduction 2

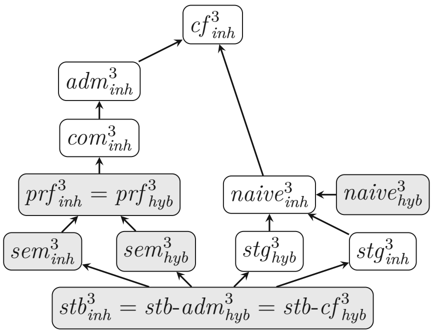

(c) Reduction 3

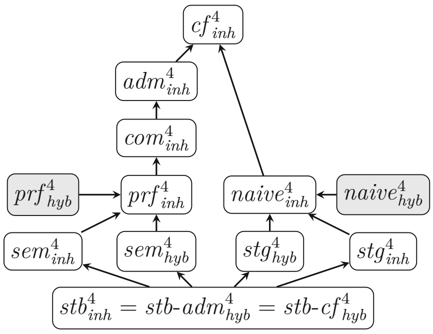

(d) Reduction 4

Figure 4.6: Relations between PCAF-semantics. If there is an arrow from σ to τ, then $\sigma(\mathcal{P}) \subseteq \tau(\mathcal{P})$ for all PCAFs $\mathcal{P}$. Semantics highlighted in gray are I-maximal.

may introduce new attacks in contrast to Reduction 3. Reduction 1 on the other hand does not preserve conflict-freeness. In fact, it has been deemed problematic for exactly this reason when applied to AFs (Amgoud and Vesic 2014), although it is still discussed and considered in the literature alongside the other reductions (Kaci et al. 2021).

Our results support the decision-making process when choosing how preferences should be resolved (i.e., which preference reduction should be used). For example, if Reduction 3 is chosen then no attention has to be paid to the existence of several variants for preferred or stable semantics, since all the variants coincide. What is more, we know that these semantics are I-maximal and therefore behave ‘as expected’ on the claim level. If on the other hand Reduction 1 is chosen, then one must be aware that the different variants for stable and preferred semantics may deliver different extensions, and that none of them (except hybrid preferred semantics) are I-maximal.

4.4 Complexity

In this section, we investigate the impact of preferences on the computational complexity of claim-based reasoning. To this end, we define the three main decision problems for PCAFs analogously to those for CAFs (cf. Definition 2.26), except that we take a PCAF instead of a CAF as input and appeal to PCAF-semantics σ_μ^i instead of CAF-semantics σ_μ.

Definition 4.27 (Decision problems for PCAFs)**.** *We consider the following decision problems pertaining to a PCAF-semantics σ_μ^i:*

- Credulous Acceptance *($Cred_{\sigma_\mu^i}^{PCAF}$): Given a PCAF $\mathcal{P}$ and claim α, is α contained in some $C \in \sigma_\mu^i(\mathcal{P})$?*
- Skeptical Acceptance *($Skept_{\sigma_\mu^i}^{PCAF}$): Given a CAF $\mathcal{P}$ and claim α, is α contained in each $C \in \sigma_\mu^i(\mathcal{P})$?*
- Verification *($Ver_{\sigma_\mu^i}^{PCAF}$): Given a CAF $\mathcal{P}$ and a set of claims C, is $C \in \sigma_\mu^i(\mathcal{P})$?*

Membership results for PCAFs can be inferred from results for general CAFs (recall that the preference reductions from PCAFs to CAFs can be done in polynomial time), and hardness results from results for wfCAFs. Thus, except for $naive_{hyb}^i$, the complexity of credulous and skeptical acceptance follows immediately from known results for CAFs and wfCAFs (cf. Table 2.3):

Observation 4.28. *Let $i \in \{1, 2, 3, 4\}$ and let σ_μ^i be any PCAF-semantics considered in this chapter. $Cred_{\sigma_\mu^i}^{PCAF}$ has the same complexity as $Cred_{\sigma_\mu}^{wfCAF}$. $Skept_{\sigma_\mu^i}^{PCAF}$ has the same complexity as $Skept_{\sigma_\mu}^{wfCAF}$, except for $\sigma_\mu^i = naive_{hyb}^i$.*

The computational complexity of the verification problem, on the other hand, is one level higher on the polynomial hierarchy for general CAFs compared to wfCAFs (cf. Table 2.3), i.e., the bounds that existing results yield for PCAFs are not tight. We address this open problem and comprehensively analyze $Ver_{\sigma_\mu^i}^{PCAF}$ for each of the considered reductions and semantics. Moreover, we investigate the complexity of $Skept_{naive_{hyb}^i}^{PCAF}$.

Regarding conflict-free and naive semantics, the fact that Reductions 2–4 do not remove conflicts straightforwardly implies that the properties of wfCAFs are preserved.

Proposition 4.29. *$Ver_{\sigma_\mu^i}^{PCAF}$ is in P for $\sigma_\mu \in \{cf_{inh}, naive_{inh}, naive_{hyb}\}$ and $i \in \{2, 3, 4\}$.*

Proof. Let $\mathcal{P} = (A, R, cl, \succ)$ be a PCAF, C a set of claims, and $i \in \{2, 3, 4\}$. To check whether $C \in cf_{inh}^i(\mathcal{P})$, by Lemma 4.13, it suffices to check whether $C \in cf_{inh}((A, R, cl))$. This can be done in polynomial time on wfCAFs (cf. Table 2.3). Analogously for $naive_{inh}^i(\mathcal{P})$ and $naive_{hyb}^i(\mathcal{P})$. □

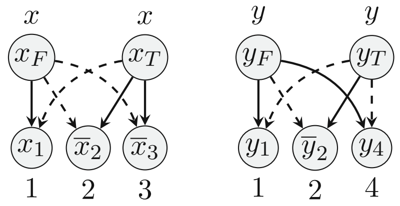

Figure 4.7: Reduction of 3-SAT-instance $\omega_1 = \{x, y\}$, $\omega_2 = \{\neg x, \neg y\}$, $\omega_3 = \{\neg x, z\}$, $\omega_4 = \{y, \neg z\}$, to an instance $(\mathcal{P}, C)$ of $Ver^{PCAF}_{cf^1_{inh}}$. Dashed arrows are attacks deleted in $\mathcal{R}_1(\mathcal{P})$, i.e., they are edges in $wfp(\mathcal{R}_1(\mathcal{P}))$.

Proposition 4.30. *$Skept^{PCAF}_{naive^i_{hyb}}$ is* coNP*-complete for $\sigma^i_\mu = naive^i_{hyb}$ and $i \in \{2, 3, 4\}$.*

Proof. coNP-hardness follows from known results for wfCAFs (see Table 2.3). Regarding coNP-membership, let $\mathcal{P} = (A, R, cl, \succ)$ be a PCAF, $\alpha \in cl(A)$, and $i \in \{2, 3, 4\}$. To decide whether α is skeptically accepted in $\mathcal{P}$ under $naive^i_{hyb}$-semantics, by Lemma 4.13, it suffices to decide whether α is skeptically accepted in the underlying CAF (A, R, cl) of $\mathcal{P}$. This can be done in coNP-time on wfCAFs (cf. Table 2.3). □

4.4.1 Hardness under Reduction 1

Since Reduction 1 *does* remove conflicts between arguments, we cannot apply the same reasoning as above when analyzing the complexity of conflict-free and naive semantics under Reduction 1. Indeed, it turns out we lose the benfits of wfCAFs for these semantics (as well as $stb\text{-}cf_{hyb}$). In the proof of Proposition 4.31 we make use of Reduction 1's ability to remove conflicts in order to show hardness.

Proposition 4.31. *$Ver^{PCAF}_{\sigma^i_\mu}$ is* NP*-hard for $\sigma^i_\mu \in \{cf^1_{inh}, naive^1_{inh}, stb\text{-}cf^1_{hyb}\}$, even if we restrict ourselves to PCAFs with transitive preference relations.*

Proof. Let φ be an arbitrary instance of 3-SAT given as a set $\Omega = \{\omega_1, \dots, \omega_m\}$ of clauses over variables X. Without loss of generality, we can assume that every variable appears both positively and negatively in φ. We construct a PCAF $\mathcal{P} = (A, R, cl, \succ)$ as well as a set of claims C:

- $A = V \cup \overline{V} \cup H$ where $V = \{x_i \mid x \in \omega_i, 1 \leq i \leq m\}$, $\overline{V} = \{\overline{x}_i \mid \neg x \in \omega_i, 1 \leq i \leq m\}$, and $H = \{x_T, x_F \mid x \in X\}$;
- $R = \{(x_T, x_i), (x_F, x_i) \mid x_i \in V\} \cup \{(x_T, \overline{x}_i), (x_F, \overline{x}_i) \mid \overline{x}_i \in \overline{V}\}$;
- $cl(x_i) = cl(\overline{x}_i) = i$ for all $x_i, \overline{x}_i \in V \cup \overline{V}$, $cl(x_T) = cl(x_F) = x$ for all $x \in X$;
- $x_i \succ x_T$ for all $x_i \in V$ and $\overline{x}_i \succ x_F$ for all $\overline{x}_i \in \overline{V}$;
- $C = \{1, \dots, m\} \cup X$.

Figure 4.7 illustrates the above construction. Note that the preferences $x_i \succ x_T$ remove all conflicts between the 'true' variable arguments x_T and their unnegated occurrences x_i. Likewise for preferences of the form $\overline{x}_i \succ x_F$. Now let $\mathcal{F} = \mathcal{R}_1(\mathcal{P}) = (A, R', cl)$. We must show that φ is satisfiable iff $C \in \sigma_\mu(\mathcal{F})$ for $\sigma_\mu \in \{cf_{inh}, naive_{inh}, stb\text{-}cf_{hyb}\}$.

Assume φ is satisfiable. Then there is an interpretation I such that $I \models \varphi$. Let $S = \{x_T \in H \mid x \in I\} \cup \{x_F \in H \mid x \notin I\} \cup \{x_i \in V \mid x \in I\} \cup \{\overline{x}_i \in \overline{V} \mid x \notin I\}$. It can be easily verified that S is conflict free in (A, R') and that $cl(S) = C$. Note that C contains all claims in $\mathcal{F}$, i.e., $C = cl(A)$. Thus, $C \in stb\text{-}cf_{hyb}(\mathcal{F})$. Moreover, $C \in naive_{inh}(\mathcal{F})$ and $C \in cf_{inh}(\mathcal{F})$ since $stb\text{-}cf_{hyb}(\mathcal{F}) \subseteq naive_{inh}(\mathcal{F}) \subseteq cf_{inh}(\mathcal{F})$.

Assume $C \in cf_{inh}(\mathcal{F})$. Then there is some $S \subseteq A$ such that $S \in cf((A, R'))$ and $cl(S) = C$. Let x be any variable in X. Since $x \in cl(S)$ it must be that either $x_T \in S$ or $x_F \in S$. Thus, for all i, j, we have $x_i \in S \implies \overline{x}_j \notin S$ and $\overline{x}_i \in S \implies x_j \notin S$ (otherwise, we would need both $x_T \notin S$ and $x_F \notin S$ for S to be conflict-free). Furthermore, for any $i \in \{1, \ldots, m\}$, there must be some x such that $x_i \in S$ or $\overline{x}_i \in S$. Let $I = \{x \mid x_i \in S \text{ for some } i\}$. Then for every i there is some x such that either $x \in \omega_i$ and $x \in I$ or $\neg x \in \omega_i$ and $x \notin I$. Thus, I satisfies all clauses $\omega_1, \ldots, \omega_m$ which means that φ is satisfiable. The proof works likewise if we assume $C \in naive_{inh}(\mathcal{F})$ or $C \in stb\text{-}cf_{hyb}(\mathcal{F})$ since $stb\text{-}cf_{hyb}(\mathcal{F}) \subseteq naive_{inh}(\mathcal{F}) \subseteq cf_{inh}(\mathcal{F})$. □

Note that the above construction does not work for admissible-based semantics, since the variable-arguments x_i resp. $\overline{x}_i$ in the extension S would remain undefended. The existing hardness proof for general CAFs (Dvořák and Woltran 2020, Proposition 2) cannot be used either, as the constructed CAFs are not in $\mathcal{R}_1$-**CAF**. Specifically, there are symmetric attacks between arguments whose claims occur multiple times, which leads to cycles in the wf-problematic part of the constructed CAF. Instead, we show hardness via a more involved construction in which symmetric attacks are avoided.

Proposition 4.32. *$Ver^{PCAF}_{\sigma^i_\mu}$ is* NP*-hard for $\sigma^i_\mu \in \{stb^1_{inh}, stb\text{-}adm^1_{hyb}, com^1_{inh}, adm^1_{inh}\}$, even if we restrict ourselves to PCAFs with transitive preference relations.*

Proof. Let φ be an arbitrary 3-SAT-instance given as a set $\Omega = \{\omega_1, \ldots, \omega_m\}$ of clauses over variables X. For convenience, we directly construct a CAF $\mathcal{F} = (A, R, cl)$ with $\mathcal{F} \in \mathcal{R}_1\text{-}\mathbf{CAF}_{tr}$ instead of providing a PCAF $\mathcal{P}$ such that $\mathcal{R}_1(\mathcal{P}) = \mathcal{F}$. This is legitimate, as, by our characterization of $\mathcal{R}_1\text{-}\mathbf{CAF}_{tr}$ (see Proposition 4.12), we can obtain $\mathcal{P}$ by simply adding all edges in $wfp(\mathcal{F})$ to R and defining $\succ$ accordingly. We also construct a set of claims C.

- $A = V \cup \overline{V} \cup H$ where $V = \{x_i \mid x \in \omega_i, 1 \leq i \leq m\}$, $\overline{V} = \{\overline{x}_i \mid \neg x \in \omega_i, 1 \leq i \leq m\}$, and $H = \{x^k_{i,j}, \hat{x}^k_{i,j} \mid 1 \leq k \leq 4, x_i \in V, \overline{x}_j \in \overline{V}\}$;
- $R = \{(x_i, x^1_{i,j}), (x^1_{i,j}, x^2_{i,j}), (x^2_{i,j}, \overline{x}_j), (\overline{x}_j, x^3_{i,j}), (x^3_{i,j}, x^4_{i,j}), (x^4_{i,j}, x_i) \mid x_i \in V, \overline{x}_j \in \overline{V}\}$;
- $cl(x_i) = cl(\overline{x}_i) = i$ for all $x_i, \overline{x}_i$ and $cl(x^k_{i,j}) = cl(\hat{x}^k_{i,j}) = x^k_{i,j}$ for all $x^k_{i,j}, \hat{x}^k_{i,j}$;

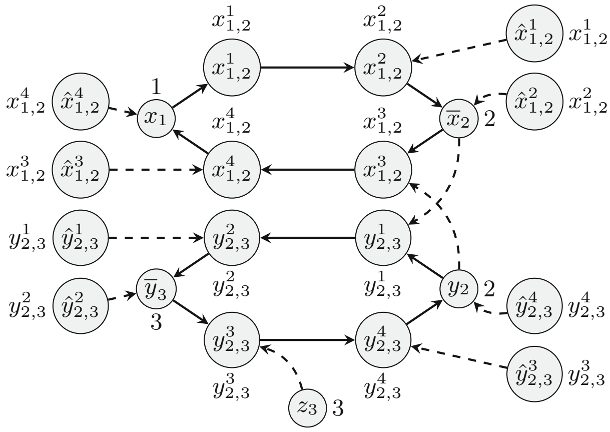

Figure 4.8: Reduction of 3-SAT-instance $\omega_1 = \{x\}$, $\omega_2 = \{\neg x, y\}$, $\omega_3 = \{\neg y, z\}$, to an instance $(\mathcal{P}, C)$ of $Ver^{PCAF}_{stb^1_{inh}}$. Dashed arrows are attacks deleted in $\mathcal{R}_1(F')$, i.e., they are edges in $wfp(\mathcal{R}_1(\mathcal{P}))$.

- $C = \{1, \ldots, m\} \cup \{cl(a) \mid a \in H\}$.

Figure 4.8 illustrates the above construction. In general, every $\hat{x}^k_{i,j}$ only has outgoing edges in the wf-problematic part, and no incoming or outgoing attacks in R. Every $x^k_{i,j}$ only has incoming edges in the wf-problematic part. Finally, there can be no edges in the wf-problematic part between any x_i (or $\overline{x}_i$) and any other x_j (or $\overline{x}_j$). From this, and by the above construction, we can infer that (A, R, cl) fulfills all of the conditions to be in $\mathcal{R}_1$-$\mathbf{CAF}_{tr}$ (cf. Proposition 4.12). It remains to show the correctness of the above construction.

Assume φ is satisfiable. Then there is an interpretation I such that $I \models \varphi$. Let $S = \{x_i \in V \mid x \in I\} \cup \{\overline{x}_i \in \overline{V} \mid x \notin I\} \cup \{x^2_{i,j}, x^3_{i,j} \mid x_i, \overline{x}_j \in A, x \in I\} \cup \{x^1_{i,j}, x^4_{i,j} \mid x_i, \overline{x}_j \in A, x \notin I\} \cup \{\hat{x}^k_{i,j} \mid \hat{x}^k_{i,j} \in A\}$. It can be verified that $S \in stb((A, R))$ and that $cl(S) = C$. Thus, $C \in stb_{inh}(\mathcal{F})$. Moreover, $C \in stb\text{-}adm_{hyb}(\mathcal{F})$, $C \in com_{inh}(\mathcal{F})$, and $C \in adm_{inh}(\mathcal{F})$, since $stb_{inh}(\mathcal{F}) \subseteq stb\text{-}adm_{hyb}(\mathcal{F}) \subseteq com_{inh}(\mathcal{F}) \subseteq adm_{inh}(\mathcal{F})$.

Assume $C \in adm_{inh}(\mathcal{F})$. Then there is some $S \subseteq A$ such that $S \in adm((A, R))$ and $cl(S) = C$. Thus, for any $i \in \{1, \ldots, m\}$, there must be some x such that $x_i \in S$ or $\overline{x}_i \in S$. Consider the case that $x_i \in S$. Since S is admissible, $x^1_{i,j} \notin S$ for any j such that $\overline{x}_j \in A$. This further means that $\overline{x}_j \notin S$ for any $\overline{x}_j \in A$, since we would need $x^1_{i,j} \in S$ to defend $\overline{x}_j$ from the attack by $x^2_{i,j}$. Likewise, if $\overline{x}_i \in S$, then $x_j \notin S$ for all $x_j \in A$. Let $I = \{x \mid x_i \in S \text{ for some } i\}$. Then for every i there is some x such that either $x \in \omega_i$ and $x \in I$ or $\neg x \in \omega_i$ and $x \notin I$. Thus, I satisfies all clauses $\omega_1, \ldots, \omega_m$ which means that φ is satisfiable. The proof works likewise if we assume $C \in stb_{inh}(\mathcal{F})$, $C \in stb\text{-}adm_{hyb}(\mathcal{F})$, or $C \in com_{inh}(\mathcal{F})$, since $stb_{inh}(\mathcal{F}) \subseteq stb\text{-}adm_{hyb}(\mathcal{F}) \subseteq com_{inh}(\mathcal{F}) \subseteq adm_{inh}(\mathcal{F})$. □

As for semi-stable and inherited preferred semantics, we can build upon the standard

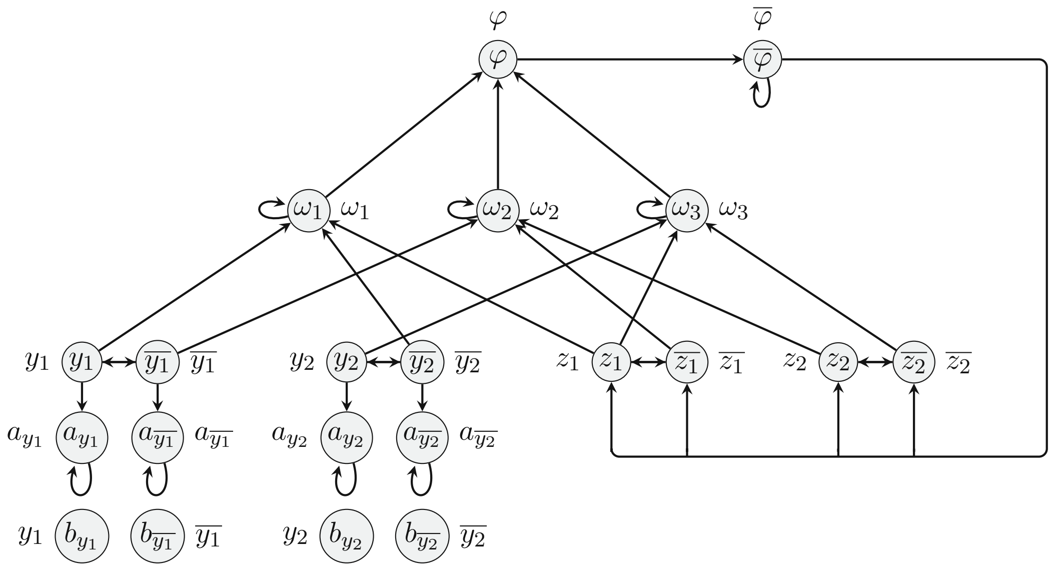

Figure 4.9: Reduction of the $\text{QBF}^2_\forall$ instance $\Phi = \forall y_1, y_2 \exists z_1, z_2 \varphi$ with φ given by clauses $\omega_1 = \{y_1, \neg y_2, z_1\}$, $\omega_2 = \{\neg y_1, \neg z_1, z_2\}$, $\omega_3 = \{y_2, z_1, \neg z_2\}$ to an instance of $Ver^{PCAF}_{sem^1_{inh}}$.

translation for skeptical acceptance of preferred-semantics (Dvořák and Dunne 2018, Reduction 3.7). We introduce helper arguments and avoid symmetric attacks between arguments of the same claim.

Proposition 4.33. *$Ver^{PCAF}_{\sigma^i_\mu}$ is Σ^{P}_2-hard for $\sigma^i_\mu \in \{prf^1_{inh}, sem^1_{inh}, sem^1_{hyb}, stg^1_{inh}, stg^1_{hyb}\}$, even if we restrict ourselves to PCAFs with transitive preference relations.*

Proof. We show hardness for $\sigma_\mu \in \{prf^1_{inh}, sem^1_{inh}, sem^1_{hyb}\}$. The remaining cases can be found in the appendix (Lemma A.4). Let $\Phi = \forall Y \exists Z \varphi$ be an instance of $\text{QBF}^2_\forall$, where φ is given by a set Ω of clauses over atoms $X = Y \cup Z$. We provide a reduction to the complementary problem of $Ver^{PCAF}_{\sigma^1_\mu}$. In particular, we construct the CAF $\mathcal{F} = (A, R, cl)$ with underlying AF $F = (A, R)$ and a set of claims C:

- $A = \{\varphi, \overline{\varphi}\} \cup \Omega \cup X \cup \overline{X} \cup Y_a \cup \overline{Y}_a \cup Y_b \cup \overline{Y}_b$, where $\overline{X} = \{\overline{x} \mid x \in X\}$, $Y_a = \{a_y \mid y \in Y\}$, $\overline{Y}_a = \{a_{\overline{y}} \mid y \in Y\}$, $Y_b = \{b_y \mid y \in Y\}$, $\overline{Y}_b = \{b_{\overline{y}} \mid y \in Y\}$;
- $R = \{(x, \overline{x}), (\overline{x}, x) \mid x \in X\} \cup \{(\omega, \omega), (\omega, \varphi) \mid \omega \in \Omega\} \cup \{(\varphi, \overline{\varphi}), (\overline{\varphi}, \overline{\varphi})\} \cup$ $\{(x, \omega) \mid x \in \omega, \omega \in \Omega\} \cup \{(\overline{x}, \omega) \mid \neg x \in \omega, \omega \in \Omega\} \cup$ $\{(a_v, a_v), (v, a_v) \mid v \in Y \cup \overline{Y}\} \cup \{(\overline{\varphi}, z), (\overline{\varphi}, \overline{z}) \mid z \in Z\}$;
- $cl(b_v) = v$ for $b_v \in Y_b \cup \overline{Y}_b$ and $cl(v) = v$ else;
- $C = Y \cup \overline{Y}$.

Figure 4.9 illustrates the above construction. Note that $\mathcal{F} \in \mathcal{R}_1$-$\mathbf{CAF}_{tr}$ (cf. Proposition 4.12) since all paths in $\mathit{wfp}(\mathcal{F}) = \{(b_a, v) \mid v \in Y \cup \overline{Y}\}$ are of length 1 (only arguments in $Y_b \cup \overline{Y}_b$ have outgoing edges in $\mathit{wfp}(\mathcal{F})$). It remains to verify the correctness of the reduction, i.e., we will show that Φ is valid iff $C \notin \sigma_\mu(\mathcal{F})$.

" $\Longrightarrow$ ": Assume Φ is valid. Consider any $S \subseteq A$ such that $S \in \mathit{adm}(F)$ and $\mathit{cl}(S) = C$. Then $S \subseteq Y \cup \overline{Y} \cup Y_b \cup \overline{Y}_b$. Let $Y' = S \cap Y$. Since Φ is valid, there is $Z' \subseteq Z$ such that $M = Y' \cup Z'$ is a model of φ. Let $T = M \cup \{\overline{x} \mid x \in X \setminus M\} \cup Y_b \cup \overline{Y}_b \cup \{\varphi\}$. Note that $S \subset T$ and $T \in \mathit{cf}(F)$ by construction. Moreover, $T \in \mathit{adm}(F)$ since φ defends $v \in Z' \cup \{\overline{z} \mid z \in Z \setminus Z'\}$ against $\overline{\varphi}$; moreover, each argument $v \in X$ defends itself against $\overline{v}$ and vice verca; also, $M \cup \{\overline{x} \mid x \in X \setminus M\}$ defends φ against each attack from clause-arguments $\omega \in \Omega$ since $M \models \varphi$: for each clause $\omega \in \Omega$, there is either $v \in M$ with $v \in \omega$ or $\neg v \in \omega$ for some $v \notin M$. In the first case, $(v, \omega) \in R$ and $v \in S$, in the latter, $(\overline{v}, \omega) \in R$ and $\overline{v} \in S$. Thus, $S \notin \mathit{prf}(F)$. Since S was chosen as an arbitrary admissible set such that $\mathit{cl}(S) = C$, we can conclude that $C \notin \mathit{prf}_{inh}(\mathcal{F})$. Moreover, $C \notin \mathit{sem}_{inh}(\mathcal{F})$ and $C \notin \mathit{sem}_{hyb}(\mathcal{F})$, since $\mathit{sem}_{inh}(\mathcal{F}) \subseteq \mathit{prf}_{inh}(\mathcal{F})$ and $\mathit{sem}_{hyb}(\mathcal{F}) \subseteq \mathit{prf}_{inh}(\mathcal{F})$.

" $\Longleftarrow$ ": Assume $C \notin \mathit{sem}_{inh}(\mathcal{F})$ (resp. $C \notin \mathit{sem}_{hyb}(\mathcal{F})$). Consider any $Y' \subseteq Y$. We will show that there is some $Z' \subseteq Z$ such that $Y' \cup Z' \models \varphi$. Let $S = Y' \cup \{\overline{y} \mid y \in Y \setminus Y'\} \cup Y_b \cup \overline{Y}_b$. Observe that $\mathit{cl}(S) = C$ and $S \in \mathit{adm}(F)$. Since $C \notin \mathit{sem}_{inh}(\mathcal{F})$ (resp. $C \notin \mathit{sem}_{hyb}(\mathcal{F})$), there is $T \in \mathit{adm}(F)$ with $T \cup T_F^+ \supset S \cup S_F^+$ (resp. $\mathit{cl}(T) \cup T_\mathcal{F}^* \supset \mathit{cl}(S) \cup S_\mathcal{F}^*$).

In particular, we have $Y' \cup \{\overline{y} \mid y \in Y \setminus Y'\} \subseteq T$ since each $a_v \in S_F^+$ (resp. $a_v \in S_\mathcal{F}^*$) with $v \in Y \cup \overline{Y}$ has precisely one non-self-attacking attacker (namely the argument v). Moreover, we can assume that T contains each argument $v \in Y_b \cup \overline{Y}_b$ since each such v is unattacked and does not attack any other argument. We can conclude that $T \supset S$.

It follows that $\varphi \in T$: since $S \subset T$, there is some $v \in A \setminus S$ such that $v \in T$. Clearly, $v \in \{\varphi\} \cup Z \cup \overline{Z}$ since each remaining argument is either self-attacking or attacked by S (and thus also by T). In case $v = \varphi$, we are done; in case $v \in Z \cup \overline{Z}$, we have $\varphi \in T$ by admissibility of T (observe that φ is the only attacker of $\overline{\varphi}$). Consequently, T defends φ against each attack from each clause-argument $\omega \in \Omega$.

Now, let $Z' = Z \cap T$. We show that $M = Y' \cup Z'$ is a model of φ. Consider some arbitrary clause $\omega \in \Omega$. Since $\varphi \in T$, there is some $v \in T$ such that $(v, \omega) \in R$ by admissibility of T. In case $v \in X$, we have $v \in M$ and $v \in \omega$ by construction of $\mathcal{F}$; similarly, in case $v \in \overline{X}$ we have $v \notin M$ and $\neg v \in \omega$. Thus, ω is satisfied by M. Since ω was chosen arbitrarily it follows that $M \models \varphi$. We can conclude that Φ is valid. The proof works likewise if we assume $C \notin \mathit{prf}_{inh}(\mathcal{F})$ since $\mathit{sem}_{inh}(\mathcal{F}) \subseteq \mathit{prf}_{inh}(\mathcal{F})$ (resp. $\mathit{sem}_{hyb}(\mathcal{F}) \subseteq \mathit{prf}_{inh}(\mathcal{F})$). □

Just like inherited preferred/naive semantics, hybrid preferred/naive semantics are DP-complete and thus preserve the high complexity of general CAFs. To show this for prf^1_{hyb}, we adapt an existing reduction from SatUnsat to general CAFs (Dvořák et al. 2023).

Proposition 4.34. *$\mathit{Ver}^{PCAF}_{\sigma^i_\mu}$ is DP-hard for $\sigma^i_\mu \in \{\mathit{prf}^1_{hyb}, \mathit{naive}^1_{hyb}\}$, even if we restrict ourselves to PCAFs with transitive preference relations.*

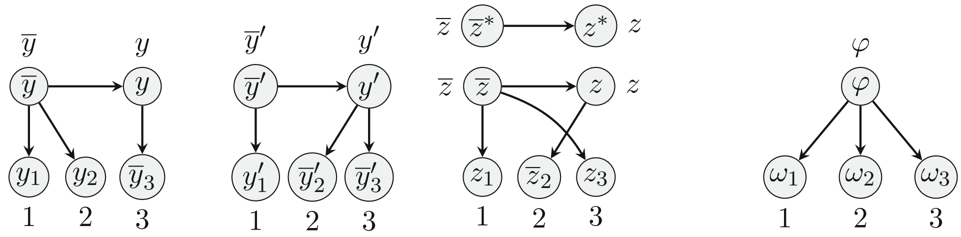

Figure 4.10: Reduction of the $\text{QBF}_\forall^2$ instance $\Phi = \forall y, y'\exists z\varphi$ with φ given by $\omega_1 = \{y, y', z\}$, $\omega_2 = \{y, \neg y', \neg z\}$, $\omega_3 = \{y, \neg y', z\}$, to an instance of $Skept^{PCAF}_{naive^1_{hyb}}$.

Proof. We show hardness for $\sigma_\mu = prf^1_{hyb}$. The proof for $\sigma_\mu = naive^1_{hyb}$ can be found in the appendix (Lemma A.5). Let (φ_1, φ_2) be an arbitrary instance of SatUnsat, where φ_i is given over a set of clauses Ω_i and a set of variables X_i such that $X_1 \cap X_2 = \emptyset$. Given X_i, we define $\overline{X_i} = \{\overline{x} \mid x \in X_i\}$. Instead of constructing a PCAF, we directly construct a CAF $\mathcal{F} = (A, R, cl) \in \mathcal{R}_1\text{-}\mathbf{CAF}_{tr}$:

- $A = A_1 \cup A_2$, where $A_i = X_i \cup \overline{X_i} \cup \Omega_i \cup \{\varphi_i\} \cup \{d_x \mid x \in X_i \cup \overline{X_i}\}$;
- $R = R_1 \cup R_2$, where $R_i = \{(x, \omega) \mid \omega \in \Omega_i, x \in \Omega_i\} \cup \{(\overline{x}, \omega) \mid \omega \in \Omega_i, \neg x \in \Omega_i\} \cup \{(x, \overline{x}), (\overline{x}, x) \mid x \in X_i\} \cup \{(\omega, \varphi_i), (\omega, \omega) \mid \omega \in \Omega_i\}$;
- $cl(d_x) = x$ for $x \in X_i \cup \overline{X_i}$, $cl(x) = x$ for all other arguments in A.

The claim-set to be verified is $C = X_1 \cup \overline{X_1} \cup X_2 \cup \overline{X_2} \cup \{\varphi_1\}$. Observe that $C = cl(A) \setminus (\Omega_1 \cup \Omega_2 \cup \{\varphi_2\})$, i.e., all claims except those of clauses and φ_2 are contained in C. We show that (φ_1, φ_2) is a yes-instance of SatUnsat if and only if $C \in prf_{hyb}(\mathcal{F})$:

Assume (φ_1, φ_2) is a yes-instance of SatUnsat. Then there is an interpretation I such that $I \models \varphi_1$, but there is no interpretation that satisfies φ_2. Thus, φ_2 cannot be part of any admissible extension, since it must be defended against all clause arguments from Ω_2. Let $S = I \cup \{\overline{x} \mid x \in (X_1 \cup X_2) \setminus I\} \cup \{d_x \mid x \in X_1 \cup X_2\} \cup \{\varphi_1\}$. It can be verified that $cl(S) = C$, $S \in adm((A, R))$, and that there is no $S' \in adm((A, R))$ with $cl(S') \supset cl(S)$.

Assume (φ_1, φ_2) is a no-instance of SatUnsat. There are two cases: (1) φ_1 is unsatisfiable. Then φ_1 cannot be part of any admissible extension, i.e., $C \notin adm_{inh}(\mathcal{F})$. (2) Both φ_1 and φ_2 are satisfiable. Since φ_1 and φ_2 share no variables, there is an interpretation I such that $I \models \varphi_1$ and $I \models \varphi_2$. Let $S = I \cup \{\overline{x} \mid x \in (X_1 \cup X_2) \setminus I\} \cup \{d_x \mid x \in X_1 \cup X_2\} \cup \{\varphi_1, \varphi_2\}$. Note that $cl(S) \supset C$ and $S \in adm((A, R))$. Thus, $C \notin prf_{hyb}(\mathcal{F})$. □

It only remains to investigate skeptical acceptance for $naive^1_{hyb}$, which, as we show, also preserves the higher complexity of general CAFs. This means that Reduction 1 loses the computational benefits of wfCAFs for all semantics considered in this chapter.

Proposition 4.35. *$Skept^{PCAF}_{\sigma^i_\mu}$ is Π_2^{P}-hard for $\sigma^i_\mu = naive^1_{hyb}$, even if we restrict ourselves to PCAFs with transitive preference relations.*

Proof. Let $\Phi = \forall Y \exists Z \varphi$ be an instance of $\text{QBF}^2_\forall$, where φ is given by a set $\Omega = \{\omega_1, \ldots, \omega_m\}$ of clauses over atoms $X = Y \cup Z$. We construct $\mathcal{F} = (A, R, cl)$ with

- $A = \{\varphi\} \cup \Omega \cup \{x, \overline{x} \mid x \in X\} \cup \{z^*, \overline{z}^* \mid z \in Z\} \cup$
  $\{x_i \mid x \in X, x \in \omega_i, 1 \le i \le m\} \cup \{\overline{x}_i \mid x \in X, \neg x \in \omega_i, 1 \le i \le m\}$;
- $R = \{(\varphi, \omega) \mid \omega \in \Omega\} \cup \{(\overline{x}, x) \mid x \in X\} \cup \{(\overline{z}^*, z^*) \mid z \in Z\} \cup$
  $\{(x, \overline{x_i}) \mid x \in X, \neg x \in \omega_i, 1 \le i \le m\} \cup \{(\overline{x}, x_i) \mid x \in X, x \in \omega_i, 1 \le i \le m\}$;
- $cl(x_i) = cl(\overline{x}_i) = cl(\omega_i) = i$ for $1 \le i \le m$,
  $cl(z^*) = z$, $cl(\overline{z}^*) = cl(\overline{z})$, and
  $cl(v) = cl(v)$ for all other $v \in A$.

Figure 4.10 illustrates the above construction. Note that $\mathcal{F} \in \mathcal{R}_1\text{-}\mathbf{CAF}_{tr}$ (cf. Proposition 4.12): the only arguments with both incoming and outgoing edges in $\mathit{wfp}(\mathcal{F})$ are the z^* arguments, with $z \in Z$. The edge leading to z^* comes from $\overline{z}$, and the edge going out of z^* leads to some $\overline{z}_i$. Thus, there is no cycle in $\mathit{wfp}(\mathcal{F})$. Moreover, the only path in $\mathit{wfp}(\mathcal{F})$ with more than one edge is from $\overline{z}$ to some $\overline{z}_i$, while $(\overline{z}, \overline{z}_i) \notin R$. It remains to verify the correctness of the reduction: we show that Φ is valid iff $\varphi \in C$ for all $C \in \mathit{naive}_{hyb}(\mathcal{F})$.

"$\Longrightarrow$": Assume Φ is valid. Let $C \in \mathit{naive}_{hyb}(\mathcal{F})$. Note that, for each $y \in Y$, we cannot have $y \in C$ and $\overline{y} \in C$ at the same time. Consider the argument set $S_Y = (Y \cap C) \cup \{\overline{y} \mid y \in Y \setminus C\}$. Note that $cl(S_Y) \supseteq (C \cap \{y, \overline{y} \mid y \in Y\})$. Let $Y' = S_Y \cap Y$. Since Φ is valid there is $Z' \subseteq Z$ such that $I \models \varphi$ for $I = Y' \cup Z'$. Let $S_Z = \{z, \overline{z}^* \mid z \in Z'\} \cup \{\overline{z}, z^* \mid z \in Z \setminus Z'\}$. Note that $cl(S_Z) = \{z, \overline{z} \mid z \in Z\} \supseteq (C \cap \{z, \overline{z} \mid z \in Z\})$. Now let $S_X = S_Y \cup S_Z$ and finally $S = S_X \cup \{x_i \mid x \in S_X, x_i \in A\} \cup \{\overline{x}_i \mid \overline{x} \in S_X, \overline{x}_i \in A\} \cup \{\varphi\}$. Note that $S \in \mathit{cf}((A, R))$ by construction. Moreover, since I satisfies all clauses in Ω we have $cl(S) \supseteq \{1, \ldots, m\} \supseteq (C \cap \{1, \ldots, m\})$. Since also $\varphi \in cl(S)$ we can conclude that $cl(S) \supseteq C$. But $C \in \mathit{naive}_{hyb}(\mathcal{F})$, i.e., it cannot be that $cl(S) \supset C$. Thus, $cl(S) = C$ and therefore $\varphi \in C$.

"$\Longleftarrow$": Assume Φ is not valid. Then there is $Y' \subseteq Y$ such that for all $Z' \subseteq Z$ we have $Y' \cup Z' \not\models \varphi$. Let $S = Y' \cup \{\overline{y} \mid y \in Y \setminus Y'\} \cup \{z, \overline{z}^* \mid z \in Z\} \cup \{\omega_1, \ldots, \omega_m\}$. Towards a contradiction, assume there is $T \in \mathit{cf}((A, R))$ such that $cl(T) \supset cl(S)$. Since for every $y \in Y$ we already have $y \in S$ or $\overline{y} \in S$, and since y and $\overline{y}$ are in conflict, we have $y \in T$ iff $y \in S$ and $\overline{y} \in T$ iff $\overline{y} \in S$. Moreover, since $(\{z, \overline{z}^* \mid z \in Z\} \cup \{1, \ldots, m\}) \subseteq cl(S) \subset cl(T)$ it must be that $\varphi \in T$. This further implies that $\omega_i \notin T$ for all $\omega_i \in \Omega$. Note that for every $z \in Z$ we must have $z, \overline{z}^* \in T$ or $\overline{z}, z^* \in T$ since $\{z, \overline{z}^* \mid z \in Z\} \subset cl(T)$. Let $Z' = T \cap Z$. Since $\{1, \ldots, m\} \subset cl(T)$, we can infer that every clause $\omega_i \in \Omega$ is satisfied by $Y' \cup Z'$, i.e., $Y' \cup Z' \models \varphi$. Contradiction. □

4.4.2 Efficient Algorithms for Reductions 2–4

We have already seen that the computational benefits of wfCAFs are preserved when using Reductions 2–4 and considering conflict-free/naive semantics (cf. Propositions 4.29

and 4.30). In this subsection we show that the benefits of wfCAFs are in fact retained under Reductions 2–4 for the vast majority of admissible-based semantics, with the only exception being complete semantics under Reductions 2 and 4. To do so, we require a more involved argument than in the case of conflict-free/naive semantics, since Reductions 2–4 may very well cause certain arguments to be undefended. Consider for example the PCAF $\mathcal{P} = (A, R, cl, \succ)$ with two arguments $A = \{x, y\}$, attacks $R = \{(x, y), (y, x)\}$, claims $cl(x) = x$ and $cl(y) = y$, and the preference $x \succ y$. The preferred claim-extensions before resolving preferences are $prf_{inh}((A, R, \mathcal{L})) = \{\{x\}, \{y\}\}$ while the only preferred claim-extension after resolving preferences is $prf_{inh}(\mathcal{R}_i(\mathcal{P})) = \{\{x\}\}$.

Given a wfCAF $\mathcal{F}$ and a set of claims C, a set of arguments S can be constructed in polynomial time such that S is the unique maximal admissible set in $\mathcal{F}$ with claim $cl(S) = C$ (Dvořák and Woltran 2020). Making use of the fact that Reductions 2–4 do not alter conflicts between arguments, we can construct such a maximal set of arguments also for PCAFs: given a PCAF $\mathcal{P}$ and set C of claims, we define the set $E_0(C)$ containing all arguments of $\mathcal{P}$ with a claim in C; the set $E^i_1(C)$ is obtained from $E_0(C)$ by removing all arguments attacked by $E_0(C)$ in the underlying CAF of $\mathcal{P}$; finally, the set $E^i_*(C)$ is obtained by repeatedly removing all arguments not defended by $E^i_1(C)$ in $\mathcal{R}_i(\mathcal{P})$ until a fixed point is reached.

Definition 4.36. *Given a PCAF $\mathcal{P} = (A, R, cl, \succ)$, a set of claims C, and $i \in \{2, 3, 4\}$, let*

$$
\begin{aligned}
E_0(C) =& \{a \in A \mid cl(a) \in C\};\\
E^i_1(C) =& E_0(C) \setminus E_0(C)^+_{(A,R)};\\
E^i_k(C) =& \{x \in E^i_{k-1}(C) \mid x \textit{ is defended by } E^i_{k-1}(C) \textit{ in } \mathcal{R}_i(\mathcal{P})\} \textit{ for } k \geq 2;\\
E^i_*(C) =& E^i_k \textit{ for } k \geq 2 \textit{ such that } E^i_k(C) = E^i_{k-1}(C).
\end{aligned}
$$

The above definition is based on (Dvořák and Woltran 2020, Definition 5), but with the crucial differences that undefended arguments are (i) computed w.r.t. $\mathcal{R}_i(\mathcal{P})$ and (ii) are iteratively removed until a fixed point is reached.

Lemma 4.37. *Let $\mathcal{P}$ be a PCAF, C a set of claims, and $i \in \{2, 3, 4\}$. The following holds:*

- *$C \in cf^i_{inh}(\mathcal{P})$ iff $cl(E^i_1(C)) = C$. Moreover, if $C \in cf^i_{inh}(\mathcal{P})$ then $E^i_1(C)$ is the unique maximal conflict-free set S in $\mathcal{R}_i(\mathcal{P})$ such that $cl(S) = C$;*
- *$C \in adm^i_{inh}(\mathcal{P})$ iff $cl(E^i_*(C)) = C$. If $C \in adm^i_{inh}(\mathcal{P})$ then $E^i_*(C)$ is the unique maximal admissible set S in $\mathcal{R}_i(\mathcal{P})$ such that $cl(S) = C$.*

Proof. We consider the two statements separately:

- Conflict-freeness: let $\mathcal{P} = (A, R, cl, \succ)$ be a PCAF. From (Dvořák and Woltran 2020, Lemma 1) we know that $C \in cf_{inh}((A, R, cl))$ iff $cl(E_1^i(C)) = C$, as well as that, if $C \in cf_{inh}((A, R, cl))$ then $E_1^i(C)$ is the unique maximal conflict-free set S in (A, R, cl) with $cl(S) = C$. From this and our Lemma 4.13, our result follows immediately.

- Admissibility: let $\mathcal{P} = (A, R, cl, \succ)$ be a PCAF, C a set of claims, and $i \in \{2, 3, 4\}$.

  Assume $cl(E_*^i(C)) = C$. By construction, $E_*^i(C) \in adm(\mathcal{R}_i(\mathcal{P}))$, and thus $C \in adm_{inh}^i(\mathcal{P})$.

  Now assume $C \in adm_{inh}^i(\mathcal{P})$. Then there exists $S \subseteq A$ such that $cl(S) = C$ and $S \in adm(\mathcal{R}_i(\mathcal{P}))$. Furthermore, $C \in cf_{inh}^i(\mathcal{P})$ and therefore $S \subseteq E_1^i(C)$. By construction, $E_*^i(C) \subseteq E_1^i(C)$. Moreover, any $x \in S$ is defended by S in $\mathcal{R}_i(\mathcal{P})$ and therefore also by $E_1^i(C)$. Thus, by definition, $x \in E_2^i(C)$. By the same argument, if $x \in S$ and $x \in E_k^i(C)$ then $x \in E_{k+1}^i(C)$. We can conclude that $S \subseteq E_*^i(C) \subseteq E_1^i(C)$ and thus $cl(E_*^i(C)) = C$. By the above we have that $E_*^i(C)$ is admissible and each $S \subseteq A$ such that $cl(S) = C$ is a subset of $E_*^i(C)$. In other words $E_*^i(C)$ is the unique maximal admissible set S in $\mathcal{R}_i(\mathcal{P})$ such that $cl(S) = C$. □

By computing the maximal conflict-free (resp. admissible) extensions $E_1^i(C)$ (resp. $E_*^i(C)$) for a claim set C, verification becomes easier for most semantics.

Proposition 4.38. *$Ver_{\sigma_\mu^i}^{PCAF}$ is in* P *for $\sigma_\mu \in \{adm_{inh}, stb_{inh}, stb\text{-}adm_{hyb}, stb\text{-}cf_{hyb}\}$ and $i \in \{2, 3, 4\}$, as well as for $\sigma_\mu^i = com_{inh}^3$.*

Proof. Let $\mathcal{P} = (A, R, cl, \succ)$ be a PCAF, let C be a set of claims, and let $i \in \{2, 3, 4\}$. Our goal is to verify that $C \in \sigma_\mu^i(\mathcal{P})$. Note that we can compute $\mathcal{R}_i(\mathcal{P}) = (A, R', cl)$, $E_1^i(C)$, and $E_*^i(C)$ in polynomial time.

- $\sigma_\mu^i = adm_{inh}^i$: by Lemma 4.37, it suffices to test whether $cl(E_*^i(C)) = C$.

- $\sigma_\mu^i \in \{stb_{inh}^i, stb\text{-}adm_{hyb}^i\}$: note that $stb_{inh}^i(\mathcal{P}) = stb\text{-}adm_{hyb}^i(\mathcal{P})$ under Reductions 2,3,4 (cf. Proposition 4.14), i.e., we must only verify that $C \in stb_{inh}^i(\mathcal{P})$. We first check whether $C \in adm_{inh}^i(\mathcal{P})$. If not, $C \notin stb_{inh}^i(\mathcal{P})$. If yes, then $cl(E_*^i(C)) = C$ by Lemma 4.37. We can check in polynomial time if $E_*^i(C) \in stb((A, R'))$. If yes, we are done. If no, then there is an argument x that is not in $E_*^i(C)$ but is also not attacked by $E_*^i(C)$ in $\mathcal{R}_i(\mathcal{P})$. Moreover, there can be no other $S \in stb((A, R'))$ with $cl(S) = C$ since for any such S we would have $S \subseteq E_*^i(C)$ by Lemma 4.37, which would imply that S does not attack x and that $x \notin S$.

- $\sigma_\mu^i = stb\text{-}cf_{hyb}^i$: we first check whether $C \in cf_{inh}^i(\mathcal{P})$. If not, $C \notin stb\text{-}cf_{hyb}^i(\mathcal{P})$. If yes, then, by Lemma 4.37, $cl(E_1^i(C)) = C$. We can check in polynomial time if $E_1^i(C)^{\circledast}_{(A,R')} = cl(A)$. If yes, then $C \in stb\text{-}cf_{hyb}(\mathcal{R}_i(\mathcal{P}))$ and we are done. If no, then there is an argument x such that $x \notin E_1^i(C)$, $cl(x) \notin C$, and x is not attacked by

$E_1^i(C)$ in $\mathcal{R}_i(\mathcal{P})$. Moreover, there can be no other $S \in cf((A, R'))$ with $cl(S) = C$ and $S^{\circledast}_{(A,R')} = cl(A)$ since for any such S we would have $S \subseteq E_1^i(C)$ by Lemma 4.37, which would imply that S does not attack x.

- $\sigma_\mu^i = com^3_{inh}$: we first check if $C \in adm^3_{inh}(\mathcal{P})$. If not, $C \notin com^3_{inh}(\mathcal{P})$. If yes, then $cl(E^3_*(C)) = C$. We can check in polynomial time if $E^3_*(C) \in com(\mathcal{R}_3(\mathcal{P}))$. If no, then $E^3_*(C)$ defends some $x \notin E^3_*(C)$ in $\mathcal{R}_3(\mathcal{P})$. Towards a contradiction, assume there is some $S \subseteq A$ such that $S \in com(\mathcal{R}_3(\mathcal{P}))$ and $cl(S) = C$. By Lemma 4.37, $S \subseteq E^3_*(C)$, which implies $x \notin S$. Then S cannot defend x in $\mathcal{R}_3(\mathcal{P})$, i.e., there must be y and z such that $y \in E^3_*(C)$, $y \notin S$, $(z, x) \in \mathcal{R}_3(\mathcal{P})$, and $(y, z) \in \mathcal{R}_3(\mathcal{P})$. Then also $(y, z) \in \mathcal{P}$ by the definition of Reduction 3. But there must also be some $y' \in S$ with $cl(y') = cl(y)$, and since the underlying CAF of $\mathcal{P}$ is well-formed there must be $(y', z) \in \mathcal{P}$. Since there cannot be $(y', z) \in \mathcal{R}_3(\mathcal{P})$, otherwise S would defend x, it has to be that $z \succ y'$. For Reduction 3 this further requires $(z, y') \in \mathcal{P}$. Crucially, $(z, y') \in \mathcal{R}_3(\mathcal{P})$. But then S must be defended from z, i.e., there must be some $w \in S$ such that $(w, z) \in \mathcal{R}_3(\mathcal{P})$. But this means that S defends x, i.e., S is not complete. Contradiction. □

Proposition 4.39. *$Ver^{PCAF}_{\sigma_\mu^i}$ is in* coNP *for $\sigma_\mu \in \{prf_{inh}, prf_{hyb}, sem_{inh}, sem_{hyb}, stg_{inh}, stg_{hyb}\}$ and $i \in \{2, 3, 4\}$.*

Proof. We show that the complementary problem is in NP. Let $\mathcal{P} = (A, R, cl, \succ)$ be a PCAF with $\mathcal{R}_i(\mathcal{P}) = (A, R', cl)$ for $i \in \{2, 3, 4\}$. Let $C \subseteq cl(A)$ be a set of claims. Our algorithm must verify that $C \notin \sigma^i_\mu(\mathcal{P})$ in NP-time. Note that the argument-sets $E^i_1(C)$ and $E^i_*(C)$ can be computed in polynomial time with respect to $\mathcal{P}$ (cf. Definition 4.36).

- $\sigma^i_\mu \in \{prf^i_{inh}, prf^i_{hyb}, sem^i_{inh}, sem^i_{hyb}\}$: first, guess a set of claims $D \subseteq cl(A)$. Then, check whether $cl(E^i_*(C)) = C$. If no, then, by Lemma 4.37, $C \notin adm_{inh}(\mathcal{R}_i(\mathcal{P}))$ and we are done. If yes, we proceed differently depending on which semantics we consider:
  - $\sigma^i_\mu = prf^i_{inh}$: verify that $D \in adm^i_{inh}(\mathcal{P})$ and $E^i_*(C) \subset E^i_*(D)$. Since $E^i_*(C)$ is the unique maximal admissible set in $\mathcal{R}_i(\mathcal{P})$ with claim C (cf. Lemma 4.37), we have $S \subseteq E^i_*(C) \subset E^i_*(D)$ for every $S \in adm((A, R'))$ with $cl(S) = C$. Hence, $C \notin prf^i_{inh}(\mathcal{P})$.
  - $\sigma^i_\mu = prf^i_{hyb}$: verify that $D \in adm^i_{inh}(\mathcal{P})$ and $C \subset D$. Then $C \notin prf^i_{hyb}(\mathcal{P})$.
  - $\sigma^i_\mu = sem^i_{inh}$: verify that $D \in adm^i_{inh}(\mathcal{P})$ and $E^i_*(C)^{\oplus}_{(A,R')} \subset E^i_*(D)^{\oplus}_{(A,R')}$. As above, we have $S \subseteq E^i_*(C)$ and therefore also $S^{\oplus}_{(A,R')} \subseteq E^i_*(C)^{\oplus}_{(A,R')} \subset E^i_*(D)^{\oplus}_{(A,R')}$ for every $S \in adm((A, R'))$ with $cl(S) = C$. Hence, $C \notin sem^i_{inh}(\mathcal{P})$.
  - $\sigma^i_\mu = sem^i_{hyb}$: verify that $D \in adm^i_{inh}(\mathcal{P})$ and $E^i_*(C)^{\circledast}_{\mathcal{R}_i(P)} \subset E^i_*(D)^{\circledast}_{\mathcal{R}_i(P)}$. As above, we have $S \subseteq E^i_*(C)$ and therefore also $S^{\circledast}_{\mathcal{R}_i(P)} \subseteq E^i_*(C)^{\circledast}_{\mathcal{R}_i(P)} \subset$

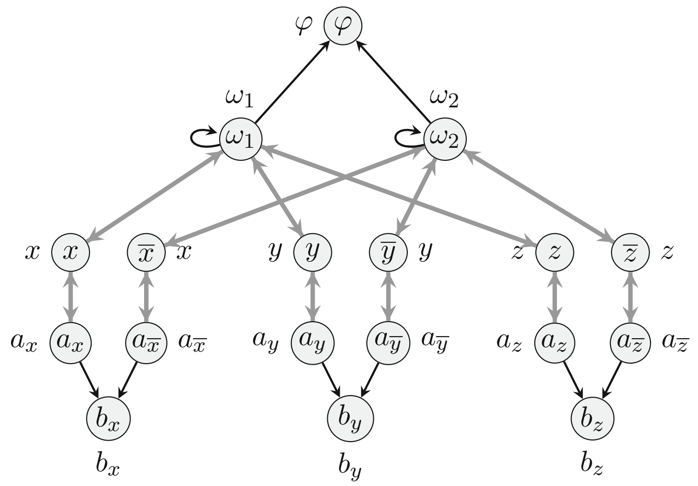

Figure 4.11: $\mathcal{R}_4(\mathcal{P})$ from the proof of Proposition 4.40, with φ given by clauses $\omega_1 = \{x, y, z\}$, $\omega_2 = \{\neg x, \neg y, \neg z\}$. Symmetric attacks (gray/thick) have been introduced by Reduction 4.

$E^i_*(D)^{\circledast}_{\mathcal{R}_i(P)}$ for every $S \in adm((A, R'))$ with $cl(S) = C$. Hence, $C \notin sem^i_{hyb}(\mathcal{P})$.

- $\sigma^i_\mu \in \{stg^i_{inh}, stg^i_{hyb}\}$: first, guess a set of claims $D \subseteq cl(A)$. Then, check whether $cl(E^i_1(C)) = C$. If no, then, by Lemma 4.37, $C \notin cf_{inh}(\mathcal{R}_i(\mathcal{P}))$ and we are done. If yes, we proceed differently depending on which semantics we consider:
  - $\sigma^i_\mu = stg^i_{inh}$: verify that $D \in cf^i_{inh}(\mathcal{P})$ and $E^i_1(C)^{\oplus}_{(A,R')} \subset E^i_1(D)^{\oplus}_{(A,R')}$. Since $E^i_1(C)$ is the unique maximal conflict-free set in $\mathcal{R}_i(\mathcal{P})$ with claim C (cf. Lemma 4.37), we have $S \subseteq E^i_1(C)$ and therefore also $S^{\oplus}_{(A,R')} \subseteq E^i_1(C)^{\oplus}_{(A,R')} \subset E^i_1(D)^{\oplus}_{(A,R')}$ for every $S \in cf((A, R'))$ with $cl(S) = C$. Hence, $C \notin stg^i_{inh}(\mathcal{P})$.
  - $\sigma^i_\mu = stg^i_{hyb}$: verify that $D \in cf^i_{inh}(\mathcal{P})$ and $E^i_1(C)^{\circledast}_{\mathcal{R}_i(P)} \subset E^i_1(D)^{\circledast}_{\mathcal{R}_i(P)}$. As above, we have $S \subseteq E^i_1(C)$ and therefore also $S^{\circledast}_{\mathcal{R}_i(P)} \subseteq E^i_1(C)^{\circledast}_{\mathcal{R}_i(P)} \subset E^i_1(D)^{\circledast}_{\mathcal{R}_i(P)}$ for every $S \in cf((A, R'))$ with $cl(S) = C$. Hence, $C \notin stg^i_{hyb}(\mathcal{P})$. □

For complete semantics, only Reduction 3 retains the benefits of wfCAFs. Here, the fact that Reductions 2 and 4 can introduce new attacks leads to an increase in complexity.

Proposition 4.40. *$Ver^{PCAF}_{\sigma^i_\mu}$ is* NP*-hard for $\sigma_\mu = com_{inh}$ and $i \in \{2, 4\}$, even if we restrict ourselves to PCAFs with transitive preference relations.*

Proof. We show NP-hardness for $\sigma^i_\mu = com^4_{inh}$. The proof for $\sigma^i_\mu = com^2_{inh}$ is similar and can be found in the appendix (Lemma A.6).

Let φ be an arbitrary instance of 3-SAT given as a set Ω of clauses over variables X and let $\overline{X} = \{\overline{x} \mid x \in X\}$. We construct a PCAF $\mathcal{P} = (A, R, cl, \succ)$ as well as a set of claims C:

- $A = \{\varphi\} \cup \Omega \cup X \cup \overline{X} \cup \{a_x \mid x \in X \cup \overline{X}\} \cup \{b_x \mid x \in X\}$;
- $R = \{(\omega, \varphi) \mid \omega \in \Omega\} \cup \{(\omega, \omega) \mid \omega \in \Omega\} \cup$
  $\{(\omega, x) \mid x \in \omega, \omega \in \Omega\} \cup \{(\omega, \overline{x}) \mid \neg x \in \omega, \omega \in \Omega\} \cup$
  $\{(a_x, x) \mid x \in X \cup \overline{X}\} \cup \{(a_x, b_x), (a_{\overline{x}}, b_x) \mid x \in X\}$;
- $cl(x) = cl(\overline{x}) = x$ for $x \in X$, $cl(v) = v$ otherwise;
- $x \succ \omega$, $x \succ a_x$ for all $x \in X \cup \overline{X}$ and all $\omega \in \Omega$;
- $C = X \cup \{\varphi\}$.

Figure 4.11 illustrates the above construction. It remains to show that φ is satisfiable if and only if $C \in com_{inh}(\mathcal{R}_4(\mathcal{P}))$.

Assume φ is satisfiable. Then there is an interpretation I such that $I \models \varphi$. Let $S = \{x \mid x \in X, x \in I\} \cup \{\overline{x} \mid x \in X, x \notin I\} \cup \{\varphi\}$. Clearly, $cl(S) = C$. Furthermore, S defends φ in $\mathcal{R}_4(\mathcal{P})$ since each clause is satisfied by I, and thus each clause argument ω_j is attacked by some x (or $\overline{x}$) in S. Each variable $x \in X$ clearly defends itself. Moreover, if $x \in S$, then $\overline{x} \notin S$ and none of b_x, $\overline{x}$, or $a_{\overline{x}}$ is defended by S. Analogously for the case that $\overline{x} \in S$. Thus, S is admissible, and contains all arguments it defends, i.e., $S \in com(\mathcal{R}_4(\mathcal{P}))$.

Assume $C \in com_{inh}(\mathcal{R}_4(\mathcal{P}))$. Then there is $S \subseteq A$ such that $cl(S) = C$ and $S \in com(\mathcal{R}_4(\mathcal{P}))$. For each $x \in X$, at least one of $x, \overline{x}$ must be contained in S. In fact, it cannot be that $x \in S$ and $\overline{x} \in S$, otherwise b_x would be defended by S and we would have $cl(S) \neq C$. Thus, for each $x \in X$, there is either $x \in S$ or $\overline{x} \in S$, but not both. Furthermore, S defends φ, i.e., S attacks all clause arguments ω_j. Thus, $I \models \varphi$ for $I = X \cap S$. □

4.4.3 Summary and Impact of Complexity Results

When using Reduction 1 we obtain the same complexity as for general CAFs, i.e., the benefits of wfCAFs are lost. On the other hand, Reductions 2–4 preserve the lower complexity of wfCAFs for almost all semantics. Intuitively, this can be explained by the fact that these reductions do not remove conflicts between arguments. This in turn means that Reductions 2–4 retain enough of the structure of wfCAFs in order to, given a claim, efficiently compute a subset-maximal admissible argument set with that claim. The only outlier is complete semantics, for which verification remains hard under Reductions 2 and 4 but not Reduction 3. Here, the fact that Reductions 2 and 4 can introduce new attacks leads to an increase in complexity. We conclude:

Table 4.1: Complexity of PCAFs. Results in boldface had to be proven explicitly. All other results follow directly from known properties (cf. Observation 4.28).

<table>
<tr><th rowspan="2">σ_μ^i</th><th>$Cred_{\sigma_\mu^i}^{PCAF}$</th><th colspan="2">$Skept_{\sigma_\mu^i}^{PCAF}$</th><th colspan="3">$Ver_{\sigma_\mu^i}^{PCAF}$</th></tr>
<tr><th>$i \in \{1,2,3,4\}$</th><th>$i=1$</th><th>$i \in \{2,3,4\}$</th><th>$i=1$</th><th>$i \in \{2,4\}$</th><th>$i=3$</th></tr>
<tr><td>cf_{inh}</td><td>in P</td><td colspan="2">trivial</td><td>NP-c</td><td colspan="2">in P</td></tr>
<tr><td>adm_{inh}</td><td>NP-c</td><td colspan="2">trivial</td><td>NP-c</td><td colspan="2">in P</td></tr>
<tr><td>com_{inh}</td><td>NP-c</td><td colspan="2">P-c</td><td colspan="2">NP-c</td><td>in P</td></tr>
<tr><td>stb_{inh}
$stb\text{-}adm_{hyb}$
$stb\text{-}cf_{hyb}$</td><td>NP-c</td><td colspan="2">coNP-c</td><td>NP-c</td><td colspan="2">in P</td></tr>
<tr><td>$naive_{inh}$</td><td rowspan="2">in P</td><td colspan="2">coNP-c</td><td>NP-c</td><td rowspan="2" colspan="2">in P</td></tr>
<tr><td>$naive_{hyb}$</td><td>Π_2^P-c</td><td>coNP-c</td><td>DP-c</td></tr>
<tr><td>prf_{inh}</td><td rowspan="2">NP-c</td><td rowspan="2" colspan="2">Π_2^P-c</td><td>Σ_2^P-c</td><td rowspan="2" colspan="2">coNP-c</td></tr>
<tr><td>prf_{hyb}</td><td>DP-c</td></tr>
<tr><td>sem_{inh}
sem_{hyb}</td><td>Σ_2^P-c</td><td colspan="2">Π_2^P-c</td><td>Σ_2^P-c</td><td colspan="2">coNP-c</td></tr>
<tr><td>stg_{inh}
stg_{hyb}</td><td>Σ_2^P-c</td><td colspan="2">Π_2^P-c</td><td>Σ_2^P-c</td><td colspan="2">coNP-c</td></tr>
</table>

Theorem 4.41. *The complexity results in Table 4.1 hold, even if we restrict ourselves to PCAFs with transitive preference relations.*

The lower complexity of the verification problem is crucial for enumerating claim-extensions in wfCAFs (Dvořák and Woltran 2020). Indeed, this is also true for PCAFs using Reductions 2–4, as we will now show. If claim sets can be verified in polynomial time we can simply iterate through all claim sets. For preferred, semi-stable, and stage semantics the algorithm builds heavily on the existence and polynomial-time computability of unique maximal realizations for conflict-free and admissible claim-sets, i.e., $E_1^i(C)$ and $E_*^i(C)$ (cf. Definition 4.36).

Proposition 4.42. *Consider PCAFs $\mathcal{P} = (A, R, cl, \succ)$ with $|A| \leq n$ and $|cl(A)| \leq k$.*

- *If $Ver_{\sigma_\mu^i}^{PCAF}$ is in* P *for a PCAF-semantics σ_μ^i, then there is a polynomial $poly(\cdot)$ such that $\sigma_\mu^i(\mathcal{P})$ can be enumerated in $O(2^k \cdot poly(n))$ time.*
- *For σ_μ^i with $\sigma_\mu \in \{prf_{inh}, prf_{hyb}, sem_{inh}, sem_{hyb}, stg_{inh}, stg_{hyb}\}$ and $i \in \{2,3,4\}$ there is a polynomial $poly(\cdot)$ such that $\sigma_\mu^i(\mathcal{P})$ can be enumerated in $O(4^k \cdot poly(n))$ time.*

Proof. If $Ver^{PCAF}_{\sigma^i_\mu}$ is in P we can iterate through all 2^k claim-sets $C \subseteq cl(A)$ and check whether $C \in \sigma^i(\mathcal{F})$ in polynomial time. This procedure runs in $O(2^k \cdot poly(n))$ time.

For σ^i_μ with $\sigma_\mu \in \{prf_{inh}, prf_{hyb}, sem_{inh}, sem_{hyb}, stg_{inh}, stg_{hyb}\}$ and $i \in \{2,3,4\}$, recall the proof of Proposition 4.39. There, to decide that $C \notin \sigma^i_\mu(\mathcal{P})$, we guessed a claim-set $D \subseteq cl(A)$ and performed some checks in polynomial time. Instead of guessing D, we can iterate through all 2^k claim-sets $D \subseteq cl(A)$. If $C \in adm_{inh}(\mathcal{P})$ (resp. $C \in cf_{inh}(\mathcal{P})$ in case $\sigma^i_\mu \in \{stg^i_{inh}, stg^i_{hyb}\}$), and if no D that witnesses $C \notin \sigma^i_\mu(\mathcal{P})$ is found, we have $C \in \sigma^i_\mu(\mathcal{P})$. Therefore, to enumerate $\sigma^i_\mu(\mathcal{P})$ we can iterate through all $(2^k)^2 = 4^k$ pairs (C, D) of claim-sets. This procedure runs in $O(4^k \cdot poly(n))$ time. □

Proposition 4.42 directly implies that deciding the main decision problems is tractable if the number of claims is bounded by a constant k, i.e., these problems are fixed parameter tractable (FPT).

Corollary 4.43. *For all PCAF-semantics σ^i_μ considered in this chapter, except for those such that $i = 1$ and except for com^2_{inh} and com^4_{inh}, there is a polynomial $poly(\cdot)$ such that $Cred^{PCAF}_{\sigma^i_\mu}$, $Skept^{PCAF}_{\sigma^i_\mu}$, and $Ver^{PCAF}_{\sigma^i_\mu}$ can be solved in time $O(4^k \cdot poly(n))$ for PCAFs $(A, R, cl, \succ)$ with $|cl(A)| \leq k$.*

4.5 Conclusion

Many approaches to argumentation assume that arguments with the same claims attack the same arguments. This gives rise to the natural class of wfCAFs, which enjoy several desired semantic and computational properties (Dvořák et al. 2023; Dvořák, Rapberger, and Woltran 2023). However, in formalisms in which preferences are used, well-formedness cannot be assumed in general. In this chapter, we analyzed whether the desired properties of wfCAFs still hold when preferences are taken into account. To this end, we introduced Preference-based CAFs (PCAFs) and investigated the impact of the four commonly used preference reductions on PCAFs.

We examined and characterized resulting CAF-classes, yielding insights into the expressiveness of argumentation formalisms that can be instantiated as CAFs and allow for preference incorporation. Furthermore, we investigated PCAFs with respect to semantic properties, computational complexity, and their relationship to structured formalisms. Preserving semantic properties such as I-maximality can be desirable since it implies intuitive behavior of maximization-based semantics, while the complexity of the verification problem is crucial for the enumeration of claim-extensions. Note that insights in terms of both semantical and computational properties often provide the necessary theoretical foundations towards practical realizations. This has been demonstrated before in similar research endeavors, e.g., for incomplete AFs (Baumeister et al. 2021; Fazzinga, Flesca, and Furfaro 2020).

Our results show that (i) Reduction 3 exhibits the same properties as wfCAFs regarding computational complexity, and mostly preserves semantic properties such as I-maximality;

(ii) Reductions 2 and 4 retain the advantages of wfCAFs regarding complexity for all but complete semantics, but do not preserve I-maximality; (iii) under Reduction 1, neither complexity properties nor semantic properties are preserved. The above results hold even if we restrict ourselves to transitive preferences. It is worth noting that Reduction 3 behaves favorably on standard Dung-style AFs as well, fulfilling many principles for preference-based semantics (Kaci et al. 2021).

Regarding future work, one could consider different methods for handling preferences, and examine the effect of these methods in (well-formed) CAFs. In this work, we dealt with preferences via preference reductions that modify the attack relation. Another approach is to lift orderings over arguments to sets of arguments and select extensions in this way (Alfano et al. 2022, 2023; Amgoud and Vesic 2014; Brewka, Truszczynski, and Woltran 2010; Kaci, van der Torre, and Villata 2018). These two paradigms interpret the meaning of preferences between arguments differently: using reductions, $x \succ y$ expresses that x is stronger than y, while in the second approach $x \succ y$ expresses that it is preferred to have outcomes with x rather than with y. Interestingly, under Reduction 3, the admissible/complete/stable extensions of a preference-based AF are also extensions in the underlying AF (Kaci et al. 2021). Thus, Reduction 3 selects the 'best' extensions from the underlying AF in these cases. A similar dichotomy concerning preference handling can be observed in related areas such as logic programs, where preferences can be incorporated either on the syntactic level or by ranking the outcome (Brewka et al. 2015; Brewka, Niemelä, and Syrjänen 2004; Delgrande, Schaub, and Tompits 2003; Sakama and Inoue 2000).

Another possibility for future work is to lower the level of abstraction used here, e.g., by incorporating more structure into arguments, by allowing arguments to act in support of other arguments as is done in bipolar AFs (Amgoud et al. 2008), or by preserving more information about the claims of arguments. Regarding the latter point, recent research (Wakaki 2020) has shown that formalisms which permit strong negation require careful examination with regards to consistency.

CHAPTER 5

Preferred Model Entailment in Choice Logics

In this chapter, we investigate preferred model entailment in choice logics, a crucial notion that allows us to draw conclusions from available knowledge about hard- and soft-constraints. Specifically, a choice logic theory T, i.e., a set of choice logic formulas, entails a classical formula F if and only if F is true in all preferred models of T. Preferred model entailment is closely related to notions of entailment in other non-monotonic systems such as circumscription (McCarthy 1980), default logic (Reiter 1980), autoepistemic logic (Moore 1985), or conditional knowledge bases (Kraus, Lehmann, and Magidor 1990; Lehmann and Magidor 1992). However, while these related formalisms have been extensively studied in the literature with respect to various properties (Bonatti and Olivetti 2002; Eiter and Gottlob 1993; Eiter and Lukasiewicz 2000), only little is known about preferred model entailment in choice logics. While Brewka, Benferhat, and Berre (2004) investigated Qualitative Choice Logic (QCL) regarding properties for defeasible entailment laid out by Kraus, Lehmann, and Magidor (1990), the subsequently introduced Concjunctive Choice Logic (CCL) (Boudjelida and Benferhat 2016) and Lexicographic Choice Logic (LCL) (Bernreiter 2020) have not been studied in this respect. Moreover, the computational and proof-theoretic properties of preferred model entailment in choice logics are entirely unknown.

When addressing these open questions regarding preferred model entailment in choice logics, we must consider two axes of generalization. Firstly, there are several choice logics such as QCL or CCL defined in the literature. Indeed, the general choice logic framework (cf. Section 2.4) allows us to study choice logics that belong to this framework but have not been defined explicitly yet. Secondly, the preferred models of a choice logic theory can be determined in several ways, with the lexicographic and inclusion-based approaches being the most commonly used such preferred model semantics in the literature (Brewka, Benferhat, and Berre 2004).

Contributions. We study the semantic, computational, and proof-theoretic properties of preferred model entailment in choice logics. We consider large classes of choice logics, with a particular focus on QCL, CCL, and LCL. Moreover, we investigate several preferred model semantics, including the lexicographic and inclusion-based approaches from the literature, but also two newly introduced semantics, namely the minmax and log-lexicographic approaches. In detail, the contributions of this chapter are as follows:

- We examine preferred model entailment with respect to some key properties for defeasible entailment laid out by Kraus, Lehmann, and Magidor (1990), namely cautious monotonicity, cumulative transitivity (called cut by Kraus, Lehmann, and Magidor (1990)), and rational monotonicity. Given finite theories, cautious monotonicity and cumulative transitivity are satisfied for all considered preferred model semantics and all choice logics. Moreover, rational monotonicity is satisfied under all preferred model semantics except for the inclusion-based approach.
- We study the complexity of preferred model checking, i.e., verifying whether a given interpretation is a preferred model, and find this problem to be coNP-complete for all preferred model semantics and all choice logics that feature more than two satisfaction degrees, which includes QCL, CCL, and LCL. We then investigate the complexity of deciding preferred model entailment, and find that it is located on the second level of the polynomial hierarchy. The exact complexity depends both on the considered preferred model semantics and the considered choice logic. For QCL and CCL, the complexity ranges from Θ_2^P-complete (minmax semantics) to Π_2^P-complete (inclusion-based semantics), with the other semantics being located inbetween ($\Delta_2^\mathsf{P}[O(\log^2 n)]$-complete for the log-lexicographic semantics and Δ_2^P-complete for the lexicographic semantics). For LCL, the complexity ranges from Δ_2^P-complete (minmax and (log-)lexicographic semantics) to Π_2^P-complete (inclusion-based semantics).
- Lastly, we introduce sequent calculi for preferred model entailment. We consider various choice logics (QCL, CCL, and LCL) and preferred model semantics (lexicographic, inclusion-based, and minmax). Each calculus for preferred model entailment is based on two labeled calculi, a monotonic calculus and a refutation calculus, together with a non-monotonic rule. We show that all considered calculi are sound and complete.

Publications. This chapter is based on the papers (Bernreiter, Maly, and Woltran 2022), (Bernreiter et al. 2022b), and (Bernreiter et al. 2024a).

Outline. In Section 5.1 we define the notion of preferred model entailment and the various preferred model semantics formally. In Section 5.2 we investigate preferred model entailment with respect to the properties of Kraus, Lehmann, and Magidor (1990). Section 5.3 contains the complexity analysis of preferred model entailment, while our sequent calculi are presented in Section 5.4. We conclude in Section 5.5.

Required preliminaries. Before reading this chapter, it is recommended to read Section 2.1 (propositional logic), Section 2.2 (computational complexity), and Section 2.4 (choice logics).

5.1 Formal Definition

Formally, if $\mathcal{L}$ is a choice logic, then a set of $\mathcal{L}$-formulas is called an $\mathcal{L}$-theory. An $\mathcal{L}$-theory T entails a classical formula F, written as $T \mid\!\sim F$, if F is true in all preferred models of T. However, we first need to define what the preferred models of a choice logic theory are. There are several approaches for how to determine the preferred models of a choice logic theory. In the original QCL paper (Brewka, Benferhat, and Berre 2004), a lexicographic and an inclusion-based approach were introduced. In addition, we introduce the simple yet previously not considered minmax-approach.

Definition 5.1. *Let $\mathcal{L}$ be a choice logic, $\mathcal{I}$ an interpretation, and T an $\mathcal{L}$-theory. $\mathcal{I}$ is a model of T, written as $\mathcal{I} \in Mod_{\mathcal{L}}(T)$, iff $deg_{\mathcal{L}}(\mathcal{I}, F) < \infty$ for all $F \in T$. Moreover, $\mathcal{I}^k_{\mathcal{L}}(T)$ denotes the set of formulas in T satisfied to a degree of k by $\mathcal{I}$, that is, $\mathcal{I}^k_{\mathcal{L}}(T) = \{F \in T \mid deg_{\mathcal{L}}(\mathcal{I}, F) = k\}$.*

- *$\mathcal{I}$ is a minmax preferred model of T, denoted $\mathcal{I} \in Prf^{mm}_{\mathcal{L}}(T)$, iff $\mathcal{I} \in Mod_{\mathcal{L}}(T)$ and there is no $\mathcal{J} \in Mod_{\mathcal{L}}(T)$ such that $max\{deg_{\mathcal{L}}(\mathcal{I}, F) \mid F \in T\} > max\{deg_{\mathcal{L}}(\mathcal{J}, F) \mid F \in T\}$.*

- *$\mathcal{I}$ is a lexicographically preferred model of T, denoted $\mathcal{I} \in Prf^{lex}_{\mathcal{L}}(T)$, iff $\mathcal{I} \in Mod_{\mathcal{L}}(T)$ and there is no $\mathcal{J} \in Mod_{\mathcal{L}}(T)$ such that, for some $k \in \mathbb{N}$ and all $l < k$, $|\mathcal{I}^k_{\mathcal{L}}(T)| < |\mathcal{J}^k_{\mathcal{L}}(T)|$ and $|\mathcal{I}^l_{\mathcal{L}}(T)| = |\mathcal{J}^l_{\mathcal{L}}(T)|$ holds.*

- *$\mathcal{I}$ is an inclusion-based preferred model of T, denoted $\mathcal{I} \in Prf^{inc}_{\mathcal{L}}(T)$, iff $\mathcal{I} \in Mod_{\mathcal{L}}(T)$ and there is no $\mathcal{J} \in Mod_{\mathcal{L}}(T)$ such that, for some $k \in \mathbb{N}$ and all $l < k$, $\mathcal{I}^k_{\mathcal{L}}(T) \subset \mathcal{J}^k_{\mathcal{L}}(T)$ and $\mathcal{I}^l_{\mathcal{L}}(T) = \mathcal{J}^l_{\mathcal{L}}(T)$ holds.*

Intuitively, under the minmax approach a finite $\mathcal{L}$-theory $T = \{A_1, \ldots, A_n\}$ can be seen as the $\mathcal{L}$-formula $A_1 \wedge \cdots \wedge A_n$. On the other hand, the lexicographic and inclusion-based approaches choose those models as preferred models that satisfy as many formulas as possible in the theory to a degree of 1. If there is a tie between two interpretations with regards to degree 1, then it is determined which interpretation satisfies more formulas to a degree of 2, and so forth. The differences between the two approaches is how the phrase 'as many degrees as possible' is understood: either in terms of cardinality (lexicographic approach) or in terms of subset-maximization (inclusion-based approach).

We now provide an example for the various preferred model semantics. Note that we will only give the interpretations that are relevant to a given theory T, i.e., only the interpretations that contain variables that actually occur in T.

Example 5.2. *Consider the* QCL*-theory*

$$T = \{a \vec{\times} c, b \vec{\times} c, \neg(a \wedge b)\}.$$

Regarding the classical models of T we have $Mod_{\text{QCL}}(T) = \{\{c\}, \{a, c\}, \{b, c\}\}$.

Regarding the minmax semantics, note that for all three models $\mathcal{I} \in Mod_{\text{QCL}}(T)$ we have $max\{deg_{\mathcal{L}}(\mathcal{J}, F) \mid F \in T\} = 2$. Thus, $Prf^{mm}_{\text{QCL}}(T) = \{\{a, c\}, \{b, c\}, \{c\}\}$.

For the lexicographic approach, observe that

$$\{c\}^1_{\text{QCL}}(T) = \{\neg(a \wedge b)\} \text{ and } \{c\}^2_{\text{QCL}}(T) = \{a \vec{\times} c, b \vec{\times} c\},$$
$$\{a, c\}^1_{\text{QCL}}(T) = \{a \vec{\times} c, \neg(a \wedge b)\} \text{ and } \{a, c\}^2_{\text{QCL}}(T) = \{b \vec{\times} c\},$$
$$\{b, c\}^1_{\text{QCL}}(T) = \{b \vec{\times} c, \neg(a \wedge b)\} \text{ and } \{b, c\}^2_{\text{QCL}}(T) = \{a \vec{\times} c\}.$$

Thus, $\{c\}$ satisfies less formulas to a degree of 1 than $\{a, c\}$ or $\{b, c\}$. Formally,

$$|\{c\}^1_{\text{QCL}}(T)| = 1 \text{ and } |\{c\}^2_{\text{QCL}}(T)| = 2,$$
$$|\{a, c\}^1_{\text{QCL}}(T)| = 2 \text{ and } |\{a, c\}^2_{\text{QCL}}(T)| = 1,$$
$$|\{b, c\}^1_{\text{QCL}}(T)| = 2 \text{ and } |\{b, c\}^2_{\text{QCL}}(T)| = 1.$$

Therefore, $Prf^{lex}_{\text{QCL}}(T) = \{\{a, c\}, \{b, c\}\}$, i.e., $\{c\} \notin Prf^{lex}_{\text{QCL}}(T)$.

Regarding the inclusion-based semantics, note that $\{c\}^1_{\text{QCL}}(T) \subset \{a, c\}^1_{\text{QCL}}(T)$. Moreover, $\{a, c\}^1_{\text{QCL}}(T) \not\subset \{b, c\}^1_{\text{QCL}}(T)$ and $\{b, c\}^1_{\text{QCL}}(T) \not\subset \{a, c\}^1_{\text{QCL}}(T)$. We can conclude that $Prf^{inc}_{\text{QCL}}(T) = \{\{a, c\}, \{b, c\}\}$.

Note that under the minmax semantics the ranking of an interpretation depends only on a single formula, namely that with the highest degree under the given interpretation. In contrast, under the lexicographic and inclusion-based semantics all formulas in a theory may influence an interpretation's ranking. We will now propose a preferred model semantics that constitutes a middle ground in this matter: given an $\mathcal{L}$-theory $T = \{A_1, \ldots, A_n\}$ and an interpretation $\mathcal{I}$, the $\log(n)$ formulas with the highest degree will influence the ranking of $\mathcal{I}$. The resulting preferred model semantics is defined below, and is especially interesting for our complexity analysis in Section 5.3.

Definition 5.3. *Let $\mathcal{L}$ be a choice logic and $T = \{A_1, \ldots, A_n\}$ a finite $\mathcal{L}$-theory. The log-worst formulas of T relative to an interpretation $\mathcal{I}$ is a set $L^{\mathcal{I}}_{\mathcal{L}}(T)$ such that $|L^{\mathcal{I}}_{\mathcal{L}}(T)| = \lceil \log(n) \rceil$ and such that, for all $A \in T \setminus L^{\mathcal{I}}_{\mathcal{L}}(T)$, we have that $deg_{\mathcal{L}}(\mathcal{I}, A) \leq min\{deg_{\mathcal{L}}(\mathcal{I}, B) \mid B \in L^{\mathcal{I}}_{\mathcal{L}}(T)\}$. An interpretation $\mathcal{I}$ is a log-lexicographically preferred model of T, written as $\mathcal{I} \in Prf^{log}_{\mathcal{L}}(T)$, iff $\mathcal{I} \in Mod_{\mathcal{L}}(T)$ and if there is no $\mathcal{J} \in Mod_{\mathcal{L}}(T)$ such that, for some $k \in \mathbb{N}$ and all $l > k$, $|\mathcal{I}^k_{\mathcal{L}}(L^{\mathcal{I}}_{\mathcal{L}}(T))| > |\mathcal{J}^k_{\mathcal{L}}(L^{\mathcal{J}}_{\mathcal{L}}(T))|$ and $|\mathcal{I}^l_{\mathcal{L}}(L^{\mathcal{I}}_{\mathcal{L}}(T))| = |\mathcal{J}^l_{\mathcal{L}}(L^{\mathcal{J}}_{\mathcal{L}}(T))|$ holds.*

In the log-lexicographic approach, satisfaction degrees are considered in a top-down manner, i.e., we strive to minimize the number of formulas satisfied to high degrees. As a result, $\mathcal{I} \in \mathit{Prf}_{\mathcal{L}}^{log}(T)$ implies $\mathcal{I} \in \mathit{Prf}_{\mathcal{L}}^{mm}(T)$, meaning that the log-lexicographic semantics is a refinement of the minmax semantics.

Moreover, note that $L_{\mathcal{L}}^{\mathcal{I}}(T)$ from Definition 5.3 is not necessarily unique for a given T. However, for the log-lexicographic semantics it is of no importance which exact $L_{\mathcal{L}}^{\mathcal{I}}(T)$ is considered, as we only care about how many formulas are satisfied to certain degrees, not *which* formulas.

Example 5.4. *Let* $T = \{a \overrightarrow{\times} c, b \overrightarrow{\times} c, \neg(a \wedge b)\}$*, just as in Example 5.2. Note that* $\lceil \log_2(3) \rceil = 2$*, i.e., given an interpretation, using the log-lexicographic semantics we are only interested in those two formulas that are satisfied to a maximal degree. For* $\{c\}$ *this is* $L_{\text{QCL}}^{\{c\}}(T) = \{a \overrightarrow{\times} c, b \overrightarrow{\times} c\}$ *with both formulas satisfied to a degree of* 2*. For* $\{a, c\}$ *this can be* $L_{\text{QCL}}^{\{a,c\}}(T) = \{\neg(a \wedge b), b \overrightarrow{\times} c\}$ *or* $L_{\text{QCL}}^{\{a,c\}}(T) = \{a \overrightarrow{\times} c, b \overrightarrow{\times} c\}$ *with* $b \overrightarrow{\times} c$ *satisfied to a degree of* 2 *but* $\neg(a \wedge b)$ *and* $a \overrightarrow{\times} b$ *satisfied to a degree of* 1*. Analogously for* $L_{\text{QCL}}^{\{b,c\}}(T)$*. Thus,* $\mathit{Prf}_{\text{QCL}}^{log}(T) = \{\{a, c\}, \{b, c\}\}$*.*

Now that we introduced some ways to determine the preferred models of a choice logic theory, we formally define the notion of preferred model entailment and provide a small example.

Definition 5.5. *Let* $\mathcal{L}$ *be a choice logic,* T *an* $\mathcal{L}$*-theory,* F *a classical formula, and* σ *a preferred model semantics, e.g.,* $\sigma \in \{mm, lex, inc, log\}$*. Then* $T \mid\!\sim_{\mathcal{L}}^{\sigma} F$ *if and only if* $\mathcal{I} \in \mathit{Prf}_{\mathcal{L}}^{\sigma}(T)$ *implies* $\mathcal{I} \models F$*.*

Example 5.6. *Let* T *be the same theory as in Example 5.2, and recall that* $\mathit{Prf}_{\text{QCL}}^{lex}(T) = \{\{a, c\}, \{b, c\}\}$*. Thus,* $T \mid\!\sim_{\text{QCL}}^{lex} c \wedge (a \vee b)$*. However,* $T \not\mid\!\sim_{\text{QCL}}^{lex} a$ *and* $T \not\mid\!\sim_{\text{QCL}}^{lex} b$*. Analogously for* $\mid\!\sim_{\text{QCL}}^{inc}$ *and* $\mid\!\sim_{\text{QCL}}^{log}$*. Moreover, recall that* $\mathit{Prf}_{\text{QCL}}^{mm}(T) = \{\{a, c\}, \{b, c\}, \{c\}\}$*. Thus,* $T \not\mid\!\sim_{\text{QCL}}^{mm} c \wedge (a \vee b)$*.*

In the following sections, we examine the notion of preferred models with regards to logical properties, computational complexity, and proof systems.

5.2 Logical Properties

In this section, we investigate preferred model entailment with respects to some key properties for defeasible inference laid out by Kraus, Lehmann, and Magidor (1990). Specifically, we consider the properties of cautious monotonicity, cumulative transitivity (called cut by Kraus, Lehmann, and Magidor (1990)), and rational monotonicity. First however, we show that preferred model entailment is non-monotonic for all preferred model semantics considered in this chapter (mm, lex, inc, log) and for all choice logics in which more than two satisfaction degrees can be obtained (which of course includes QCL, CCL, and LCL). Recall that $\mathcal{D}_\mathcal{L}$ denotes the obtainable degrees in a choice logic $\mathcal{L}$ (see Definition 2.39).

Proposition 5.7. *Let $\mathcal{L}$ be a choice logic such that $\mathcal{D}_\mathcal{L} \neq \{1, \infty\}$ and let $\sigma \in \{mm, lex, inc, log\}$. The preferred model entailment $\mathrel{|\!\sim}^\sigma_\mathcal{L}$ is non-monotonic, i.e., $T \mathrel{|\!\sim}^\sigma_\mathcal{L} B$ does not necessarily imply $T \cup \{A\} \mathrel{|\!\sim}^\sigma_\mathcal{L} B$.*

Proof. Let $k \in \mathcal{D}_\mathcal{L} \setminus \{1, \infty\}$ and let a be a propositional variable. By Proposition 2.40 we know that there is an $\mathcal{L}$-formula F such that $deg_\mathcal{L}(\mathcal{I}, F) = 1$ if $a \in \mathcal{I}$ and $deg_\mathcal{L}(\mathcal{I}, F) = k$ if $a \notin \mathcal{I}$. Then under all considered semantics $\mathcal{I}$ is a preferred model of F if and only if it contains a. Therefore, we have $\{F\} \mathrel{|\!\sim}^\sigma_\mathcal{L} a$. Now consider $\{F\} \cup \{\neg a\}$. We observe that $\mathcal{I} \in Mod_\mathcal{L}(\{F\} \cup \{\neg a\})$ if and only if $a \notin \mathcal{I}$. It follows that for all preferred model semantics $\{F\} \cup \{\neg a\} \mathrel{|\!\not\sim}^\sigma_\mathcal{L} a$. □

The first property we examine is that of cautious monotonicity, where $T \mathrel{|\!\sim} A$ and $T \mathrel{|\!\sim} B$ implies $T \cup \{A\} \mathrel{|\!\sim} B$. In the original QCL-paper, Brewka, Benferhat, and Berre (2004) show that $\mathrel{|\!\sim}^{lex}_{\text{QCL}}$ satisfies this property. In fact, cautious monotonicity is satisfied by all choice logics under all preferred model semantics considered in this chapter:

Proposition 5.8. *Let $\mathcal{L}$ be a choice logic and $\sigma \in \{mm, lex, inc, log\}$. The inference relation $\mathrel{|\!\sim}^\sigma_\mathcal{L}$ satisfies cautious monotonicity for finite theories, i.e., $T \mathrel{|\!\sim}^\sigma_\mathcal{L} A$ and $T \mathrel{|\!\sim}^\sigma_\mathcal{L} B$ implies $T \cup \{A\} \mathrel{|\!\sim}^\sigma_\mathcal{L} B$ for all finite $\mathcal{L}$-theories T and all classical formulas A, B.*

Proof. Assume $T \mathrel{|\!\sim}^\sigma_\mathcal{L} A$ and $T \mathrel{|\!\sim}^\sigma_\mathcal{L} B$. Note that A and B are classical formulas. Let $\mathcal{I} \in Prf^\sigma_\mathcal{L}(T \cup \{A\})$. Then $\mathcal{I} \in Mod_\mathcal{L}(T)$ and $deg_\mathcal{L}(\mathcal{I}, A) = 1$. Towards a contradiction, assume $\mathcal{I} \notin Prf^\sigma_\mathcal{L}(T)$. Since $\mathcal{I} \in Mod_\mathcal{L}(T)$, and since T is finite, there must be $\mathcal{J} \in Prf^\sigma_\mathcal{L}(T)$ that is more preferable than $\mathcal{I}$ with respect to T. By $T \mathrel{|\!\sim}^\sigma_\mathcal{L} A$ also $deg_\mathcal{L}(\mathcal{J}, A) = 1$. We claim that then $\mathcal{J}$ is more preferable than $\mathcal{I}$ for $T \cup \{A\}$ for all considered semantics, which is a contradiction:

- For $\sigma = mm$ the minmax semantics, observe that $max\{deg_\mathcal{L}(\mathcal{I}, F) \mid F \in T\} = max\{deg_\mathcal{L}(\mathcal{I}, F) \mid F \in T \cup \{A\}\}$. The same holds for $\mathcal{J}$.
- For $\sigma \in \{lex, inc\}$, observe that for $l \neq 1$ we have $\mathcal{I}^l_\mathcal{L}(T \cup \{A\}) = \mathcal{I}^l_\mathcal{L}(T)$ and for $l = 1$ we have $\mathcal{I}^1_\mathcal{L}(T \cup \{A\}) = \mathcal{I}^1_\mathcal{L}(T) \cup \{A\}$. The same holds for $\mathcal{J}$.
- For $\sigma = log$ there exists a $k \in \mathbb{N}$ such that for all $l > k$, $|\mathcal{I}^k_\mathcal{L}(L^\mathcal{I}_\mathcal{L}(T))| > |\mathcal{J}^k_\mathcal{L}(L^\mathcal{J}_\mathcal{L}(T))|$ and $|\mathcal{I}^l_\mathcal{L}(L^\mathcal{I}_\mathcal{L}(T))| = |\mathcal{J}^l_\mathcal{L}(L^\mathcal{J}_\mathcal{L}(T))|$. As $|L^\mathcal{I}_\mathcal{L}(T)| = |L^\mathcal{J}_\mathcal{L}(T)|$ it is not possible that $k = 1$. Therefore, we must have $|\mathcal{I}^k_\mathcal{L}(L^\mathcal{I}_\mathcal{L}(T \cup \{A\}))| > |\mathcal{J}^k_\mathcal{L}(L^\mathcal{J}_\mathcal{L}(T \cup \{A\}))|$ and $|\mathcal{I}^l_\mathcal{L}(L^\mathcal{I}_\mathcal{L}(T \cup \{A\}))| = |\mathcal{J}^l_\mathcal{L}(L^\mathcal{J}_\mathcal{L}(T \cup \{A\}))|$ for all $l > k$.

Thus, $\mathcal{I} \in Prf^\sigma_\mathcal{L}(T)$ and by $T \mathrel{|\!\sim}^\sigma_\mathcal{L} B$ also $\mathcal{I} \models B$. □

Note that we only considered finite theories in Proposition 5.8. Brewka, Benferhat, and Berre (2004) did not explicitly make this assumption when investigating $\mid\!\sim^{lex}_{\mathrm{QCL}}$, but we believe it was implicitly assumed. In fact, $\mid\!\sim^{lex}_{\mathrm{QCL}}$ does not satisfy cautious monotonicity if infinite theories are allowed. The following result can likely be generalized for other preferred model semantics and choice logics, but we do not consider this here.

Proposition 5.9. *The inference relation $\mid\!\sim^{lex}_{\mathrm{QCL}}$ does not satisfy cautious monotonicity for infinite theories, i.e., there is an infinite* QCL*-theory T and classical formulas A, B such that $T \mid\!\sim^{lex}_{\mathrm{QCL}} A$, $T \mid\!\sim^{lex}_{\mathrm{QCL}} B$, but $T \cup \{A\} \not\mid\!\sim^{lex}_{\mathrm{QCL}} B$.*

Proof. In the following, by $\top^i$ we denote i occurrences of $\top$ connected by $\vec{\times}$. For example, $\top^1 = \top$, $\top^2 = \top\vec{\times}\top$, and $\top^3 = \top\vec{\times}\top\vec{\times}\top$. For $i \in \mathbb{N}$, let

$$A_i = \{a_i\vec{\times}\top^1, a_i\vec{\times}\top^2, \ldots, a_i\vec{\times}\top^i\}.$$

The additional occurrences of $\top$ are added simply to duplicate the formula $a_i\vec{\times}\top$. Observe that $\{a_i\}$ satisfies all i formulas in A_i to a degree of 1, and that an interpretation that sets a_i to false satisfies all i formulas in A_i to a degree of 2. Furthermore, we define

$$R = \{\bot\vec{\times}\neg(a_i \wedge a_j) \mid i, j \in \mathbb{N}, i \neq j\}$$

and finally

$$T = R \cup \bigcup_{i \in \mathbb{N}} A_i.$$

The formulas in R can at best be satisfied to a degree of 2. Furthermore, R enforces that, if $i \neq j$, then a_i and a_j cannot be set to true by the same interpretation. For any $i \in \mathbb{N}$, the interpretation $\{a_i\}$ satisfies all formulas in T to some finite degree. Moreover, $|\{a_i\}^1_{\mathrm{QCL}}(T)| = i$. This means that, for any $i \in \mathbb{N}$, $\{a_i\}$ is not a preferred model of T since there is always $j > i$ such that $|\{a_i\}^1_{\mathrm{QCL}}(T)| < |\{a_j\}^1_{\mathrm{QCL}}(T)|$. In fact, T has no preferred models, since if some $\mathcal{I}$ does not satisfy any a_i, then $|\mathcal{I}^1_{\mathrm{QCL}}(T)| = 0$. Thus, $T \mid\!\sim^{lex}_{\mathrm{QCL}} a_1$ and $T \mid\!\sim^{lex}_{\mathrm{QCL}} a_2$ hold vacuously.

Now we consider the extension of T by a_1, i.e. $T \cup \{a_1\}$. To satisfy $T \cup \{a_1\}$, a_1 must be set to true, but a_1 cannot be set to true at the same time as any other a_i with $i \neq 1$. Thus, $\{a_1\}$ is the only model of $T \cup \{a_1\}$ (without loss of generality, we can assume $\mathcal{I} \subseteq \{a_1, a_2, \ldots\}$ for all interpretations $\mathcal{I}$ we are dealing with). Therefore, $\{a_1\}$ is the single preferred model of $T \cup \{a_1\}$ which means that $T \cup \{a_1\} \not\mid\!\sim^{lex}_{\mathrm{QCL}} a_2$, □

The next property we examine is that of cumulative transitivity, which is known to hold for $\mid\!\sim^{lex}_{\mathrm{QCL}}$ (Brewka, Benferhat, and Berre 2004). As before with cautious monotonicity, we can generalize this result for the other preferred model semantics and all choice logics. Note that the following proposition holds for infinite theories, except for the log-lexicographic semantics which are only defined for finite theories.

Proposition 5.10. *Let $\mathcal{L}$ be a choice logic and $\sigma \in \{mm, lex, inc, log\}$. The inference relation $\mid\!\sim^{\sigma}_{\mathcal{L}}$ satisfies cumulative transitivity, i.e., $T \mid\!\sim^{\sigma}_{\mathcal{L}} A$ and $T \cup \{A\} \mid\!\sim^{\sigma}_{\mathcal{L}} B$ implies $T \mid\!\sim^{\sigma}_{\mathcal{L}} B$ for all $\mathcal{L}$-theories T and all classical formulas A, B.*

Proof. Assume $T \mid\!\sim^{\sigma}_{\mathcal{L}} A$ and $T \cup \{A\} \mid\!\sim^{\sigma}_{\mathcal{L}} B$. Note that A and B are classical formulas. Let $\mathcal{I} \in \mathit{Prf}^{\sigma}_{\mathcal{L}}(T)$. Then $\mathcal{I} \in \mathit{Mod}_{\mathcal{L}}(T)$ and, by $T \mid\!\sim^{\sigma}_{\mathcal{L}} A$, $\mathit{deg}_{\mathcal{L}}(\mathcal{I}, A) = 1$. It is easy to see that $\mathcal{I} \in \mathit{Prf}^{\sigma}_{\mathcal{L}}(T \cup \{A\})$ which, by $T \cup \{A\} \mid\!\sim^{\sigma}_{\mathcal{L}} B$, implies $\mathcal{I} \models B$. □

Lastly, (Brewka, Benferhat, and Berre 2004) also considered the property of rational monotonicity and showed that it is satisfied by $\mid\!\sim^{lex}_{\mathrm{QCL}}$. The result is achieved by transforming all formulas in a given QCL-theory into a normal form, and to then further transform the theory into a stratified knowledge base. We now show that rational monotonicity is satisfied in all choice logics and under all considered preferred model semantics except the inclusion-based approach. Note that we give a direct proof, and do not require a translation to another formalism.

Proposition 5.11. *Let $\mathcal{L}$ be a choice logic and $\sigma \in \{mm, lex, log\}$. The inference relation $\mid\!\sim^{\sigma}_{\mathcal{L}}$ satisfies rational monotonicity, i.e., $T \cup \{A\} \not\mid\!\sim^{\sigma}_{\mathcal{L}} B$ and $T \not\mid\!\sim^{\sigma}_{\mathcal{L}} \neg A$ implies $T \not\mid\!\sim^{\sigma}_{\mathcal{L}} B$ for all $\mathcal{L}$-theories T and all classical formulas A, B.*

Proof. Assume $T \cup \{A\} \not\mid\!\sim^{\sigma}_{\mathcal{L}} B$ and $T \not\mid\!\sim^{\sigma}_{\mathcal{L}} \neg A$. Then there is $\mathcal{I} \in \mathit{Prf}^{\sigma}_{\mathcal{L}}(T \cup \{A\})$ such that $\mathcal{I} \not\models B$ and $\mathcal{J} \in \mathit{Prf}^{\sigma}_{\mathcal{L}}(T)$ such that $\mathcal{J} \models A$. Note that $\mathit{deg}_{\mathcal{L}}(\mathcal{I}, A) = \mathit{deg}_{\mathcal{L}}(\mathcal{J}, A) = 1$. Towards a contradiction, assume that $\mathcal{I} \notin \mathit{Prf}^{\sigma}_{\mathcal{L}}(T)$. Then $\mathcal{J}$ is more preferable than $\mathcal{I}$ for T, which also means that $\mathcal{J}$ is more preferable than $\mathcal{I}$ for $T \cup \{A\}$, i.e., $\mathcal{I} \notin \mathit{Prf}^{\sigma}_{\mathcal{L}}(T \cup \{A\})$. Contradiction. Thus, $\mathcal{I} \in \mathit{Prf}^{\sigma}_{\mathcal{L}}(T)$. Since $\mathcal{I} \not\models B$ we have $T \not\mid\!\sim^{\sigma}_{\mathcal{L}} B$. □

Proposition 5.12. *Let $\mathcal{L}$ be a choice logic such that $\mathcal{D}_{\mathcal{L}} \neq \{1, \infty\}$. The inference relation $\mid\!\sim^{inc}_{\mathcal{L}}$ does not satisfy rational monotonicity, i.e., there is an $\mathcal{L}$-theory T and classical formulas A, B such that $T \cup \{A\} \not\mid\!\sim^{\sigma}_{\mathcal{L}} B$, $T \not\mid\!\sim^{\sigma}_{\mathcal{L}} \neg A$, but $T \mid\!\sim^{\sigma}_{\mathcal{L}} B$.*

Proof. Let $k \in \mathcal{D}_{\mathcal{L}} \setminus \{1, \infty\}$. Let a, b, c, be propositional variables. By Proposition 2.40 we know that there are $\mathcal{L}$-formulas F, G, H over $\{a, b, c\}$ such that

$$\begin{gathered}
\mathit{deg}_{\mathcal{L}}(\{a\}, F) = 1, \mathit{deg}_{\mathcal{L}}(\{a\}, G) = k, \mathit{deg}_{\mathcal{L}}(\{a\}, H) = k \\
\mathit{deg}_{\mathcal{L}}(\{b\}, F) = k, \mathit{deg}_{\mathcal{L}}(\{b\}, G) = 1, \mathit{deg}_{\mathcal{L}}(\{b\}, H) = 1 \\
\mathit{deg}_{\mathcal{L}}(\{c\}, F) = k, \mathit{deg}_{\mathcal{L}}(\{c\}, G) = k, \mathit{deg}_{\mathcal{L}}(\{c\}, H) = 1,
\end{gathered}$$

and $\mathit{deg}_{\mathcal{L}}(\mathcal{I}, F) = \mathit{deg}_{\mathcal{L}}(\mathcal{I}, G) = \mathit{deg}_{\mathcal{L}}(\mathcal{I}, H) = \infty$ for all other $\mathcal{I} \subseteq \{a, b, c\}$. Let $T = \{F, G, H\}$, $A = \neg b$, and $B = a \vee b$. It can be verified that $\mathit{Prf}^{inc}_{\mathcal{L}}(T) = \{\{a\}, \{b\}\}$ and $\mathit{Prf}^{inc}_{\mathcal{L}}(T \cup \{A\}) = \{\{a\}, \{c\}\}$. Thus, $T \cup \{A\} \not\mid\!\sim^{\sigma}_{\mathcal{L}} B$, $T \not\mid\!\sim^{\sigma}_{\mathcal{L}} \neg A$, but $T \mid\!\sim^{\sigma}_{\mathcal{L}} B$. □

In conclusion, we have shown that preferred model entailment in choice logics satisfies the key properties of cautious monotonicity and cumulative transitivity, which Lehmann and Magidor (1992) include in their list of properties that consequence relations are "expected to satisfy". Moreover, preferred model entailment under all except the inclusion-based preferred model semantics satisfies rational monotonicity.

5.3 Complexity

In this section, we analyze the computational complexity of preferred model entailment in choice logics. We consider the minmax, lexicographic, inclusion-based, and log-lexicographic preferred model semantics. While membership results will apply to all tractable choice logics (cf. Definition 2.41), hardness results will be less general and mainly apply to specific choice logics. In particular, we investigate QCL, CCL, LCL, and, as a base line, classical propositional logic (PL). The considered decision problems are checking whether a given interpretation is a preferred model, and deciding preferred model entailment:

Definition 5.13. *Let $\mathcal{L}$ be a choice logic and σ a preferred-model semantics for $\mathcal{L}$-theories, e.g., $\sigma \in \{mm, lex, inc, log\}$. We define the decision problems*

- *$\mathcal{L}$-PMCHECKING[σ]: Given a finite $\mathcal{L}$-theory T and an interpretation $\mathcal{I}$, decide whether $\mathcal{I} \in Prf^{\sigma}_{\mathcal{L}}(T)$;*
- *$\mathcal{L}$-ENTAILMENT[σ]: Given a finite $\mathcal{L}$-theory T and a classical formula F, decide whether $T \mathrel{|\!\sim}^{\sigma}_{\mathcal{L}} F$.*

A summary of results for the above problems can be found at the end of this section in Table 5.1. Regarding $\mathcal{L}$-PMCHECKING[σ] for choice logic theories, we see no complexity increase compared to $\mathcal{L}$-PMCHECKING for single formulas (cf. Table 2.5).

Proposition 5.14. *For all $\sigma \in \{mm, lex, inc, log\}$, $\mathcal{L}$-PMCHECKING[σ] is in* P *for $\mathcal{L} = \mathrm{PL}$ and in* coNP *for any tractable choice logic $\mathcal{L}$. $\mathcal{L}$-PMCHECKING[σ] is* coNP-*complete for any tractable choice logic $\mathcal{L}$ for which $\mathcal{D}_{\mathcal{L}} \neq \{1, \infty\}$ holds.*

Proof. P-membership for PL is straightforward, since we only need to verify whether the given interpretation $\mathcal{I}$ satisfies all formulas in the given theory T. Furthermore, coNP-hardness for tractable choice logics with $\mathcal{D}_{\mathcal{L}} \neq \{1, \infty\}$ follows from the complexity of $\mathcal{L}$-PMCHECKING for single choice logic formulas (see Table 2.5). As for coNP-membership of tractable choice logics, it is easy to see that the complementary problem is in NP: given an $\mathcal{L}$-theory T and an interpretation $\mathcal{I}$, guess an interpretation $\mathcal{J}$ and check whether $\mathcal{J}$ is more preferred than $\mathcal{I}$. This can be done in polynomial time for all preferred-model semantics we are concerned with ($\sigma \in \{mm, lex, inc, log\}$). □

Note that in classical propositional logic the problem PL-Entailment$[\sigma]$ with $\sigma \in \{mm, lex, inc, log\}$ is simply the problem of classical entailment and is thus coNP-complete by well-known properties of PL. We now turn our attention to $\mathcal{L}$-Entailment$[mm]$ in the general case. For convenience, we write $opt_{\mathcal{L}}(T)$ for $max\{opt_{\mathcal{L}}(F) \mid F \in T\}$. Moreover, recall that $opt_{\mathcal{L}}(F) < 2^{(|F|^2)}$ (see Lemma 2.43). Thus, showing that a function is logarithmic in $opt_{\mathcal{L}}(F)$ also shows that the function is polynomial in $|F|$. Finally, recall that in QCL and CCL, but not LCL, we have $opt_{\mathcal{L}}(F) \leq |F|$ (see Section 2.4).

Proposition 5.15. *$\mathcal{L}$-Entailment$[mm]$ is in Δ_2^{P} and coNP-hard for all tractable choice logics. $\mathcal{L}$-Entailment$[mm]$ is in Θ_2^{P} for a tractable choice logic $\mathcal{L}$ if for some constant c and all $\mathcal{L}$-formulas F it holds that $opt_{\mathcal{L}}(F) \in O(|F|^c)$.*

Proof. coNP-hardness follows from coNP-hardness of PL. We show membership for the complementary problem, where we ask whether a classical formula F evaluates to false under some minmax preferred model of a given theory $T = \{A_1, \ldots, A_n\}$: first, we conduct a binary search over $(1, \ldots, opt_{\mathcal{L}}(T))$, in each step using an NP-oracle to answer whether there is an interpretation $\mathcal{I}$ such that $max\{deg_{\mathcal{L}}(\mathcal{I}, A) \mid A \in T\} \leq k$, where k is the current position of the binary search. In this way, we find the minimum m such that $max\{deg_{\mathcal{L}}(\mathcal{J}, A) \mid A \in T\} = m$ for any interpretation $\mathcal{J}$. If $m = \infty$, we already have a no-instance since $Prf_{\mathcal{L}}^{mm}(T) = \emptyset$, i.e., F is true in all preferred models of T. If $m < \infty$, we conduct a final NP-oracle call in which we guess an interpretation $\mathcal{I}$ and ask whether $max\{deg_{\mathcal{L}}(\mathcal{I}, A) \mid A \in T\} = m$ and whether $\mathcal{I} \not\models F$. This procedure runs in polynomial time, making use of $O(\log(opt_{\mathcal{L}}(T)))$ NP-oracle calls. If $opt_{\mathcal{L}}(F) \in O(|F|^c)$ for every $\mathcal{L}$-formula F, then we require only $O(\log(|F'|))$ NP-oracle calls, where F' is the formula in T with the highest optionality, which means we have Θ_2^{P}-membership in these cases. □

Completeness results for specific choice logics follow directly by the above proposition as well as from known results for single formulas (cf. Table 2.5).

Proposition 5.16. *$\mathcal{L}$-Entailment$[mm]$ is Θ_2^{P}-complete for $\mathcal{L} \in \{\text{QCL}, \text{CCL}\}$ and Δ_2^{P}-complete for $\mathcal{L} = \text{LCL}$.*

Proof. Membership is established in Proposition 5.15. Hardness follows from the Θ_2^{P}-hardness of $\mathcal{L}$-PMContainment for $\mathcal{L} \in \{\text{QCL}, \text{CCL}\}$ and from the Δ_2^{P}-hardness of $\mathcal{L}$-PMContainment for $\mathcal{L} = \text{LCL}$ (cf. Table 2.5). Specifically, an instance (F, a) of $\mathcal{L}$-PMContainment can be reduced to an instance $(\{F\}, \neg a)$ of co-$\mathcal{L}$-Entailment$[mm]$. Then a is true in some preferred model $\mathcal{I}$ of F if and only if $deg_{\mathcal{L}}(\mathcal{I}, \neg a) = \infty$ for some $\mathcal{I} \in Prf_{\mathcal{L}}^{mm}(\{F\})$ iff $\neg a$ is not entailed by $\{F\}$ under the minmax preferred model semantics. □

As the minmax semantics is equivalent to taking the conjunction of all formulas in the theory, it is not surprising that we see the same complexity landscape as for $\mathcal{L}$-

PMCONTAINMENT. This is different for $\mathcal{L}$-ENTAILMENT$[lex]$, which cannot be represented in terms of single formulas.

Proposition 5.17. *$\mathcal{L}$-ENTAILMENT[lex] is in Δ_2^{P} and* coNP*-hard for every tractable choice logic $\mathcal{L}$.*

Proof. The coNP-hardness of $\mathcal{L}$-ENTAILMENT$[lex]$ follows from the coNP-hardness for $\mathcal{L} = \text{PL}$. We now show Δ_2^{P}-Membership of co-$\mathcal{L}$-ENTAILMENT$[lex]$: given a theory $T = \{A_1, \ldots, A_n\}$ and a classical formula F, we first conduct a binary search over $(1, \ldots, opt_{\mathcal{L}}(T))$ to find the smallest k_1 such that some formula $A \in T$ is satisfied by a model of T to a degree of k_1. We then conduct a binary search over $(1, \ldots, n)$ to find the maximum number m_1 such that $|\mathcal{I}_{\mathcal{L}}^{k_1}(T)| = m_1$ for some $\mathcal{I} \in Mod_{\mathcal{L}}(T)$. Observe that any $\mathcal{J} \in Prf_{\mathcal{L}}^{lex}(T)$ must satisfy exactly m_1 formulas in T to a degree of k_1. The above procedure makes $O(\log(opt_{\mathcal{L}}(T)) + \log(n))$ NP-oracle calls.

We proceed inductively: for $i > 1$, we conduct a binary search over $(k_{i-1}+1, \ldots, opt_{\mathcal{L}}(T))$ to find the minimum degree k_i with $k_i > k_{i-1}$ such that $deg_{\mathcal{L}}(\mathcal{I}, A) = k_i$ for some $A \in T$ and $\mathcal{I} \in Mod_{\mathcal{L}}(T)$ with $|\mathcal{I}_{\mathcal{L}}^{k_j}(T)| = m_j$ for all $j < i$. Then we conduct a binary search over $(1, \ldots, n - \sum_{j=1}^{i-1} m_j)$ to find the maximum number m_i such that $|\mathcal{I}_{\mathcal{L}}^{k_i}(T)| = m_i$ for some $\mathcal{I} \in Mod_{\mathcal{L}}(T)$ with $|\mathcal{I}_{\mathcal{L}}^{k_j}(T)| = m_j$ for all $j < i$. Again, this requires $O(\log(opt_{\mathcal{L}}(T)) + \log(n))$ NP-oracle calls.

The above procedure has to be executed at most n times to find the 'degree-profile' for preferred models of T, i.e., every preferred model of T must satisfy exactly m_j formulas in T to a degree of k_j. Thus, $O(n \cdot (\log(opt_{\mathcal{L}}(T)) + \log(n)))$ NP-oracle calls are required so far.

Lastly, we make a final NP-oracle call to guess an interpretation $\mathcal{I}$ and, using k_j and m_j, check whether $\mathcal{I} \in Prf_{\mathcal{L}}^{lex}(T)$ and $\mathcal{I} \not\models F$. □

Crucially, $\mathcal{L}$-ENTAILMENT$[lex]$ is Δ_2^{P}-hard for all logics considered here. We show this via a reduction from LEXMAXSAT (Creignou, Pichler, and Woltran 2018).

Definition 5.18. *LEXMAXSAT is the decision problem where, given a* PL*-formula F and an ordering $a_1 > \cdots > a_n$ over all variables in F, we ask whether a_n is true in the lexicographically largest model of F.*

Proposition 5.19. *$\mathcal{L}$-ENTAILMENT[lex] is Δ_2^{P}-complete for $\mathcal{L} \in \{\text{QCL}, \text{CCL}, \text{LCL}\}$.*

Proof. Δ_2^{P}-membership for each $\mathcal{L} \in \{\text{QCL}, \text{CCL}, \text{LCL}\}$ is established in Proposition 5.17. Hardness for $\mathcal{L} = \text{LCL}$ follows from LCL-PMCONTAINMENT (cf. Table 2.5 and the proof of Proposition 5.16). We now show Δ_2^{P}-hardness for the complementary problem and $\mathcal{L} \in \{\text{QCL}, \text{CCL}\}$: consider an instance $(F, (a_1 \ldots, a_n))$ of LEXMAXSAT. For every $1 \leq i \leq n$ we construct an $\mathcal{L}$-formula A_i such that, for any interpretation $\mathcal{I}$, $deg_{\mathcal{L}}(\mathcal{I}, A_i) = i$ if $a_i \in \mathcal{I}$ and $deg_{\mathcal{L}}(\mathcal{I}, a_i) = n + 1$ if $a_i \notin \mathcal{I}$:

- For $\mathcal{L} = \text{QCL}$ we realize this with

$$A_i = \bot \vec{\times} \ldots \vec{\times} \bot \vec{\times} a_i \vec{\times} \bot \vec{\times} \ldots \vec{\times} \bot \vec{\times} \top,$$

  where $\bot$ occures $i-1$ times before a_i and $n-i$ times after a_i. For example, if $n = 4$, then $A_1 = a_1 \vec{\times} \bot \vec{\times} \bot \vec{\times} \bot \vec{\times} \top$, $A_2 = \bot \vec{\times} a_2 \vec{\times} \bot \vec{\times} \bot \vec{\times} \top$, and so on.

- For $\mathcal{L} = \text{CCL}$ we construct

$$A_i = \top \vec{\odot} a_i \vec{\odot} \ldots \vec{\odot} a_i \vec{\odot} \bot \vec{\odot} \ldots \vec{\odot} \bot,$$

  where a_i occurs $n-i+1$ times and $\bot$ occurs $i-1$ times. For example, if $n = 4$, then $A_1 = \top \vec{\odot} a_1 \vec{\odot} a_1 \vec{\odot} a_1 \vec{\odot} a_1$, $A_2 = \top \vec{\odot} a_2 \vec{\odot} a_2 \vec{\odot} a_2 \vec{\odot} \bot$, and so on.

We now construct an instance $(T, \neg a_n)$ of co-$\mathcal{L}$-Entailment[lex] with $T = \{F, A_1, \ldots, A_n\}$. It remains to show that a_n is true in the lexicographically maximal model of F with respect to $a_1 > \cdots > a_n$ if and only if $\neg a_n$ is false in some lexicographically preferred model of T.

Assume a_n is true in the lexicographically maximal model $\mathcal{I}$ of F with respect to $a_1 > \cdots > a_n$. Since each A_i is always satisfied to some degree, and since $\mathcal{I} \models F$, we have $\mathcal{I} \in \mathit{Mod}_{\mathcal{L}}(T)$. Moreover, there can be no other $\mathcal{J} \in \mathit{Mod}_{\mathcal{L}}(T)$ such that, for some k, $|\mathcal{I}^k_{\mathcal{L}}(T)| < |\mathcal{J}^k_{\mathcal{L}}(T)|$ and, for all $l < k$, $|\mathcal{I}^l_{\mathcal{L}}(T)| = |\mathcal{J}^l_{\mathcal{L}}(T)|$. Otherwise, there would be some variable a_k such that $\mathcal{J} \models a_k$, $\mathcal{I} \not\models a_k$, and $\mathcal{J} \models a_l \iff \mathcal{I} \models a_l$ for all $l < k$. But then $\mathcal{I}$ would not be the lexicographically maximal model of F. Thus, $\mathcal{I} \in \mathit{Prf}^{lex}_{\mathcal{L}}(T)$. Furthermore, $\mathcal{I} \not\models \neg a_n$.

Assume a_n is not true in the lexicographically maximal model of F. If F is not satisfiable, then neither is T, i.e., T has no preferred model and we have a no-instance of co-$\mathcal{L}$-Entailment[lex]. If F is satisfiable, then, by the same argument as above, $\mathcal{I} \in \mathit{Prf}^{lex}_{\mathcal{L}}(T)$ for the lexicographically largest model $\mathcal{I}$ of F with respect to $a_1 > \cdots > a_n$. In fact, $\mathcal{I}$ is the unique lexicographically preferred model of T. Furthermore, $\mathcal{I} \models \neg a_n$. □

The proofs for log-lexicographic semantics are similar to those of regular lexicographic semantics. However, the complexity for choice logics with polynomially-bounded optionality is actually located *inbetween* Θ^{P}_2 and Δ^{P}_2, namely in $\Delta^{\mathsf{P}}_2[O(\log^2 n)]$. The only truly natural $\Delta^{\mathsf{P}}_2[O(\log^2 n)]$-complete problem we are aware of is model checking for a specific temporal logic (Schnoebelen 2003). A less natural, but useful $\Delta^{\mathsf{P}}_2[O(\log^2 n)]$-complete problem is Log²LexMaxSat, which is defined analogously to LexMaxSat (cf. Definition 5.18) and LogLexMaxSat (cf. Section 2.2) except that we are given a lexicographic order over $\log^2(n)$ variables (Segoufin and ten Cate 2013).

Proposition 5.20. *$\mathcal{L}$-Entailment[log] is in Δ^{P}_2 and* coNP*-hard for all tractable choice logics. $\mathcal{L}$-Entailment[log] is in $\Delta^{\mathsf{P}}_2[O(\log^2 n)]$ for tractable $\mathcal{L}$ if for some constant c and all $\mathcal{L}$-formulas F it holds that $opt_{\mathcal{L}}(F) \in O(|F|^c)$.*

Proof. The coNP-hardness of $\mathcal{L}$-Entailment[*log*] follows from coNP-hardness for $\mathcal{L} = \mathrm{PL}$. Regarding membership, we can determine the degree-profile of T's log-lexicographically preferred models analogously to the proof of Proposition 5.17, except that we do not need to execute the binary searches n times. In fact, by the definition of the log-lexicographic preferred model semantics, we need to execute the binary searches only $O(\log(n))$ times, giving us a decision procedure that in total makes use of $O(\log(n) \cdot (\log(opt_{\mathcal{L}}(T)) + \log(n)))$ NP-oracle calls. □

Definition 5.21. *Log²LexMaxSat is the decision problem where, given a* PL*-formula F and an ordering $a_1 > \cdots > a_j$ over $l = \log(n)^2$ of the n variables in F, we ask whether a_l is true in the lexicographically largest interpretation $\mathcal{J} \subseteq \{a_1, \ldots, a_l\}$ that can be extended to a model of F.*

Proposition 5.22. *$\mathcal{L}$-Entailment[log] is $\Delta_2^{\mathsf{P}}[O(\log^2 n)]$-complete for $\mathcal{L} \in \{$*QCL, CCL*$\}$ and Δ_2^{P}-complete for $\mathcal{L} =$* LCL*.*

Proof. Membership for each $\mathcal{L} \in \{\mathrm{QCL}, \mathrm{CCL}, \mathrm{LCL}\}$ is established in Proposition 5.17. Hardness for LCL follows from LCL-PMContainment (cf. Table 2.5 and the proof of Proposition 5.16). We show hardness for co-$\mathcal{L}$-Entailment[*log*] and $\mathcal{L} = \mathrm{QCL}$ by a reduction from Log²LexMaxSat, i.e., we are given a classical formula F over variables $X = \{x_1, \ldots, x_n\}$ and an ordering $x_1 > \ldots > x_l$ over $l = \log^2(n)$ variables. Observe that we cannot inspect all $2^{\log(n)^2}$ interpretations over the variables $\{x_1, \ldots, x_l\}$ as this would not be polynomial in n. Instead, we break up the ordering $x_1 > \ldots > x_l$ into $\log(n)$ parts and, for each part, inspect the relevant $2^{\log(n)} = n$ interpretations. Formally, for every $1 \leq i \leq \log(n)$, let $X_i = \{x_{(i-1)\cdot\log(n)+1}, \ldots, x_{i\cdot\log(n)}\}$ and let $\mathcal{J}_k^i \subseteq X_i$ be the k-th largest interpretation with respect to the ordering $x_{(i-1)\cdot\log(n)+1} > \cdots > x_{i\cdot\log(n)}$. We then construct the formula

$$\varphi(\mathcal{J}_k^i) = \Big(\bigwedge_{x \in \mathcal{J}_k^i} x\Big) \wedge \Big(\bigwedge_{x \in \{X_i \setminus \mathcal{J}_k^i\}} \neg x\Big).$$

Then, for every $1 \leq i \leq \log(n)$, we construct

$$A_i = \varphi(\mathcal{J}_1^i) \vec{\times} \cdots \vec{\times} \varphi(\mathcal{J}_{2^{\log(n)}}^i).$$

Note that, for any interpretation $\mathcal{I} \subseteq \{x_1, \ldots, x_n\}$, $deg_{\mathcal{L}}(\mathcal{I}, A_i) = k$ iff $\mathcal{I}$ is the lexicographically k-th largest interpretation with respect to the ordering $x_{(i-1)\cdot\log(n)+1} > \cdots > x_{i\cdot\log(n)}$. Moreover, observe that $|\varphi(\mathcal{J}_k^i)| \in O(\log(n))$ and that therefore $|A_i| \in O(\log(n) \cdot 2^{\log(n)}) = O(\log(n) \cdot n)$. Now, for every $1 \leq i \leq \log(n)$, we construct

$$B_i = \bot \vec{\times} \cdots \vec{\times} \bot \vec{\times} A_i,$$

where $\bot$ appears $(\log(n) - i) \cdot 2^{\log(n)}$ times before A_i. Observe that $|B_i|$ is polynomial in $|A_i|$ since $|B_i| \in O(\log(n) \cdot 2^{\log(n)} + |A_i|) = O(\log(n) \cdot n + |A_i|)$. Lastly, we construct

$n - \log(n) - 1$ formulas C_j such that $deg_{\mathcal{L}}(\mathcal{J}, C_j) = 1$ for all interpretations $\mathcal{J}$. We can now define our $\mathcal{L}$-theory

$$T = \{F, B_1, \ldots, B_{\log(n)}, C_1, \ldots, C_{n-\log(n)-1}\}.$$

Observe that T contains exactly n formulas. Moreover, $B_{\log(n)}$ can be satisfied to a degree between 1 and $2^{\log(n)}$, $B_{\log(n)-1}$ to a degree between $2^{\log(n)} + 1$ and $2 \cdot 2^{\log(n)}$, and so on until finally B_1 can be satisfied to a degree between $(\log(n) - 1) \cdot 2^{\log(n)} + 1$ and $\log(n) \cdot 2^{\log(n)}$. The formulas C_j are always satisfied to a degree of 1 and therefore do not influence the log-lexicographically preferred models of T. Furthermore, the lexicographically largest model $\mathcal{I}$ of F with respect to the ordering $x_1 > \cdots > x_l$ satisfies B_1 optimally among the models of F, and is therefore preferred in T to all models of F that are lexicographically smaller regarding $x_1 > \cdots > x_{\log(n)}$. $\mathcal{I}$ also satisfies B_2 optimally among those models of F that satisfy B_1 optimally. By extension, we can conclude that an interpretation $\mathcal{J}$ is a lexicographically maximal model of F with respect to $x_1 > \cdots > x_l$ if and only if $\mathcal{J} \in Prf^{log}_{\mathcal{L}}(T)$. □

Lastly, we examine inclusion-based semantics. Here, we cannot show Δ^{P}_2-membership in general. Indeed, as it turns out $\mathcal{L}$-Entailment[*inc*] is Π^{P}_2-complete for all studied choice logic except PL.

Proposition 5.23. *$\mathcal{L}$-Entailment[inc] is in Π^{P}_2 and coNP-hard for every tractable choice logic $\mathcal{L}$.*

Proof. Note that coNP-hardness follows from the fact that the problem is coNP-hard already for $\mathcal{L} = \mathrm{PL}$. Regarding Π^{P}_2-membership, it is straightforward to see that the complementary problem of $\mathcal{L}$-Entailment[*inc*] is in Σ^{P}_2: given a theory $T = \{A_1, \ldots, A_n\}$ and a classical formula F, we guess an interpretation $\mathcal{I}$ and, in coNP, check whether $\mathcal{I} \in Prf^{inc}_{\mathcal{L}}(T)$ and $\mathcal{I} \not\models F$. □

To show Π^{P}_2-completeness for specific choice logics, we can make use of an already existing translation from propositional circumscription (McCarthy 1980) to QCL (Brewka, Benferhat, and Berre 2004, Proposition 10). In fact, this existing translation starts from *prioritized* circumscription. Our construction, however, considers unprioritized circumscription and is therefore slightly simpler. Note that entailment for propositional circumscription is known to be Π^{P}_2-complete (Eiter and Gottlob 1993).

Definition 5.24. *Let T be a classical propositional theory, F a classical formula, and $(P; R)$ a circumscription policy, where P are atoms to be minimized, and R are fixed atoms with $P \cap R = \emptyset$. A model $\mathcal{I}$ of T is $(P; R)$-minimal for T if there is no other model $\mathcal{J}$ of T such that $\mathcal{I} \cap R = \mathcal{J} \cap R$ and $\mathcal{J} \cap P \subset \mathcal{I} \cap P$.*

CircEntailment is the decision problem where, given a a classical theory T, a classical formula F, and a circumscription policy $(P; R)$, we ask whether F is true in all $(P; R)$-minimal models of T.

Table 5.1: Complexity of choice logic theories ($\sigma \in \{mm, lex, inc, log\}$).

	Tract.	PL	QCL/CCL	LCL
$\mathcal{L}$-PMChecking[σ]	in coNP	in P	coNP-c	coNP-c
$\mathcal{L}$-Entailment[mm]	coNP-h/in Δ_2^{P}	coNP-c	Θ_2^{P}-c	Δ_2^{P}-c
$\mathcal{L}$-Entailment[log]	coNP-h/in Δ_2^{P}	coNP-c	$\Delta_2^{\mathsf{P}}[O(\log^2 n)]$-c	Δ_2^{P}-c
$\mathcal{L}$-Entailment[lex]	coNP-h/in Δ_2^{P}	coNP-c	Δ_2^{P}-c	Δ_2^{P}-c
$\mathcal{L}$-Entailment[inc]	coNP-h/in Π_2^{P}	coNP-c	Π_2^{P}-c	Π_2^{P}-c

Proposition 5.25. *$\mathcal{L}$-Entailment[inc] is Π_2^{P}-complete for $\mathcal{L} \in \{\mathrm{QCL}, \mathrm{CCL}, \mathrm{LCL}\}$.*

Proof. Membership for each $\mathcal{L} \in \{\mathrm{QCL}, \mathrm{CCL}, \mathrm{LCL}\}$ follows from Proposition 5.23. We now show hardness for $\mathcal{L} = \mathrm{QCL}$. Consider an arbitrary instance $(T, F, (P; R))$ of CircEntailment. Let

$$T' = T \cup \{\neg p \overrightarrow{\times} p \mid p \in P\} \cup \{r \overrightarrow{\times} \neg r, \neg r \overrightarrow{\times} r \mid r \in R\}.$$

It remains to show that F is true in all $(P; R)$-minimal models of T if and only if F is true in all $Prf_{\mathcal{L}}^{inc}(T')$. We show this by proving that $\mathcal{I}$ is a $(P; R)$-minimal models of T if and only if $\mathcal{I} \in Prf_{\mathcal{L}}^{inc}(T')$.

Assume $\mathcal{I}$ is a $(P; R)$-minimal model of T. Then definitely $\mathcal{I} \in Mod_{\mathrm{QCL}}(T')$, since all formulas in $T' \setminus T$ are always satisfied, either to a degree of 1 or 2. Let $\mathcal{J}$ be any other model of T'. For any $r \in R$, if $r \in \mathcal{I}$ but $r \notin \mathcal{J}$, then $\mathcal{I}$ and $\mathcal{J}$ are incomparable with respect to the *inc*-semantics since $deg_{\mathrm{QCL}}(\mathcal{I}, r \overrightarrow{\times} \neg r) = 1$, $deg_{\mathrm{QCL}}(\mathcal{J}, r \overrightarrow{\times} \neg r) = 2$, $deg_{\mathrm{QCL}}(\mathcal{I}, \neg r \overrightarrow{\times} r) = 2$, and $deg_{\mathrm{QCL}}(\mathcal{J}, \neg r \overrightarrow{\times} r) = 1$. Likewise if $r \notin \mathcal{I}$ but $r \in \mathcal{J}$. Thus, assume $r \in \mathcal{I} \iff r \in \mathcal{J}$ for all $r \in R$. Then $\mathcal{I} \cap P \subseteq \mathcal{J} \cap P$, since $\mathcal{I}$ is $(P; R)$-minimal. Therefore, for all $p \in P$, $deg_{\mathrm{QCL}}(\mathcal{I}, \neg p \overrightarrow{\times} p) \leq deg_{\mathrm{QCL}}(\mathcal{J}, \neg p \overrightarrow{\times} p)$. We can conclude that $\mathcal{I} \in Prf_{\mathrm{QCL}}^{inc}(T')$.

Assume $\mathcal{I}$ is not a $(P; R)$-minimal model of T. If T is not satisfiable, then neither is T', and we are done. If T is satisfiable, then there must be an interpretation $\mathcal{J}$ with $\mathcal{J} \cap R = \mathcal{I} \cap R$ and $\mathcal{J} \cap P \subset \mathcal{I} \cap P$. By the same argument as above, for all $p \in P$, $deg_{\mathrm{QCL}}(\mathcal{J}, \neg p \overrightarrow{\times} p) \leq deg_{\mathrm{QCL}}(\mathcal{I}, \neg p \overrightarrow{\times} p)$. Furthermore, for at least one $q \in P$, $deg_{\mathrm{QCL}}(\mathcal{J}, \neg q \overrightarrow{\times} q) < deg_{\mathrm{QCL}}(\mathcal{I}, \neg q \overrightarrow{\times} q)$. Thus, $\mathcal{I} \notin Prf_{\mathrm{QCL}}^{inc}(T')$. □

Table 5.1 summarizes our complexity results for choice logic theories. Maybe the most interesting point here is that in entailment for QCL and CCL the complexity rises when going from minmax to (log-)lexicographic semantics. However, for LCL, all three problems are equally hard. Thus, the additional expressiveness of the (log-)lexicographic semantics makes entailment harder for choice logics with polynomially bounded optionality. For inclusion-based semantics we see an additional jump in complexity to Π_2^{P}-completeness for all considered choice logics (except PL). We observe that there are two ways in which the complexity of $\mathcal{L}$-Entailment$[\sigma]$ is determined: on the one hand by the choice logic (e.g. QCL vs. LCL), and on the other hand by the preferred model semantics (e.g. minmax vs. lexicographic vs. inclusion-based). Furthermore, with $\mathcal{L}$-Entailment$[\mathit{log}]$ ($\mathcal{L} \in \{\mathrm{QCL}, \mathrm{CCL}\}$) we introduced fairly natural $\Delta_2^{\mathsf{P}}[O(\log^2 n)]$-complete problems.[9]

5.4 Sequent Calculi

In this section we propose the first sound and complete calculus for preferred model entailment in choice logics. We first introduce a calculus for QCL and the minmax, lexicographic, and inclusion-based preferred model semantics. We then show how this calculus can be adapted for CCL and LCL.

Although preferred model entailment in choice logics is related to other forms of non-monotonic entailment for which proof calculi have been developed (Bonatti and Olivetti 2002; Geibinger and Tompits 2020), choice logics are different from other non-monotonic logics in the way non-monotonicity is introduced. Specifically, the non-standard part of our logics are the choice connectives, such as ordered disjunction, which are fully embedded in the logical language. This has to be kept in mind when designing our calculus. Indeed, our calculi for choice logics will differ from most other calculi for non-monotonic logics: our calculi do not use non-standard inference rules as in default logic (Reiter 1980), modal operators expressing consistency or belief as in autoepistemic logic (Moore 1985), or predicates whose extensions are minimized as in circumscription (McCarthy 1980). However, one method that can also be applied to choice logics is the use of a refutation calculus (also known as rejection or antisequent calculus) axiomatising invalid formulas, i.e., non-theorems (Bonatti 1993; Łukasiewicz 1951; Słupecki, Bryll, and Wybraniec-Skardowska 1971; Tiomkin 1988). Specifically, by combining a refutation calculus with an appropriate sequent calculus, elegant proof systems for the central non-monotonic formalisms such as default logic, autoepistemic logic, circumscription, and paraconsistent logics were obtained (Bonatti and Olivetti 2002; Geibinger and Tompits 2020).

Another aspect of choice logics semantics we must account for are satisfaction degrees and their similarity to many-valued logics. There are several kinds of sequent calculus systems for many-valued logics, where the representation as a hypersequent calculus (Avron 1996; Geibinger and Tompits 2020) plays a prominent role. However, there are crucial differences between choice logics and many-valued logics in the usual sense. Firstly, choice logic interpretations are classical, i.e., they set propositional variables to either true or false. Secondly, non-classical satisfaction degrees only arise when choice connectives, e.g. ordered disjunction in QCL, occur in a formula. Thirdly, when applying a choice connective $\circ$ to two formulas A and B, the degree of $A \circ B$ does not only depend on the

[9] We consider $\mathcal{L}$-Entailment$[\mathit{log}]$ to be at least as natural as the Θ_2^{P}-complete LogLexMaxSat and certainly more natural than Log^{2}LexMaxSat.

degrees of A and B, but also on their optionality. Therefore, techniques used in proof systems for conventional many-valued logics cannot be applied directly to choice logics.

In (Governatori and Rotolo 2006) a sequent calculus based system for reasoning with contrary-to-duty obligations was introduced, where a non-classical connective was defined to capture the notion of reparational obligation, which is in force only when a violation of a norm occurs. This is related to the ordered disjunction in QCL, however, based on the intended use in (Governatori and Rotolo 2006) the system was defined only for the occurrence of the new connective on the right side of the sequent sign. We aim for a proof system for reasoning with choice logic operators, and to deduce formulas from choice logic formulas. Thus, we need a calculus with left and right inference rules.

To obtain such a calculus we combine the idea of a refutation calculus with methods developed for multi-valued logics. First, in Subsection 5.4.1 we develop a (monotonic) sequent calculus for reasoning about satisfaction degrees in QCL, using a labeled calculus, a method developed for (finite) many-valued logics (Baaz, Lahav, and Zamansky 2013; Carnielli 1987; Kaminski and Francez 2021). Secondly, in Subsection 5.4.2 we define a labeled refutation calculus for reasoning about invalidity in terms of satisfaction degrees in QCL. Finally, also in Subsection 5.4.2, we introduce a new, non-monotonic inference rule that joins the two previously introduced labeled calculi to obtain a sequent calculus for preferred model entailment in QCL. In Subsection 5.4.3 we show how the calculi for QCL can be adapted for CCL and LCL.

5.4.1 The Sequent Calculus L[QCL]

Towards the development of a calculus for preferred model entailment, we first propose a labeled calculus for reasoning about the satisfaction degrees of QCL formulas in sequent format and prove its soundness and completeness. An advantage of the sequent calculus format is the possibility to have symmetrical left and right rules for all connectives, in particular for the choice connectives. This is in contrast to the representation of ordered disjunction in the calculus for deontic logic (Governatori and Rotolo 2006), in which only right-hand side rules are considered.

As the calculus will be concerned with satisfaction degrees rather than preferred models, entailment will be defined in terms of satisfaction degrees. To this end, the formulas occurring in the sequents of the calculus will be labeled.

Definition 5.26 (labeled QCL-formulas)**.** *Let A be a* QCL*-formula and $k \in \mathbb{N}$, then $(A)_k$ is a labeled* QCL*-formula. The labeled* QCL*-formula $(A)_k$ is satisfied by those interpretations that satisfy A to a degree of k.*

Instead of labeling formulas with degree ∞ we will use the negated formula, i.e., instead of $(A)_\infty$ we will use $(\neg A)_1$. Observe that $(A)_k$ for $opt_{\mathcal{L}}(A) > k$ can never have a model. We will deal with such formulas by replacing them with $(\bot)_1$. For classical formulas, we may write A for $(A)_1$.

Definition 5.27 (labeled QCL-sequents)**.** *Let* $(A_1)_{k_1}, \ldots, (A_m)_{k_m}$ *and* $(B_1)_{l_1}, \ldots, (B_n)_{l_n}$ *be labeled* QCL*-formulas. Then*

$$(A_1)_{k_1}, \ldots, (A_m)_{k_m} \vdash (B_1)_{l_1}, \ldots, (B_n)_{l_n}$$

is a labeled QCL*-sequent.*
$(A_1)_{k_1}, \ldots, (A_m)_{k_m} \vdash (B_1)_{l_1}, \ldots, (B_n)_{l_n}$ *is valid if and only if every interpretation that satisfies all labeled* QCL*-formulas* $(A_1)_{k_1}, \ldots, (A_m)_{k_m}$ *to the degree specified by the label also satisfies at least one labeled* QCL*-formula out of* $(B_1)_{l_1}, \ldots, (B_n)_{l_n}$ *to the degree specified by the label.*

In contrast to preferred model entailment, the entailment in terms of satisfaction degrees, as defined above, is monotonic.

Frequently we will write $(A)_{<k}$ as shorthand for the sequence of labeled QCL-formulas

$$(A)_1, \ldots, (A)_{k-1}$$

and $(A)_{>k}$ for the sequence of labeled QCL-formulas

$$(A)_{k+1}, \ldots, (A)_{opt_{\text{QCL}}(A)}, (\neg A)_1.$$

Moreover, $\langle \Gamma, (A)_i \vdash \Delta \rangle_{i<k}$ will denote the sequence of labeled QCL-sequents

$$\Gamma, (A)_1 \vdash \Delta \ \ldots \ \Gamma, (A)_{k-1} \vdash \Delta.$$

Analogously, $\langle \Gamma, (A)_i \vdash \Delta \rangle_{i>k}$ will denote the sequence of labeled QCL-sequents

$$\Gamma, (A)_{k+1} \vdash \Delta \ \ldots \ \Gamma, (A)_{opt_{\text{QCL}}(A)} \vdash \Delta \quad \Gamma, (\neg A)_1 \vdash \Delta.$$

Below we define the sequent calculus **L**[QCL] over labeled QCL-sequents. In addition to introducing inference rules for the choice connective $\vec{\times}$ we have to modify the inference rules for conjunction and disjunction of propositional **LK**. We first state the calculus, and then explain the intuition behind the rules.

Definition 5.28 (**L**[QCL])**.** *The axioms of* **L**[QCL] *are labeled* QCL*-sequents* $\Gamma \vdash \Delta$ *such that* $\bot \in \Gamma$ *or such that* $p \in \Gamma$ *and* $p \in \Delta$ *for some atom* p*. The inference rules are given below. Whenever a labeled* QCL*-formula* $(F)_k$ *appears in the conclusion of an inference rule (except for the dol- and dor-rules) it holds that* $k \leq opt_{\mathcal{L}}(F)$*.*

The rules for the classical connectives are

$$\frac{\Gamma \vdash (cp(A))_1, \Delta}{\Gamma, (\neg A)_1 \vdash \Delta} \neg l \qquad\qquad \frac{\Gamma, (cp(A))_1 \vdash \Delta}{\Gamma \vdash (\neg A)_1, \Delta} \neg r$$

$$\frac{\Gamma, (A)_k \vdash (B)_{<k}, \Delta \qquad \Gamma, (B)_k \vdash (A)_{<k}, \Delta}{\Gamma, (A \lor B)_k \vdash \Delta} \lor l$$

$$\frac{\langle \Gamma, (A)_i \vdash \Delta \rangle_{i<k} \quad \langle \Gamma, (B)_i \vdash \Delta \rangle_{i<k} \quad \Gamma \vdash (A)_k, (B)_k, \Delta}{\Gamma \vdash (A \vee B)_k, \Delta} \vee r$$

$$\frac{\Gamma, (A)_k \vdash (B)_{>k}, \Delta \qquad \Gamma, (B)_k \vdash (A)_{>k}, \Delta}{\Gamma, (A \wedge B)_k \vdash \Delta} \wedge l$$

$$\frac{\langle \Gamma, (A)_i \vdash \Delta \rangle_{i>k} \quad \langle \Gamma, (B)_i \vdash \Delta \rangle_{i>k} \quad \Gamma \vdash (A)_k, (B)_k, \Delta}{\Gamma \vdash (A \wedge B)_k, \Delta} \wedge r$$

The rules for ordered disjunction are:

$$\frac{\Gamma, (A)_k \vdash \Delta}{\Gamma, (A \overrightarrow{\times} B)_k \vdash \Delta} \overrightarrow{\times} l_1 \qquad \frac{\Gamma, (B)_l, (\neg A)_1 \vdash \Delta}{\Gamma, (A \overrightarrow{\times} B)_{opt_{\mathrm{QCL}}(A)+l} \vdash \Delta} \overrightarrow{\times} l_2$$

$$\frac{\Gamma \vdash (A)_k, \Delta}{\Gamma \vdash (A \overrightarrow{\times} B)_k, \Delta} \overrightarrow{\times} r_1 \qquad \frac{\Gamma \vdash (\neg A)_1, \Delta \qquad \Gamma \vdash (B)_l, \Delta}{\Gamma \vdash (A \overrightarrow{\times} B)_{opt_{\mathrm{QCL}}(A)+l}, \Delta} \overrightarrow{\times} r_2$$

where $k \leq opt_{\mathcal{L}}(A)$ *and* $l \leq opt_{\mathcal{L}}(B)$.

The degree overflow rules[10] *are:*

$$\frac{\Gamma, \bot \vdash \Delta}{\Gamma, (A)_{opt_{\mathrm{QCL}}(A)+k} \vdash \Delta} dol \qquad \frac{\Gamma \vdash \Delta}{\Gamma \vdash (A)_{opt_{\mathrm{QCL}}(A)+k}, \Delta} dor$$

where $k \in \mathbb{N}$.

The rules for negation are analogous to propositional **LK**. Note that we replace A by its classical counterpart $cp(A)$ (cf. Definition 2.37). This reflects the fact that negation in choice logics erases all information about preferences, and that we therefore are only interested in the classical satisfaction of A.

The idea behind the $\vee$-left rule is that a model M of the labeled QCL-formula $(A)_k$ is only a model of the labeled QCL-formula $(A \vee B)_k$ if there is no $l < k$ s.t. M is a model of $(B)_l$. Therefore, every model of $(A \vee B)_k$ is a model of Δ if and only if

- every model of $(A)_k$ is a model of Δ or of some $(B)_l$ with $l < k$,
- every model of $(B)_k$ is a model of Δ or of some $(A)_l$ with $l < k$.

Essentially, the same idea works for $\wedge$-left but with $l > k$. For the $\vee$-right rule, in order for every model of Γ to be a model of $(A \vee B)_k$, every model of Γ must either be a model of $(A)_k$ or of $(B)_k$ and no model of Γ can be a model of $(A)_l$ for $l < k$, i.e., $\Gamma, (A)_l \vdash \bot$. Similarly for $\wedge$-right.

[10] *dol*/*dor* stands for degree overflow left/right.

Observe that, in case we are dealing with classical formulas only, the modified inference rules for $\wedge$ and $\vee$ are equivalent to the inference rules for $\wedge$ and $\vee$ of propositional **LK** without structural rules (Troelstra and Schwichtenberg 2000, Section 3.5). Consider the $\vee$-left rule in **L**[QCL]: if A and B are classical, and $k = 1$, the rule equals the $\vee$-left rule of propositional **LK**, as $(A)_{<1}$ is empty. Similarly, the $\vee$-right rule in **L**[QCL] equals the $\vee$-right rule in propositional **LK**, because $\langle \Gamma, (A)_i \vdash \Delta \rangle_{i<1}$ is empty. Moreover, as $(A)_{>1} = \neg A$ for a classical formula A, the $\wedge$-left rule of **L**[QCL] is equivalent to the $\wedge$-left rule of propositional **LK** if A and B are classical formulas and $k = 1$ (but splits the proof-tree unnecessarily). Analogously for $\wedge$-right, as $\langle \Gamma, (A)_i \vdash \Delta \rangle_{i>1}$ equals $\Gamma, \neg A \vdash \Delta$ if A is classical. Therefore, this is equivalent to the $\wedge$-right rule of propositional **LK** if A and B are classical formulas and $k = 1$ (but adds an unnecessary third condition $\Gamma \vdash A, B, \Delta$).

The rules for ordered disjunction follow straightforwardly from QCL-semantics. If $A \vec{\times} B$ is satisfied to a degree $k \leq opt_{\text{QCL}}(A)$, then we know that A must be satisfied to a degree of k. If $A \vec{\times} B$ is satisfied to a finite degree higher than $opt_{\text{QCL}}(A)$, then we know that B is satisfied but A is not.

The intuition behind the degree overflow rules *dol* and *dor* is that we sometimes need to fix sequences in which a labeled QCL-formula F is assigned a label k with $opt_{\text{QCL}}(F) < k < \infty$. This can happen after applying the rules for conjunction/disjunction. For instance, consider the $\wedge l$-rule: as the premise we have $\Gamma, (A \wedge B)_k \vdash \Delta$ with $k \leq opt_{\text{QCL}}(A \wedge B)$. Recall that the optionality of this conjunct is defined as $opt_{\text{QCL}}(A \wedge B) = max(opt_{\text{QCL}}(F), opt_{\text{QCL}}(G))$. Thus, it may be the case that, for example, $opt_{\text{QCL}}(A) < k$. The $\wedge l$ rule, however, will introduce the premise $\Gamma, (A)_k \vdash (B)_{>k}, \Delta$. Since $(A)_k$ can never be satisfied, as $opt_{\text{QCL}}(A) < k < \infty$, we have to apply the *dol*-rule which replaces $(A)_k$ by $\bot$.

We now provide some examples for valid derivations in **L**[QCL] before showing soundness and completeness.

Example 5.29. *The following is an* **L**[QCL]*-proof of a valid sequent.*[11]

$$
\dfrac{
\dfrac{\dfrac{b \vee c, \neg a, b \vdash a \wedge b, a \wedge c, b}{b \vee c, (a \vec{\times} b)_2 \vdash a \wedge b, a \wedge c, b}\;\vec{\times} l_2}{(a \vec{\times} b)_2 \vdash \neg(b \vec{\times} c), a \wedge b, a \wedge c, b}\;\neg r
\qquad
\dfrac{\dfrac{\begin{matrix}\vdots \\ a \vee b, \neg b, c \vdash a \wedge b, a \wedge c, b\end{matrix}}{a \vee b, (b \vec{\times} c)_2 \vdash a \wedge b, a \wedge c, b}\;\vec{\times} l_2}{(b \vec{\times} c)_2 \vdash \neg(a \vec{\times} b), a \wedge b, a \wedge c, b}\;\neg r
}{
\dfrac{((a \vec{\times} b) \wedge (b \vec{\times} c))_2 \vdash a \wedge b, a \wedge c, b}{\neg(a \wedge b), ((a \vec{\times} b) \wedge (b \vec{\times} c))_2 \vdash a \wedge c, b}\;\neg l
}\;\wedge l
$$

[11] Note that, once we reach sequents containing only classical formulas, we do not continue the proof. However, it can be verified that the classical sequents on the left and right branch are provable in this case. Moreover, given a labeled QCL-formula $(A)_1$ with a label of 1, the label is often omitted for readability.

Example 5.30. *The following end-sequent is similar to the end-sequent of Example 5.29, but with the exception that* $(a\overrightarrow{\times}b)\wedge(b\overrightarrow{\times}c)$ *is assigned a label of* 1. *However,* $((a\overrightarrow{\times}b)\wedge(b\overrightarrow{\times}c))_1$ *is unsatisfiable in view of* $\neg(a \wedge b)$.

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\vdots}{b \vee c, a \vdash \neg b, a \wedge b, \bot} \qquad \dfrac{\vdots}{b \vee c, a \vdash c, a \wedge b, \bot}}{b \vee c, a \vdash (b\overrightarrow{\times}c)_2, a \wedge b, \bot}\,\overrightarrow{\times}r_2}{b \vee c, (a\overrightarrow{\times}b)_1 \vdash (b\overrightarrow{\times}c)_2, a \wedge b, \bot}\,\overrightarrow{\times}l_1}{(a\overrightarrow{\times}b)_1 \vdash (b\overrightarrow{\times}c)_2, \neg(b\overrightarrow{\times}c), a \wedge b, \bot}\,\neg r \qquad (\varphi)}{((a\overrightarrow{\times}b) \wedge (b\overrightarrow{\times}c))_1 \vdash a \wedge b, \bot}\,\wedge l}{\neg(a \wedge b), ((a\overrightarrow{\times}b) \wedge (b\overrightarrow{\times}c))_1 \vdash \bot}\,\neg l$$

where φ *is*

$$\dfrac{\dfrac{\dfrac{\dfrac{\vdots}{a \vee b, b \vdash \neg a, a \wedge b, \bot} \qquad a \vee b, b \vdash b, a \wedge b, \bot}{a \vee b, b \vdash (a\overrightarrow{\times}b)_2, a \wedge b, \bot}\,\overrightarrow{\times}r_2}{a \vee b, (b\overrightarrow{\times}c)_1 \vdash (a\overrightarrow{\times}b)_2, a \wedge b, \bot}\,\overrightarrow{\times}l_1}{(b\overrightarrow{\times}c)_1 \vdash (a\overrightarrow{\times}b)_2, \neg(a\overrightarrow{\times}b), a \wedge b, \bot}\,\neg r$$

Example 5.31. *The following proof shows how the* $\wedge r$*-rule can introduce more than three premises.*

$$\dfrac{\dfrac{a, b \vdash a}{a, b, \neg a \vdash}\,\neg l \qquad \dfrac{\dfrac{a, b, c \vdash b}{a, b, c, \neg b \vdash}\,\neg l}{a, b, (b\overrightarrow{\times}c)_2 \vdash}\,\overrightarrow{\times}l_2 \qquad \dfrac{\dfrac{a, b \vdash b, c}{a, b \vdash b \vee c}\,\vee r}{a, b, \neg(b\overrightarrow{\times}c) \vdash}\,\neg l \qquad a, b \vdash a, (b\overrightarrow{\times}c)_1}{a, b \vdash (a \wedge (b\overrightarrow{\times}c))_1}\,\wedge r$$

Theorem 5.32. $\mathbf{L}[\mathrm{QCL}]$ *is sound and complete.*

Proof (Soundness). We have to prove that all rules of $\mathbf{L}[\mathrm{QCL}]$ are sound.

- For the axioms this is clearly the case.
- $(\neg r)$ and $(\neg l)$: follows from the fact that $deg_{\mathrm{QCL}}(\mathcal{I}, F) < \infty \iff \mathcal{I} \models cp(F)$ for all QCL-formulas F (see Proposition 2.38).
- $(\vee l)$: Assume that the conclusion of the rule is not valid, i.e., there is a model M of Γ and $(A \vee B)_k$ that is not a model of Δ. Then, M satisfies either A or B to degree k and neither to a degree smaller than k. Assume M satisfies A to a degree of k, the other case is symmetric. Then M is a model of Γ and $(A)_k$ but, by assumption, neither of Δ nor of $(B)_j$ for $j < k$. Hence, at least one of the premises is not valid.

- $(\land l)$: Analogous to the proof for $(\lor l)$: assume that the conclusion of the rule is not valid, i.e., that there is a model M of Γ and $(A \land B)_k$ that is not a model of Δ. Then, M satisfies either A or B to degree k and neither to a degree higher than k. Assume M satisfies A to a degree of k, the other case is symmetric. Then M is a model of Γ and $(A)_k$ but, by assumption, neither of Δ nor of $(\neg B)_1$ or $(B)_j$ for any $j > k$. Hence, at least one of the premises is not valid.
- $(\lor r)$: Assume there is a model M of Γ that is not a model of Δ or of $(A \lor B)_k$. There are two possible cases why M is not a model of $(A \lor B)_k$:
  1. M satisfies neither A nor B to degree k. But in this case the premise $\Gamma \vdash (A)_k, (B)_k, \Delta$ is not valid as M is also not a model of Δ by assumption.
  2. M satisfies either A or B to a degree smaller than k. Assume that M satisfies A to degree $j < k$ (the other case is symmetric). Then the premise $\Gamma, (A)_j \vdash \Delta$ is not valid. Indeed, M is a model of Γ and $(A)_j$ but not of Δ.
- $(\land r)$: Analogous to the proof for $(\lor r)$: assume that the conclusion of the rule is not valid, i.e. there is a model M of Γ that is not a model of Δ or of $(A \land B)_k$. There are two possible cases why M is not a model of $(A \land B)_k$:
  1. M satisfies neither A nor B to degree k. However, then the premise $\Gamma \vdash (A)_k, (B)_k, \Delta$ is not valid as M is also not a model of Δ by assumption.
  2. M satisfies either A or B to a degree j higher than k. By symmetry, it suffices to consider the case that M satisfies A to a degree j higher than k. Then either $k < j \leq opt_{\mathrm{QCL}}(A)$ or $j = \infty$. If $k < j \leq opt_{\mathrm{QCL}}(A)$ the premise $\Gamma, (A)_j \vdash \Delta$ is not valid. If $k = \infty$ the premise $\Gamma, (\neg A)_1 \vdash \Delta$ is not valid.
- $(\vec{\times} l_1)$ and $(\vec{\times} r_1)$: follows from the fact that $(A)_k$ has the same models as $(A \vec{\times} B)_k$ for $k \leq opt_{\mathcal{L}}(A)$.
- $(\vec{\times} l_2)$: Assume the conclusion of the rule is not valid and let M be the model witnessing this. Then M is a model of $(A \vec{\times} B)_{opt_{QCL}(A)+l}$. By definition, M satisfies B to degree l and is not a model of A. However, then it is also a model of Γ, $(B)_l$ and $(\neg A)_1$, which means that the premise is not valid.
- $(\vec{\times} r_2)$. Assume that both premises are valid, i.e., every model of Γ is either a model of Δ or of $(\neg A)_1$ and $(B)_l$ with $l \leq opt_{\mathcal{L}}(B)$. Now, by definition, any model that is not a model of A (and hence a model of $(\neg A)_1$) and of $(B)_l$ satisfies $A \vec{\times} B$ to degree $opt_{QCL}(A) + l$. Therefore, every model of Γ is either a model of Δ or of $(A \vec{\times} B)_{opt_{QCL}(A)+l}$, which means that the conclusion of the rule is valid.
- (dol): $\Gamma, \bot$ has no models, i.e., the premise $\Gamma, \bot \vdash \Delta$ is valid. Indeed, the sequent $\Gamma, \bot \vdash \Delta$ is an axiom in our system. Crucially, the sequent $\Gamma, (A)_{opt_{\mathrm{QCL}}(A)+k}$ has no models as well since A cannot be satisfied to a degree m with $opt_{\mathcal{L}}(A) < m < \infty$.
- (dor) is clearly sound. □

Proof (Completeness). To prove completeness, we observe that for any sequent, we can decompose every formula into atomic and hence classical formulas by applying the rules of **L**[QCL]. Moreover, we observe that if all formulas are classical and labeled with 1, then all inference rules reduce to the inference rules of the classical propositional calculus without structural rules (Troelstra and Schwichtenberg 2000, Section 3.5), which is known to be complete. Therefore, we know that a sequent containing only classical formulas is valid if and only if it is provable. It remains to show that the rules of **L**[QCL] preserve validity when read "upwards".

- (dol): Assume that a sequent of the form $\Gamma, (A)_{opt_{\mathrm{QCL}}(A)+k} \vdash \Delta$ with $k \in \mathbb{N}$ is valid. Since $\Gamma, \bot$ has no models, $\Gamma, \bot \vdash \Delta$ is valid.
- (dor): Assume that a sequent $\Gamma \vdash (A)_{opt_{\mathrm{QCL}}(A)+k}, \Delta$ is valid. $(A)_{opt_{\mathrm{QCL}}(A)+k}$ cannot be satisfied, i.e., $\Gamma \vdash \Delta$ is valid.
- $(\neg r)$ and $(\neg l)$: Assume that a sequent of the form $\Gamma \vdash (\neg A)_1, \Delta$ is valid. Then every model of Γ is either a model of $(\neg A)_1$ or of Δ. In other words, every model of Γ that is not a model of $(\neg A)_1$ (i.e., is model of $cp(A)$) is a model of Δ. Therefore, every interpretation that is a model of both Γ and $cp(A)$ must be a model of Δ. It follows that $\Gamma, cp(A) \vdash \Delta$ is valid. Similarly for $\Gamma, (\neg A)_1 \vdash \Delta$.
- $(\vee l)$ and $(\wedge l)$: Assume that a sequent of the form $\Gamma, (A \vee B)_k \vdash \Delta$ is valid, with $k \leq opt_{\mathcal{L}}(A \vee B)$. We claim that then both $\Gamma, (A)_k \vdash (B)_{<k}, \Delta$ and $\Gamma, (B)_k \vdash (A)_{<k}, \Delta$ are valid. Assume to the contrary that $\Gamma, (A)_k \vdash (B)_{<k}, \Delta$ is not valid (the other case is symmetric). Then, there is a model M of Γ and $(A)_k$ that is neither a model of $(B)_{<k}$ nor of Δ. But then M is also a model of Γ and $(A \vee B)_k$, but not of Δ, which contradicts the assumption that $\Gamma, (A \vee B)_k \vdash \Delta$ is valid. Therefore, both $\Gamma, (A)_k \vdash (B)_{<k}, \Delta$ and $\Gamma, (B)_k \vdash (A)_{<k}, \Delta$ are valid. Similarly for a sequent of the form $\Gamma, (A \wedge B)_k \vdash \Delta$.
- $(\vee r)$ and $(\wedge r)$: Assume that a sequent of the form $\Gamma \vdash (A \vee B)_k, \Delta$ is valid, with $k \leq opt_{\mathcal{L}}(A \vee B)$. We claim that then for all $i < k$ the sequents $\Gamma, (A)_i \vdash \Delta$ and $\Gamma, (B)_i \vdash \Delta$ and $\Gamma \vdash (A)_k, (B)_k, \Delta$ are valid. Assume by contradiction that there is an $i < k$ s.t. $\Gamma, (A)_i \vdash \Delta$ is not valid. Then, there is a model M of Γ and $(A)_i$ that is not a model of Δ. However, then M is a model of Γ but neither of Δ nor of $(A \vee B)_k$ (as M satisfies $A \vee B$ to degree $i \neq k$), which contradicts our assumption that $\Gamma \vdash (A \vee B)_k, \Delta$ is valid. The case that there is an $i < k$ s.t. $\Gamma, (B)_i \vdash \Delta$ is not valid is symmetric. Finally, we assume that $\Gamma \vdash (A)_k, (B)_k, \Delta$ is not valid. Then, there is a model M of Γ that is not a model of $(A)_k$, $(B)_k$ or Δ. Thus, M is model of Γ but neither of Δ nor of $(A \vee B)_k$, contradicting our assumption. Therefore, all sequents listed above must be valid. Similarly for a sequent of the form $\Gamma \vdash (A \wedge B)_k, \Delta$.
- $(\vec{\times} l_1)$ and $(\vec{\times} r_1)$: Assume a sequent of the form $\Gamma, (A \vec{\times} B)_k \vdash \Delta$ with $k \leq opt_{\mathrm{QCL}}(A)$ is valid. Then $\Gamma, (A)_k \vdash \Delta$ is also valid since $(A \vec{\times} B)_k$ and $(A)_k$ have the same models if $k \leq opt_{\mathrm{QCL}}(A)$. Analogously for sequents of the form $\Gamma \vdash (A \vec{\times} B)_k, \Delta$.

- ($\vec{\times}l_2$): Assume a sequent of the form $\Gamma, (A\vec{\times}B)_{opt_{\text{QCL}}(A)+l} \vdash \Delta$ is valid, with $l \leq opt_{\mathcal{L}}(B)$. We claim that the sequent $\Gamma, (B)_l, \neg A \vdash \Delta$ is then also valid. Indeed, if M is a model of Γ, $(B)_l$ and $\neg A$, then it is also a model of Γ and $(A\vec{\times}B)_{opt_{\text{QCL}}(A)+l}$. Hence, by assumption, M must be a model of Δ.
- ($\vec{\times}r_2$): Assume that a sequent of the form $\Gamma \vdash (A\vec{\times}B)_{opt_{QCL}(A)+l}, \Delta$ is valid, with $l \leq opt_{\mathcal{L}}(B)$. We claim that then also the sequents $\Gamma \vdash \neg A, \Delta$ and $\Gamma \vdash (B)_l, \Delta$ are valid. Assume by contradiction that the first sequent is not valid. This means that there is a model M of Γ that is not a model of neither $\neg A$ nor of Δ. However, then M is a model of A and therefore satisfies $A\vec{\times}B$ to a degree smaller than $opt_{\text{QCL}}(A)$. This contradicts our assumption that $\Gamma \vdash (A\vec{\times}B)_{opt_{\text{QCL}}(A)+l}, \Delta$ is valid. Assume now that the second sequent is not valid, i.e., that there is a model M of Γ that is neither a model of $(B)_l$ nor of Δ. Then, M cannot be a model of $(A\vec{\times}B)_{opt_{\text{QCL}}(A)+l}$ and we again have a contradiction to the assumption. □

A cut rule has not been introduced for **L**[QCL] so far. However, it is easy to see that **L**[QCL] is cut-admissible.

$$\frac{\Gamma \vdash (A)_k, \Delta \qquad \Gamma', (A)_k \vdash \Delta'}{\Gamma, \Gamma' \vdash \Delta, \Delta'}\ cut$$

Proposition 5.33. *The cut-rule is sound.*

Proof. Assume $\Gamma \vdash (A)_k, \Delta$ and $\Gamma', (A)_k \vdash \Delta'$ are valid. Let M be some model of Γ, Γ'. M must satisfy some formula in $(A)_k, \Delta$. If M satisfies $(A)_k$ then M satisfies both Γ' and $(A)_k$ and thus also some formula in Δ'. In any case, M satisfies some formula in Δ, Δ'. □

We do not prove an effective cut-elimination theorem in the sense of Gentzen, i.e. by providing an algorithm for the elimination of cut inferences in a derivation. But since we do not use a cut rule when proving the completeness of **L**[QCL] (cf. Theorem 5.32), we obtain a cut-elimination theorem for free.

Another aspect of our calculus that should be mentioned is that, although **L**[QCL] is cut-free, we do not have the subformula property. This is especially obvious when looking at the rules for negation, where we use the classical counterpart $cp(A)$ of QCL-formulas. For example, $\neg(a\vec{\times}b)$ in the conclusion of the $\neg$-left rule becomes $cp(a\vec{\times}b) = a \vee b$ in the premise.

Moreover, note that we introduced no structural rules (i.e., weakening or contraction) in **L**[QCL], as they are not needed for the completeness of the calculus. It is easy to see, however, that weakening and contraction are sound in this setting. Thus, if desired, one could extend **L**[QCL] with the following rules:

$$\frac{\Gamma \vdash \Delta}{\Gamma, (A)_k \vdash \Delta}\ wl \qquad \frac{\Gamma \vdash \Delta}{\Gamma \vdash (A)_k, \Delta}\ wr$$

$$\frac{\Gamma, (A)_k, (A)_k \vdash \Delta}{\Gamma, (A)_k \vdash \Delta}\ cl \qquad \frac{\Gamma \vdash (A)_k, (A)_k, \Delta}{\Gamma \vdash (A)_k, \Delta}\ cr$$

Towards Preferred Model Entailment While we believe that $\mathbf{L}[\mathrm{QCL}]$ is interesting in its own right, the question of how this calculus can be used to obtain a calculus for preferred model entailment arises. Essentially, an inference rule has to be added that allows for the transition from standard to preferred model inferences. As a first approach we consider theories $\Gamma \cup \{A\}$ with Γ consisting only of classical formulas and A being a QCL-formula. In this simple case, preferred models of $\Gamma \cup \{A\}$ are those models of $\Gamma \cup \{A\}$ that satisfy A to the smallest possible degree. We call the resulting calculus $\mathbf{L}[\mathrm{QCL}]^{naive}_{\mid\sim}$.

Definition 5.34 ($\mathbf{L}[\mathrm{QCL}]^{naive}_{\mid\sim}$). *The labeled sequent calculus* $\mathbf{L}[\mathrm{QCL}]^{naive}_{\mid\sim}$ *is* $\mathbf{L}[\mathrm{QCL}]$ *extended by the inference rule*

$$\frac{\langle \Gamma, (A)_i \vdash \bot \rangle_{i<k} \qquad \Gamma, (A)_k \vdash \Delta}{\Gamma, A \mid\sim^{lex}_{\mathrm{QCL}} \Delta}\ \mid\sim_{naive}$$

Intuitively, the inference rule $\mid\sim_{naive}$ states that, if there are no interpretations that satisfy Γ while also satisfying A to a degree lower than k, and if Δ follows from all models of $\Gamma, (A)_k$, then Δ is entailed by the preferred models of $\Gamma \cup \{A\}$. However, it can be shown that $\mathbf{L}[\mathrm{QCL}]^{naive}_{\mid\sim}$ is unsound.

Proposition 5.35. $\mathbf{L}[\mathrm{QCL}]^{naive}_{\mid\sim}$ *is unsound.*

Proof. Consider the invalid entailment $\neg a, a \vec{\times} b \mid\sim^{lex}_{\mathrm{QCL}} a$, which is derivable in $\mathbf{L}[\mathrm{QCL}]^{naive}_{\mid\sim}$ as follows:

$$\cfrac{\cfrac{\neg a, a \vdash a}{\neg a, (a \vec{\times} b)_1 \vdash a}\ \vec{\times} l_1}{\neg a, a \vec{\times} b \mid\sim^{lex}_{\mathrm{QCL}} a}\ \mid\sim_{naive}$$

□

Thus, an extension of $\mathbf{L}[\mathrm{QCL}]$ by $\mid\sim_{naive}$ does not yield the desired calculus, not even in this restricted setting where we consider only a single non-classical formula A. What is missing is an assertion that $\Gamma, (A)_k$ is satisfiable. Unfortunately, this cannot be formulated in $\mathbf{L}[\mathrm{QCL}]$. A way of addressing this problem is to make use of a refutation calculus, as has been done for other non-monotonic logics (Bonatti and Olivetti 2002).

5.4.2 A Calculus for Preferred Model Entailment

To obtain a calculus for preferred model entailment, we first need to introduce a refutation calculus, which we call $\mathbf{L}[\mathrm{QCL}]^-$. In the literature, such a rejection method for first-order logic with equality was first introduced by Tiomkin (1988) and proved complete w.r.t. finite model theory. The refutation calculus $\mathbf{L}[\mathrm{QCL}]^-$ used in this work is based on a simpler rejection method for propositional logic defined by Bonatti and Olivetti (2002). Using $\mathbf{L}[\mathrm{QCL}]^-$, we prove that $(A)_k$ is satisfiable by deriving the antisequent $(A)_k \nvdash \bot$.

Definition 5.36 (labeled QCL-antisequents)**.** *$\Gamma \nvdash \Delta$ is a labeled* QCL-*antisequent if and only if $\Gamma \vdash \Delta$ is a labeled* QCL-*sequent. $\Gamma \nvdash \Delta$ is valid if and only if $\Gamma \vdash \Delta$ is not valid, i.e., if at least one model that satisfies all labeled* QCL-*formulas in Γ to the degree specified by the label satisfies no labeled* QCL- *formula in Δ to the degree specified by the label.*

The rules for $\mathbf{L}[\mathrm{QCL}]^-$ can be derived from the rules of $\mathbf{L}[\mathrm{QCL}]$. For example, consider the antisequent $\Gamma, (A \vee B)_k \nvdash \Delta$. To show that this antisequent is valid, we must show that the corresponding sequent $\Gamma, (A \vee B)_k \vdash \Delta$ is not valid. This in turn means that we must show that at least one of the two premises $\Gamma, (A)_k \vdash (B)_{<k}, \Delta$ and $\Gamma, (B)_k \vdash (A)_{<k}, \Delta$ introduced by the $\vee l$-rule are not valid. In other words, we must show that either the antisequent $\Gamma, (A)_k \nvdash (B)_{<k}, \Delta$ or the antisequent $\Gamma, (B)_k \nvdash (A)_{<k}, \Delta$ is valid. We therefore introduce two rules:

$$\frac{\Gamma, (A)_k \nvdash (B)_{<k}, \Delta}{\Gamma, (A \vee B)_k \nvdash \Delta} \nvdash \vee l_1 \qquad \frac{\Gamma, (B)_k \nvdash (A)_{<k}, \Delta}{\Gamma, (A \vee B)_k \nvdash \Delta} \nvdash \vee l_2$$

Which one of these two rules should be applied must be guessed, i.e., we trade the branching of $\mathbf{L}[\mathrm{QCL}]$ for non-determinism.

The rules for other connectives are derived similarly. Binary rules are translated into two rules; one inference rule per premise. $(\vee r)$ in $\mathbf{L}[\mathrm{QCL}]$ has an unbounded number of premises, but due to the rules' structure it can be translated into three inference rules. Similarly for $(\wedge r)$, but we need to introduce two extra rules for the case that either A or B is not satisfied.

Regarding the degree overflow rules we introduce a right-hand side rule, but no left-hand side rule. The reason for this is that the antisequent $\Gamma, (A)_{opt_{\mathrm{QCL}}(A)+k} \nvdash \Delta$ is always invalid, i.e., a left-hand side degree overflow rule could never be used in the derivation of a valid antisequent.

The axioms of $\mathbf{L}[\mathrm{QCL}]^-$ are antisequents $\Gamma \nvdash \Delta$ in which Γ and Δ consist only of atoms, no atom p occurs both in Γ and Δ, and $\bot \notin \Gamma$. The reason for this is that $\Gamma \nvdash \Delta$ is valid if and only if the corresponding sequent $\Gamma \vdash \Delta$ is not valid, which is the case only if Γ and Δ are disjoint sets of atoms and $\bot \notin \Gamma$ (cf. Definition 5.28).

Definition 5.37 ($\mathbf{L}[\mathrm{QCL}]^-$)**.** *The axioms of $\mathbf{L}[\mathrm{QCL}]^-$ are labeled* QCL-*antisequents of the form $\Gamma \nvdash \Delta$, where Γ and Δ are disjoint sets of atoms and $\bot \notin \Gamma$. The inference* *rules of* $\mathbf{L}[\mathrm{QCL}]^-$ *are given below. Whenever a labeled* QCL*-formula* $(F)_k$ *appears in the conclusion of an inference rule (except for the* $\nvdash$ *dor-rule) it holds that* $k \leq opt_{\mathcal{L}}(F)$*.*

The rules for the classical connectives are:

$$\frac{\Gamma, (cp(A))_1 \nvdash \Delta}{\Gamma \nvdash (\neg A)_1, \Delta} \nvdash \neg r \qquad \frac{\Gamma \nvdash (cp(A))_1, \Delta}{\Gamma, (\neg A)_1 \nvdash \Delta} \nvdash \neg l$$

$$\frac{\Gamma, (A)_k \nvdash (B)_{<k}, \Delta}{\Gamma, (A \vee B)_k \nvdash \Delta} \nvdash \vee l_1 \qquad \frac{\Gamma, (B)_k \nvdash (A)_{<k}, \Delta}{\Gamma, (A \vee B)_k \nvdash \Delta} \nvdash \vee l_2$$

$$\frac{\Gamma, (A)_i \nvdash \Delta}{\Gamma \nvdash (A \vee B)_k, \Delta} \nvdash \vee r_1 \qquad \frac{\Gamma, (B)_i \nvdash \Delta}{\Gamma \nvdash (A \vee B)_k, \Delta} \nvdash \vee r_2$$

where $i < k$*;*

$$\frac{\Gamma \nvdash (A)_k, (B)_k, \Delta}{\Gamma \nvdash (A \vee B)_k, \Delta} \nvdash \vee r_3$$

$$\frac{\Gamma, (A)_k \nvdash (B)_{>k}, \Delta}{\Gamma, (A \wedge B)_k \nvdash \Delta} \nvdash \wedge l_1 \qquad \frac{\Gamma, (B)_k \nvdash (A)_{>k}, \Delta}{\Gamma, (A \wedge B)_k \nvdash \Delta} \nvdash \wedge l_2$$

$$\frac{\Gamma, (A)_i \nvdash \Delta}{\Gamma \nvdash (A \wedge B)_k, \Delta} \nvdash \wedge r_1 \qquad \frac{\Gamma, (\neg A)_1 \nvdash \Delta}{\Gamma \nvdash (A \wedge B)_k, \Delta} \nvdash \wedge r_2$$

where $k < i \leq opt_{\mathrm{QCL}}(A)$*;*

$$\frac{\Gamma, (B)_i \nvdash \Delta}{\Gamma \nvdash (A \wedge B)_k, \Delta} \nvdash \wedge r_3 \qquad \frac{\Gamma, (\neg B)_1 \nvdash \Delta}{\Gamma \nvdash (A \wedge B)_k, \Delta} \nvdash \wedge r_4$$

where $k < i \leq opt_{\mathrm{QCL}}(B)$*;*

$$\frac{\Gamma \nvdash (A)_k, (B)_k, \Delta}{\Gamma \nvdash (A \wedge B)_k, \Delta} \nvdash \wedge r_5$$

The rules for ordered disjunction are:

$$\frac{\Gamma, (A)_k \nvdash \Delta}{\Gamma, (A \vec{\times} B)_k \nvdash \Delta} \nvdash \vec{\times} l_1 \qquad \frac{\Gamma, (B)_l, (\neg A)_1 \nvdash \Delta}{\Gamma, (A \vec{\times} B)_{opt_{QCL}(A)+l} \nvdash \Delta} \nvdash \vec{\times} l_2$$

$$\frac{\Gamma \nvdash (A)_k, \Delta}{\Gamma \nvdash (A \vec{\times} B)_k, \Delta} \nvdash \vec{\times} r_1 \qquad \frac{\Gamma \nvdash (\neg A)_1, \Delta}{\Gamma \nvdash (A \vec{\times} B)_{opt_{\mathrm{QCL}}(A)+l}, \Delta} \nvdash \vec{\times} r_2$$

$$\frac{\Gamma \nvdash (B)_l, \Delta}{\Gamma \nvdash (A \vec{\times} B)_{opt_{\mathrm{QCL}}(A)+l}, \Delta} \nvdash \vec{\times} r_3$$

where $k \leq opt_{\mathcal{L}}(A)$ *and* $l \leq opt_{\mathcal{L}}(B)$.

The degree overflow rule is:

$$\frac{\Gamma \nvDash \Delta}{\Gamma \nvDash (A)_{opt_{\text{QCL}}(A)+k}, \Delta} \nvDash dor$$

where $k \in \mathbb{N}$.

As already mentioned, instead of branching over several premises it must be guessed non-deterministicly which rule to apply next in $\mathbf{L}[\text{QCL}]^-$, e.g., whether to apply $\wedge l_1$ or $\wedge l_2$. As a result, the proofs found in $\mathbf{L}[\text{QCL}]^-$ are polynomial in size. Moreover, in the proof of an antisequent $\Gamma \nvDash \Delta$ the axiom we encounter directly witnesses a counter example for the corresponding sequent $\Gamma \vdash \Delta$. These differences between $\mathbf{L}[\text{QCL}]^-$ and $\mathbf{L}[\text{QCL}]$ reflect the duality between validity and satisfiability in classical logic.

Example 5.38. *The following derivation is related to Example 5.29 and shows that* $\neg(a \wedge b), ((a \overrightarrow{\times} b) \wedge (b \overrightarrow{\times} c))_2$ *is satisfiable.*

$$\cfrac{\cfrac{\cfrac{\cfrac{\cfrac{\cfrac{\cfrac{\cfrac{a, a, c \nvDash b, b, \bot}{a, c, \neg b \nvDash b, \bot}\ \nvDash \neg l}{a, c \nvDash a \wedge b, b, \bot}\ \nvDash \wedge r_4}{a, c, \neg b \nvDash a \wedge b, \bot}\ \nvDash \neg l}{(a \vee b), c, \neg b \nvDash a \wedge b, \bot}\ \nvDash \vee l_1}{(a \vee b), (b \overrightarrow{\times} c)_2 \nvDash a \wedge b, \bot}\ \nvDash \overrightarrow{\times} l_2}{(b \overrightarrow{\times} c)_2 \nvDash \neg(a \overrightarrow{\times} b), a \wedge b, \bot}\ \nvDash \neg r}{((a \overrightarrow{\times} b) \wedge (b \overrightarrow{\times} c))_2 \nvDash a \wedge b, \bot}\ \nvDash \wedge l_2}{\neg(a \wedge b), ((a \overrightarrow{\times} b) \wedge (b \overrightarrow{\times} c))_2 \nvDash \bot}\ \nvDash \neg l$$

The interpretation $\{a, c\}$ *witnesses the axiom* $a, a, c \nvDash b, b, \bot$ *and also the final antisequent* $\neg(a \wedge b), ((a \overrightarrow{\times} b) \wedge (b \overrightarrow{\times} c))_2 \nvDash \bot$.

Theorem 5.39. $\mathbf{L}[\text{QCL}]^-$ *is sound and complete.*

Proof (Soundness). It is easy to see that, for each rule, the same model witnessing the validity of the premise also witnesses the validity of the conclusion. We exemplify this on hand of the $\nvDash \vee l_1$-rule: assume $\Gamma, (A)_k \nvDash (B)_{<k}, \Delta$ is valid. Then there exists a model M of $\Gamma, (A)_k$. which does not satisfy B to a degree lower than k, and does not satisfy any formula in Δ. Thus, M satisfies $\Gamma, (A \vee B)_k$ and hence $\Gamma, (A \vee B)_k \nvDash \Delta$ is valid. □

Proof (Completeness). Analogously to $\mathbf{L}[\text{QCL}]$, in $\mathbf{L}[\text{QCL}]^-$ we can decompose the formulas of any antisequent into atomic formulas by applying the rules of $\mathbf{L}[\text{QCL}]^-$. Thus, it again suffices to show that the rules of $\mathbf{L}[\text{QCL}]^-$ preserve validity when read "upwards".

Assume $\Gamma \nvdash \Delta$ is valid, i.e. $\Gamma \vdash \Delta$ is not valid. There must be a rule in $\mathbf{L}[\mathrm{QCL}]$ for which $\Gamma \vdash \Delta$ is the conclusion. This rule cannot be the *dol*-rule, since both the premise and conclusion of this rule are always valid. By the soundness of $\mathbf{L}[\mathrm{QCL}]$, the fact that $\Gamma \vdash \Delta$ is not valid implies that at least one of the premises $\Gamma^* \vdash \Delta^*$ of the rule is not valid. However, then $\Gamma^* \nvdash \Delta^*$ is valid. Now, by the construction of $\mathbf{L}[\mathrm{QCL}]^-$, there is a rule that allows us to derive $\Gamma \nvdash \Delta$ from $\Gamma^* \nvdash \Delta^*$. □

Observe that we have not introduced a cut rule for $\mathbf{L}[\mathrm{QCL}]^-$. Indeed, a counterpart of the cut rule would not be sound. One possibility is to introduce a contrapositive of cut as described in (Bonatti and Olivetti 2002).

$$\frac{\Gamma \nvdash \Delta \qquad \Gamma, (A)_k \vdash \Delta}{\Gamma \nvdash (A)_k, \Delta}\ cut2$$

Proposition 5.40. *The cut2-rule is sound.*

Proof. Assume $\Gamma \nvdash \Delta$ and $\Gamma, (A)_k \vdash \Delta$ are valid. Then there is a model M of Γ that does not satisfy any formula in Δ. This further means that M does not satisfy $(A)_k$, otherwise $\Gamma, (A)_k \vdash \Delta$ would not be valid. Thus, $\Gamma \nvdash (A)_k, \Delta$ is valid. □

It is easy to see that, as in $\mathbf{L}[\mathrm{QCL}]$, contraction is also sound in $\mathbf{L}[\mathrm{QCL}]^-$. In contrast to $\mathbf{L}[\mathrm{QCL}]$, however, weakening is not sound in $\mathbf{L}[\mathrm{QCL}]^-$. With left weakening we could, e.g., derive $a, b \nvdash b$ (which is not valid) from $a \nvdash b$ (which is valid). Likewise, with right weakening we could derive $a \nvdash a, b$ from $a \nvdash b$.

We are now ready to combine $\mathbf{L}[\mathrm{QCL}]$ and $\mathbf{L}[\mathrm{QCL}]^-$ by defining an inference rule that allows us to go from labeled QCL-sequents to non-monotonic inferences. We first consider preferred model entailment under minmax semantics (recall Definitions 5.1, 5.5).

Definition 5.41 ($\mathbf{L}[\mathrm{QCL}]^{mm}_{\mid\sim}$)**.** *The axioms of* $\mathbf{L}[\mathrm{QCL}]^{mm}_{\mid\sim}$ *are either labeled* QCL*-sequents of the form* $(p)_1 \vdash (p)_1$ *with* p *being an atom, or labeled* QCL*-antisequents of the form* $\Gamma \nvdash \Delta$*, where* Γ *and* Δ *are disjoint sets of atoms and* $\bot \notin \Gamma$*. The inference rules of* $\mathbf{L}[\mathrm{QCL}]^{mm}_{\mid\sim}$ *are the inference rules of* $\mathbf{L}[\mathrm{QCL}]$ *and* $\mathbf{L}[\mathrm{QCL}]^-$*, extended by*

$$\frac{\langle \Gamma, (A_1 \wedge \dots \wedge A_n)_i \vdash \bot \rangle_{i<k} \qquad \Gamma, (A_1 \wedge \dots \wedge A_n)_k \nvdash \bot \qquad \Gamma, (A_1 \wedge \dots \wedge A_n)_k \vdash \Delta}{\Gamma, A_1, \dots, A_n \mid\sim^{mm}_{\mathrm{QCL}} \Delta}\ \mid\sim_{mm}$$

and

$$\frac{\Gamma, cp(A_1), \dots, cp(A_n) \vdash \bot}{\Gamma, A_1, \dots, A_n \mid\sim^{mm}_{\mathrm{QCL}} \Delta}\ \mid\sim_{unsat}$$

where Γ *consists only of classical formulas and every* A_j *with* $1 \leq j \leq n$ *is a* QCL*-formula.*

Observe that the premises of the new rules $\mid\!\sim_{mm}$ and $\mid\!\sim_{unsat}$ are QCL-sequents and QCL-antisequents, while the conclusion is a $\mid\!\sim_{mm}$ sequent. Consequently, any proof in $\mathbf{L}[\mathrm{QCL}]^{mm}_{\mid\sim}$ can contain only one application of the new rules, in the very last step of the proof. The $\mid\!\sim_{mm}$-rule makes use of the fact that, under minmax semantics, a QCL-theory $T = \{A_1, \ldots, A_n\}$ is semantically equivalent to $A_1 \wedge \cdots \wedge A_n$. The rule can be explained as follows: first, an optimal degree k is guessed. The premises $\langle \Gamma, (A_1 \wedge \cdots \wedge A_n)_i \vdash \bot \rangle_{i<k}$ along with $\Gamma, (A_1 \wedge \cdots \wedge A_n)_k \nvdash \bot$ ensure that models satisfying $A_1 \wedge \cdots \wedge A_n$ to a degree of k are preferred, while the premise $\Gamma, (A_1 \wedge \cdots \wedge A_n)_k \vdash \Delta$ ensures that Δ is entailed by those preferred models. The rule $\mid\!\sim_{unsat}$ is needed in case a theory is classically unsatisfiable.

Example 5.42. *The valid entailment* $\neg(a \wedge b), (a \vec{\times} b), (b \vec{\times} c) \mid\!\sim^{mm}_{\mathrm{QCL}} a \wedge c, b$ *is provable in* $\mathbf{L}[\mathrm{QCL}]^{mm}_{\mid\sim}$ *by choosing* $k = 2$*. Let* $\Gamma = \neg(a \wedge b)$ *and* $\Delta = a \wedge c, b$*.*

$$\dfrac{(\varphi_1) \qquad \overset{(\varphi_2)}{\Gamma, ((a \vec{\times} b) \wedge (b \vec{\times} c))_2 \nvdash \bot} \qquad \overset{(\varphi_3)}{\Gamma, ((a \vec{\times} b) \wedge (b \vec{\times} c))_2 \vdash \Delta}}{\Gamma, (a \vec{\times} b), (b \vec{\times} c) \mid\!\sim^{mm}_{\mathrm{QCL}} \Delta}\ \mid\!\sim_{mm}$$

where (φ_1) *is the derivation*

$$\dfrac{\dfrac{\dfrac{\dfrac{\overset{\vdots}{\Gamma, b \vee c, a, \vdash \neg b, \bot} \qquad \overset{\vdots}{\Gamma, b \vee c, a \vdash c, \bot}}{\Gamma, b \vee c, a \vdash (b \vec{\times} c)_2, \bot}\ \vec{\times} r_2}{\Gamma, b \vee c, (a \vec{\times} b)_1 \vdash (b \vec{\times} c)_2, \bot}\ \vec{\times} l_1}{\Gamma, (a \vec{\times} b)_1 \vdash (b \vec{\times} c)_2, \neg(b \vec{\times} c), \bot}\ \neg r \qquad \dfrac{\dfrac{\overset{\vdots}{\Gamma, a \vee b, b \vdash (a \vec{\times} b)_2, \bot}}{\Gamma, a \vee b, (b \vec{\times} c)_1 \vdash (a \vec{\times} b)_2, \bot}\ \vec{\times} l_1}{\Gamma, (b \vec{\times} c)_1 \vdash (a \vec{\times} b)_2, \neg(a \vec{\times} b), \bot}\ \neg r}{\Gamma, ((a \vec{\times} b) \wedge (b \vec{\times} c))_1 \vdash \bot}\ \wedge l$$

Note that (φ_2) *is the* $\mathbf{L}[\mathrm{QCL}]^-$*-proof from Example 5.38 and* (φ_3) *is the* $\mathbf{L}[\mathrm{QCL}]$*-proof from Example 5.29.*

Theorem 5.43. $\mathbf{L}[\mathrm{QCL}]^{mm}_{\mid\sim}$ *is sound and complete.*

Proof (Soundness). Consider first the $\mid\!\sim_{mm}$-rule and assume that all premises are derivable. By the soundness of $\mathbf{L}[\mathrm{QCL}]$ and $\mathbf{L}[\mathrm{QCL}]^-$ they are also valid. From the first set of premises $\langle \Gamma, (A_1 \wedge \cdots \wedge A_n)_i \vdash \bot \rangle_{i<k}$ we can conclude that if there is some model M of Γ that satisfies $A_1 \wedge \cdots \wedge A_n$ to a degree of k, then $M \in Prf^{mm}_{\mathrm{QCL}}(\Gamma \cup \{A_1, \ldots, A_k\})$. The premise $\Gamma, (A_1 \wedge \cdots \wedge A_n)_k \nvdash \bot$ ensures that there is such a model M. By the last premise $\Gamma, (A_1 \wedge \cdots \wedge A_n)_k \vdash \Delta$, we can conclude that all models of $\Gamma \cup \{A_1, \ldots, A_k\}$ that are equally as preferred as M, i.e., all $M' \in Prf^{mm}_{\mathrm{QCL}}(\Gamma \cup \{A_1, \ldots, A_k\})$, satisfy at least one formula in Δ. Therefore, $\Gamma, A_1, \ldots, A_k \mid\!\sim^{mm}_{\mathrm{QCL}} \Delta$ is valid.

Now consider the $\mid\!\sim_{unsat}$-rule and assume that $\Gamma, cp(A_1), \dots, cp(A_k) \vdash \bot$ is derivable and therefore valid. Then $\Gamma \cup cp(A_1), \dots, cp(A_k)$ has no models. Since in general we have $deg_{\text{QCL}}(\mathcal{I}, F) < \infty \iff \mathcal{I} \models cp(F)$, also $\Gamma \cup \{A_1, \dots, A_k\}$ has no models and thus no preferred models. Then $\Gamma, A_1, \dots, A_k \mid\!\sim^{mm}_{\text{QCL}} \Delta$ is valid. □

Proof (Completeness). Assume that $\Gamma, A_1, \dots, A_k \mid\!\sim^{mm}_{\text{QCL}} \Delta$ is valid. If $\Gamma \cup \{A_1, \dots, A_k\}$ is unsatisfiable then $\Gamma, cp(A_1), \dots, cp(A_k) \vdash \bot$ is valid, i.e., we can apply the $\mid\!\sim_{unsat}$-rule.

Now consider the case that $\Gamma \cup \{A_1, \dots, A_k\}$ is satisfiable and assume that some preferred model M of $\Gamma \cup \{A_1, \dots, A_k\}$ satisfies $A_1 \wedge \dots \wedge A_n$ to a degree of k. Then, we claim that all premises of the rule are valid and, by the completeness of $\mathbf{L}$[QCL] and $\mathbf{L}[\text{QCL}]^-$, also derivable.

Assume by contradiction that one of the premises is not valid. First, consider the case that $\Gamma, (A_1 \wedge \dots \wedge A_n)_i \vdash \bot$ is not valid for some $i < k$. Then there is a model M' of Γ that satisfies $A_1 \wedge \dots \wedge A_n$ to a degree of $i < k$. However, this contradicts the assumption that M is a preferred model of $\Gamma \cup \{A_1, \dots, A_k\}$.

Next, assume that $\Gamma, (A_1 \wedge \dots \wedge A_n)_k \nvdash \bot$ is not valid. However, M satisfies $(A_1 \wedge \dots \wedge A_n)_k$ and does not satisfy $\bot$. Contradiction.

Finally, we assume that $\Gamma, (A_1 \wedge \dots \wedge A_n)_k \vdash \Delta$ is not valid. Then, there is a model M' of Γ that satisfies $A_1 \wedge \dots \wedge A_n$ to a degree of k but does not satisfy any formula in Δ. But M' is a preferred model of $\Gamma \cup \{A_1, \dots, A_k\}$, which contradicts $\Gamma, A_1, \dots, A_k \mid\!\sim^{mm}_{\text{QCL}} \Delta$ being valid. □

To obtain a calculus for preferred model entailment under lexicographic semantics, we adapt the $\mid\!\sim_{mm}$-rule of $\mathbf{L}[\text{QCL}]^{mm}_{\mid\sim}$.

Definition 5.44 ($\mathbf{L}[\text{QCL}]^{lex}_{\mid\sim}$). *Let $\leq_l$ be the order on vectors in $\mathbb{N}^k$ defined by*

- *$\vec{v} <_l \vec{w}$ if there is some $n \in \mathbb{N}$ such that $\vec{v}$ has more entries of value n and for all $1 \leq m < n$ both vectors have the same number of entries of value m.*
- *$\vec{v} =_l \vec{w}$ if, for all $n \in \mathbb{N}$, $\vec{v}$ and $\vec{w}$ have the same number of entries of value n.*

The axioms and inference rules of $\mathbf{L}[\text{QCL}]^{lex}_{\mid\sim}$ are the same as those of $\mathbf{L}[\text{QCL}]^{mm}_{\mid\sim}$, except that $\mid\!\sim_{mm}$ is replaced by

$$\frac{\langle \Gamma, (A_1)_{w_1}, \dots, (A_k)_{w_k} \vdash \bot \rangle_{\vec{w} <_l \vec{v}} \quad \Gamma, (A_1)_{v_1}, \dots, (A_k)_{v_k} \nvdash \bot \quad \langle \Gamma, (A_1)_{w_1}, \dots, (A_k)_{w_k} \vdash \Delta \rangle_{\vec{w} =_l \vec{v}}}{\Gamma, A_1, \dots, A_k \mid\!\sim^{lex}_{\text{QCL}} \Delta} \mid\!\sim_{lex}$$

where $\vec{v}, \vec{w} \in \mathbb{N}^k$, Γ consists only of classical formulas and every A_j with $1 \leq j \leq k$ is a QCL*-formula.*

Instead of guessing the degree of preferred models as in $\mathbf{L}[\mathrm{QCL}]^{mm}_{\mid\sim}$, we now guess a "degree-profile" (in form of the vector $\vec{v}$) of at least one preferred model of $\Gamma \cup \{A_1, \ldots, A_k\}$. The rule $\mid\sim_{lex}$ can be explained as follows: The premises shown in the left branch confirm that our guess is indeed optimal, i.e., that $\Gamma, (A_1)_{w_1}, \ldots, (A_k)_{w_k}$ cannot be satisfied if $\vec{w}$ is better than $\vec{v}$ with respect to the *lex*-semantics. The center premise ensures that our degree-profile is satisfiable. The right premise ensures that *all* preferred models, meaning all models with a degree profile $\vec{w}$ as good as $\vec{v}$ with respect to the *lex*-semantics, satisfy at least one formula in Δ. Note that, as for $\mathbf{L}[\mathrm{QCL}]^{mm}_{\mid\sim}$, any proof in $\mathbf{L}[\mathrm{QCL}]^{lex}_{\mid\sim}$ can contain only one application of the new rules, in the very last step of the proof. Let us provide a small example before showing soundness and completeness of $\mathbf{L}[\mathrm{QCL}]^{lex}_{\mid\sim}$.

Example 5.45. *Consider the valid entailment* $\neg(a \wedge b), (a \vec{\times} b), (b \vec{\times} c) \mid\sim^{lex}_{\mathrm{QCL}} a \wedge c, b$ *similar to Example 5.42. Let* $\Gamma = \neg(a \wedge b)$ *and* $\Delta = a \wedge c, b$. *Therefore, we can also write the entailment as* $\Gamma, (a \vec{\times} b), (b \vec{\times} c) \mid\sim^{lex}_{\mathrm{QCL}} \Delta$. *Note that it is not possible to satisfy all labeled* QCL*-formulas on the left to a degree of* 1. *Rather, it is optimal to either satisfy* $\Gamma, (a \vec{\times} b)_1, (b \vec{\times} c)_2$ *or, alternatively,* $\Gamma, (a \vec{\times} b)_2, (b \vec{\times} c)_1$. *We choose* $\vec{v} = (1, 2)$. *Observe that* $\vec{w} = (1, 1)$ *is the only vector* $\vec{w}$ *such that* $\vec{w} <_l \vec{v}$. *Moreover,* $(1, 2) =_l \vec{v}$ *and* $(2, 1) =_l \vec{v}$. *Thus, we get*

$$\dfrac{\begin{array}{ccc} \vdots & \vdots & \vdots \\ \Gamma, (a \vec{\times} b)_1, (b \vec{\times} c)_1 \vdash \bot & \Gamma, (a \vec{\times} b)_1, (b \vec{\times} c)_2 \nvdash \bot & \Gamma, (a \vec{\times} b)_1, (b \vec{\times} c)_2 \vdash \Delta \quad * \end{array}}{\Gamma, (a \vec{\times} b), (b \vec{\times} c) \mid\sim^{lex}_{\mathrm{QCL}} \Delta} \mid\sim_{lex}$$

where $*$ *is*

$$\begin{array}{c} \vdots \\ \Gamma, (a \vec{\times} b)_2, (b \vec{\times} c)_1 \vdash \Delta \end{array}$$

It can be verified that indeed all branches are provable, but we do not show this explicitly here.

Theorem 5.46. $\mathbf{L}[\mathrm{QCL}]^{lex}_{\mid\sim}$ *is sound and complete.*

Proof (Soundness). Consider the $\mid\sim_{lex}$-rule and assume that all premises are derivable. By the soundness of $\mathbf{L}[\mathrm{QCL}]$ and $\mathbf{L}[\mathrm{QCL}]^{-}$ they are also valid. From the first set of premises $\langle \Gamma, (A_1)_{w_1}, \ldots, (A_k)_{w_k} \vdash \bot \rangle_{\vec{w} <_l \vec{v}}$ we can conclude that if there is some model M of Γ that satisfies A_i to a degree of v_i for all $1 \leq i \leq k$, then $M \in \mathit{Prf}^{lex}_{\mathrm{QCL}}(\Gamma \cup \{A_1, \ldots, A_k\})$. The premise $\Gamma, (A_1)_{v_1}, \ldots, (A_k)_{v_k} \nvdash \bot$ ensures that there is such a model M. By the last set of premises $\langle \Gamma, (A_1)_{w_1}, \ldots, (A_k)_{w_k} \vdash \Delta \rangle_{\vec{w} =_l \vec{v}}$ we can conclude that all models of $\Gamma \cup \{A_1, \ldots, A_k\}$ that are equally as preferred as M, i.e., all $M' \in \mathit{Prf}^{lex}_{\mathrm{QCL}}(\Gamma \cup \{A_1, \ldots, A_k\})$, satisfy at least one formula in Δ. Therefore, $\Gamma, A_1, \ldots, A_k \mid\sim^{lex}_{\mathrm{QCL}} \Delta$ is valid. □

Proof (Completeness). Assume that $\Gamma, A_1, \dots, A_k \mid\!\sim^{lex}_{\mathrm{QCL}} \Delta$ is valid. If $\Gamma \cup \{A_1, \dots, A_k\}$ is unsatisfiable then $\Gamma, cp(A_1), \dots, cp(A_k) \vdash \bot$ is valid, i.e., we can apply the $\mid\!\sim_{unsat}$-rule. Now consider the case that $\Gamma \cup \{A_1, \dots, A_k\}$ is satisfiable and assume that some preferred model M of $\Gamma \cup \{A_1, \dots, A_k\}$ satisfies A_i to a degree of v_i for all $1 \leq i \leq k$. Then, we claim that all premises of the rule are valid and, by the completeness of $\mathbf{L}$[QCL] and $\mathbf{L}$[QCL]$^-$, also derivable.

Assume by contradiction that one of the premises is not valid. First, consider the case that $\Gamma, (A_1)_{w_1}, \dots, (A_k)_{w_k} \vdash \bot$ is not valid for some $\vec{w} <_l \vec{v}$. Then there is a model M' of Γ that satisfies A_i to a degree of w_i for all $1 \leq i \leq k$. However, this contradicts the assumption that M is a preferred model of $\Gamma \cup \{A_1, \dots, A_k\}$.

Next, assume that $\Gamma, (A_1)_{v_1}, \dots, (A_k)_{v_k} \nvdash \bot$ is not valid. However, M satisfies all formulas in $\Gamma, (A_1)_{v_1}, \dots, (A_k)_{v_k}$ and does not satisfy $\bot$. Contradiction.

Finally, we assume that $\Gamma, (A_1)_{w_1}, \dots, (A_k)_{w_k} \vdash \Delta$ is not valid for some $\vec{w} =_l \vec{v}$. Then, there is a model M' of Γ that satisfies A_i to a degree of w_i for all $1 \leq i \leq k$ but does not satisfy any formula in Δ. But M' is a preferred model of $\Gamma \cup \{A_1, \dots, A_k\}$, which contradicts $\Gamma, A_1, \dots, A_k \mid\!\sim^{lex}_{\mathrm{QCL}} \Delta$ being valid. □

Finally, a calculus for the inclusion-based approach of preferred model entailment can be obtained by simply adapting the way in which vectors over $\mathbb{N}^k$ are compared (cf. Definition 5.44).

Definition 5.47 ($\mathbf{L}[\mathrm{QCL}]^{inc}_{\mid\sim}$)**.** *The calculus* $\mathbf{L}[\mathrm{QCL}]^{inc}_{\mid\sim}$ *is defined analogously to* $\mathbf{L}[\mathrm{QCL}]^{lex}_{\mid\sim}$ *(cf. Definition 5.44) except that the order* $\leq_l$ *is replaced by the order* $\leq_i$*:*

- $\vec{v} <_i \vec{w}$ *if there is some* $n \in \mathbb{N}$ *such that every entry in* $\vec{w}$ *with value* n *also has value* n *in* $\vec{v}$*, there is an entry in* $\vec{v}$ *with value* n *that has a value higher than* n *in* $\vec{w}$*, and for all* $1 \leq m < n$ *both vectors have exactly the same entries with value* m*.*
- $\vec{v} =_i \vec{w}$ *if* $\vec{v} \nless_i \vec{w}$ *and* $\vec{w} \nless_i \vec{v}$*.*

Soundness and completeness of $\mathbf{L}[\mathrm{QCL}]^{inc}_{\mid\sim}$ are analogous to that of $\mathbf{L}[\mathrm{QCL}]^{lex}_{\mid\sim}$ (cf. Theorem 5.46). Note that we did not define a calculus for the log-lexicographic semantics (cf. Definition 5.3). Such a calculus could be defined by adapting the ordering $<_l$ used in the $\mid\!\sim_{lex}$-rule of $\mathbf{L}[\mathrm{QCL}]^{lex}_{\mid\sim}$ (cf. Definition 5.44), but we do not consider this here.

5.4.3 Beyond QCL

We will now demonstrate that the calculi for QCL introduced in the previous sections can easily be adapted for other choice logics. Indeed, to introduce a labeled calculus for some choice logic $\mathcal{L}$ other than QCL it suffices to replace the $\vec{\times}$-rules in $\mathbf{L}$[QCL] by appropriate rules for the choice connectives of $\mathcal{L}$. The rules for the classical connectives in $\mathbf{L}$[QCL] can be retained. Moreover, note that the inference rules for preferred model entailment (i.e., the rules $\mid\!\sim_{mm}$, $\mid\!\sim_{lex}$, $\mid\!\sim_{inc}$, $\mid\!\sim_{unsat}$ from Definitions 5.41, 5.44, 5.47) do not depend on any specific choice logic. Thus, once labeled calculi are developed for $\mathcal{L}$, the calculi for preferred model entailment follow immediately.

Calculi for CCL

First, we introduce $\mathbf{L}$[CCL] by defining rules for the choice connective $\vec{\odot}$ of CCL. Recall that $A \vec{\odot} B$ expresses that, if possible, both A and B should be satisfied, but if this is not possible then satisfying only A is also acceptable (cf. Section 2.4).

Definition 5.48. $\mathbf{L}$[CCL] *is* $\mathbf{L}$[QCL]*, except that the* $\vec{\times}$*-rules are replaced by the following* $\vec{\odot}$*-rules:*

$$\frac{\Gamma, (A)_1, (B)_k \vdash \Delta}{\Gamma, (A \vec{\odot} B)_k \vdash \Delta}\ \vec{\odot} l_1 \qquad \frac{\Gamma, (A)_1, (\neg B)_1 \vdash \Delta}{\Gamma, (A \vec{\odot} B)_{opt_{\mathrm{CCL}}(B)+1} \vdash \Delta}\ \vec{\odot} l_2$$

$$\frac{\Gamma, (A)_l \vdash \Delta}{\Gamma, (A \vec{\odot} B)_{opt_{\mathrm{CCL}}(B)+l} \vdash \Delta}\ \vec{\odot} l_3$$

$$\frac{\Gamma \vdash (A)_1, \Delta \qquad \Gamma \vdash (B)_k, \Delta}{\Gamma \vdash (A \vec{\odot} B)_k, \Delta}\ \vec{\odot} r_1 \qquad \frac{\Gamma \vdash (A)_1, \Delta \qquad \Gamma \vdash (\neg B)_1, \Delta}{\Gamma \vdash (A \vec{\odot} B)_{opt_{\mathrm{CCL}}(B)+1}, \Delta}\ \vec{\odot} r_2$$

$$\frac{\Gamma \vdash (A)_l, \Delta}{\Gamma \vdash (A \vec{\odot} B)_{opt_{\mathrm{CCL}}(B)+l}, \Delta}\ \vec{\odot} r_3$$

where $k \leq opt_{\mathrm{CCL}}(B)$ *and* $1 < l \leq opt_{\mathrm{CCL}}(A)$*.*

The $\vec{\odot} l_1$-rule takes care of the case in which A is optimally satisfied, and B is satisfied to some degree. In $\vec{\odot} l_2$ and $\vec{\odot} l_3$ the label m of $(A \vec{\odot} B)_m$ is higher than the optionality of B. If $m = opt_{\mathrm{CCL}}(B) + 1$ we know that B cannot be satisfied, and hence we need to apply $\vec{\odot} l_2$. If $m = opt_{\mathrm{CCL}}(B) + l$ with $l > 1$ then, by the semantics of CCL (cf. Definition 2.33), it must be that A is satisfied to a degree of l, regardless of whether B is satisfied or not.

Example 5.49. *The following is a small* $\mathbf{L}$[CCL]*-proof of a valid sequent, showcasing the application of the* $\vec{\odot} l_2$*- and* $\vec{\odot} l_3$*-rules.*

$$\cfrac{\cfrac{\vdots}{\Gamma, (a)_1, (\neg b)_1 \vdash a \land \neg b}}{\cfrac{\Gamma, (a \vec{\odot} b)_2 \vdash a \land \neg b}{((a \vec{\odot} b) \vec{\odot} c)_3 \vdash a \land \neg b}\ \vec{\odot} l_3}\ \vec{\odot} l_2$$

Theorem 5.50. $\mathbf{L}$[CCL] *is sound and complete.*

Proof (Soundness). We consider the newly introduced rules.

- For $\vec{\odot}l_1$, $\vec{\odot}l_2$, and $\vec{\odot}l_3$ this follows directly from the definition of CCL.
- $(\vec{\odot}r_1)$. Assume both premises are valid, i.e., every model of Γ is a model of Δ or of $(A)_1$ and $(B)_k$ with $k \leq opt_{\text{CCL}}(B)$. By definition, any model that satisfies $(A)_1$ and $(B)_k$ satisfies $A\vec{\odot}B$ to degree k. Thus, every model of Γ is a model of Δ or of $(A\vec{\odot}B)_k$, which means the conclusion of the rule is valid.
- $(\vec{\odot}r_2)$. Assume both premises are valid, i.e., every model of Γ is a model of Δ or of $(A)_1$ and $(\neg B)_1$. By definition, any model that satisfies $(A)_1$ and does not satisfy B (and hence satisfies $(\neg B)_1$) satisfies $A\vec{\odot}B$ to degree $opt_{\text{CCL}}(B) + 1$.
- $(\vec{\odot}r_3)$. Assume the premise is valid, i.e., every model of Γ is a model of Δ or of $(A)_l$ with $1 < l \leq opt_{\text{CCL}}(A)$. By definition, any model that satisfies $(A)_l$, regardless of what degree this model ascribes to B, satisfies $A\vec{\odot}B$ to degree $opt_{\text{CCL}}(B) + l$. □

Proof (Completeness). We adapt the completeness proof of **L**[QCL] (cf. Theorem 5.32).

- Assume that a sequent of the form $\Gamma, (A\vec{\odot}B)_k \vdash \Delta$ is valid, with $k \leq opt_{\text{CCL}}(B)$. All models that satisfy $(A\vec{\odot}B)_k$ must satisfy A to a degree of 1 and B to a degree of k. Thus, $\Gamma, (A)_1, (B)_k \vdash \Delta$ is valid. Similarly for the cases $\Gamma, (A\vec{\odot}B)_{opt_{\text{CCL}}(B)+1} \vdash \Delta$ and $\Gamma, (A\vec{\odot}B)_{opt_{\text{CCL}}(B)+l} \vdash \Delta$ with $1 < l \leq opt_{\text{CCL}}(A)$.
- Assume that a sequent of the form $\Gamma \vdash (A\vec{\odot}B)_k, \Delta$ is valid, with $k \leq opt_{\text{CCL}}(B)$. We claim that then $\Gamma \vdash (A)_1, \Delta$ and $\Gamma \vdash (B)_k, \Delta$ are valid. Assume, for the sake of a contradiction, that the first sequent is not valid. This means that there is a model M of Γ that is neither a model of $(A)_1$ nor of Δ. However, then M satisfies $A\vec{\odot}B$ to a degree higher than $opt_{\text{CCL}}(B)$. This contradicts the assumption that $\Gamma \vdash (A\vec{\odot}B)_k, \Delta$ is valid. Assume now that the second sequent is not valid, i.e., that there is a model M of Γ that is neither a model of $(B)_k$ nor of Δ. Then M cannot be a model of $(A\vec{\odot}B)_k$, contradicting the assumption. Similarly for the cases $\Gamma \vdash (A\vec{\odot}B)_{opt_{\text{CCL}}(B)+1}, \Delta$ and $\Gamma \vdash (A\vec{\odot}B)_{opt_{\text{CCL}}(B)+l}, \Delta$ with $1 < l \leq opt_{\text{CCL}}(A)$. □

We do not define the refutation calculus **L**[CCL]$^-$ here, but the necessary rules for $\vec{\odot}$ can be inferred from the $\vec{\odot}$-rules of **L**[CCL] in a similar way to how **L**[QCL]$^-$ was derived from **L**[QCL]: if a rule contains only a single premise then it suffices to replace the $\vdash$-symbol with the $\nvdash$-symbol; if a rule contains two premises then we introduce two rules in **L**[CCL]$^-$, one for each premise. Once **L**[CCL] and **L**[CCL]$^-$ are established, calculi for preferred model entailment follow immediately.

Calculi for LCL

Our methods can also be adapted for LCL in which $A \vec{\diamond} B$ expresses that it is best to satisfy A and B, second best to satisfy only A, third best to satisfy only B, and unacceptable to satisfy neither (cf. Section 2.4).

Definition 5.51. **L**[LCL] *is* **L**[QCL], *except that the* $\vec{\times}$*-rules are replaced by the following* $\vec{\diamond}$*-rules:*

$$\frac{\Gamma, (A)_k, (B)_l \vdash \Delta}{\Gamma, (A \vec{\diamond} B)_{(k-1)\cdot opt_{\mathrm{LCL}}(B)+l} \vdash \Delta}\ \vec{\diamond} l_1$$

$$\frac{\Gamma, (A)_k, (\neg B)_1 \vdash \Delta}{\Gamma, (A \vec{\diamond} B)_{opt_{\mathrm{LCL}}(A)\cdot opt_{\mathrm{LCL}}(B)+k} \vdash \Delta}\ \vec{\diamond} l_2$$

$$\frac{\Gamma, (\neg A)_1, (B)_l \vdash \Delta}{\Gamma, (A \vec{\diamond} B)_{opt_{\mathrm{LCL}}(A)\cdot opt_{\mathrm{LCL}}(B)+opt_{\mathrm{LCL}}(A)+l} \vdash \Delta}\ \vec{\diamond} l_3$$

$$\frac{\Gamma \vdash (A)_k, \Delta \qquad \Gamma \vdash (B)_l, \Delta}{\Gamma \vdash (A \vec{\diamond} B)_{(k-1)\cdot opt_{\mathrm{LCL}}(B)+l}, \Delta}\ \vec{\diamond} r_1$$

$$\frac{\Gamma \vdash (A)_k, \Delta \qquad \Gamma \vdash (\neg B)_1, \Delta}{\Gamma \vdash (A \vec{\diamond} B)_{opt_{\mathrm{LCL}}(A)\cdot opt_{\mathrm{LCL}}(B)+k}, \Delta}\ \vec{\diamond} r_2$$

$$\frac{\Gamma \vdash (\neg A)_1, \Delta \qquad \Gamma \vdash (B)_l, \Delta}{\Gamma \vdash (A \vec{\diamond} B)_{opt_{\mathrm{LCL}}(A)\cdot opt_{\mathrm{LCL}}(B)+opt_{\mathrm{LCL}}(A)+l}, \Delta}\ \vec{\diamond} r_3$$

where $k \leq opt_{\mathrm{LCL}}(A)$ *and* $l \leq opt_{\mathrm{LCL}}(B)$*.*

The labels used in the above rules might appear quite involved. However, finding the correct rule to apply given a labeled LCL-formula $(A \vec{\diamond} B)_m$ is actually a straightforward task: the values for $opt_{\mathrm{LCL}}(A)$ and $opt_{\mathrm{LCL}}(B)$ can be computed according to Definition 2.35. If $m \leq opt_{\mathrm{LCL}}(A) \cdot opt_{\mathrm{LCL}}(B)$ then the $\vec{\diamond} l_1$-rule must be applied. If $opt_{\mathrm{LCL}}(A) \cdot opt_{\mathrm{LCL}}(B) < m \leq opt_{\mathrm{LCL}}(A) \cdot opt_{\mathrm{LCL}}(B) + opt_{\mathrm{LCL}}(A)$ then the $\vec{\diamond} l_2$-rule must be applied. If $opt_{\mathrm{LCL}}(A) \cdot opt_{\mathrm{LCL}}(B) + opt_{\mathrm{LCL}}(A) < m \leq opt_{\mathrm{LCL}}(A) \cdot opt_{\mathrm{LCL}}(B) + opt_{\mathrm{LCL}}(A) + opt_{\mathrm{LCL}}(B)$ then the $\vec{\diamond} l_3$-rule must be applied.

Example 5.52. *The following is a small* **L**[LCL]*-proof of a valid sequent. Since we have a label of* 2 *in the end-sequent, and since* $opt_{\mathrm{LCL}}(a \vee b) = opt_{\mathrm{LCL}}(b \vee c) = 1$*, we know that the* $\vec{\diamond} l_2$*-rule must be applied.*

$$\dfrac{\dfrac{\vdots}{a \vee b \vdash b \vee c, a \wedge \neg b}}{\dfrac{a \vee b, \neg(b \vee c) \vdash a \wedge \neg b}{((a \vee b)\vec{\diamond}(b \vee c))_2 \vdash a \wedge \neg b}\ \vec{\diamond}l_2}\ \neg l$$

Theorem 5.53. **L**[LCL] *is sound and complete.*

Proof (Soundness). We consider the newly introduced rules.

- For $\vec{\diamond}l_1$, $\vec{\diamond}l_2$, and $\vec{\diamond}l_3$ this follows directly from the definition of LCL.
- ($\vec{\diamond}r_1$). Assume both premises are valid, i.e., every model of Γ is a model of Δ or of $(A)_k$ and $(B)_l$ with $k \leq opt_{\text{LCL}}(A)$ and $l \leq opt_{\text{LCL}}(B)$. By definition, any model that satisfies $(A)_k$ and $(B)_l$ satisfies $A\vec{\diamond}B$ to degree $(k-1) \cdot opt_{\text{LCL}}(B) + l$. Thus, every model of Γ is a model of Δ or of $(A\vec{\diamond}B)_{(k-1) \cdot opt_{\text{LCL}}(B)+l}$, which means the conclusion of the rule is valid.
- ($\vec{\diamond}r_2$). Assume both premises are valid, i.e., every model of Γ is a model of Δ or of $(A)_k$ and $(\neg B)_1$ with $k \leq opt_{\text{LCL}}(A)$. By definition, any model that satisfies $(A)_k$ and does not satisfy B (and hence satisfies $(\neg B)_1$) satisfies $A\vec{\diamond}B$ to degree $opt_{\text{LCL}}(A) \cdot opt_{\text{LCL}}(B) + k$.
- ($\vec{\diamond}r_3$). Analogous to ($\vec{\diamond}r_2$). □

Proof (Completeness). We adapt the completeness proof of **L**[QCL] (cf. Theorem 5.32).

- Assume that a sequent of the form $\Gamma, (A\vec{\diamond}B)_m \vdash \Delta$ is valid, with $m = (k-1) \cdot opt_{\text{LCL}}(B) + l$ such that $k \leq opt_{\text{LCL}}(A)$ and $l \leq opt_{\text{LCL}}(B)$. Now assume some model satisfies Γ, $(A)_k$, and $(B)_l$. Then M satisfies Γ and $(A\vec{\diamond}B)_m$, and, since $\Gamma, (A\vec{\diamond}B)_m \vdash \Delta$ is valid, M also satisfies Δ. Thus, $\Gamma, (A)_k, (B)_l \vdash \Delta$ is valid.

  The proofs for sequents of the form $\Gamma, (A\vec{\diamond}B)_{opt_{\text{LCL}}(A) \cdot opt_{\text{LCL}}(B)+k} \vdash \Delta$ as well as $\Gamma, (A\vec{\diamond}B)_{opt_{\text{LCL}}(A) \cdot opt_{\text{LCL}}(B)+opt_{\text{LCL}}(A)+l} \vdash \Delta$ are analogous.
- Assume that a sequent of the form $\Gamma \vdash (A\vec{\diamond}B)_m, \Delta$ is valid, with $m = (k-1) \cdot opt_{\text{LCL}}(B) + l$ such that $k \leq opt_{\text{LCL}}(A)$ and $l \leq opt_{\text{LCL}}(B)$. We claim that then $\Gamma \vdash (A)_k, \Delta$ and $\Gamma \vdash (B)_l, \Delta$ are valid. Assume, for the sake of a contradiction, that the first sequent is not valid. This means that there is a model M of Γ that is neither a model of $(A)_k$ nor of Δ. Following Definition 2.35, M must satisfy $A\vec{\diamond}B$ to some degree other than m. This contradicts the assumption that $\Gamma \vdash (A\vec{\diamond}B)_m, \Delta$ is valid. Assume now that the second sequent is not valid, i.e., that there is a model M of Γ that is neither a model of $(B)_l$ nor of Δ. Again, this means that M satisfies $A\vec{\diamond}B$ to some degree other than m, and this would contradict our assumption that

$\Gamma \vdash (A \vec{\diamond} B)_m, \Delta$ is valid. Thus, both $\Gamma \vdash (A)_k, \Delta$ and $\Gamma \vdash (B)_k, \Delta$ are valid and $\Gamma \vdash (A \vec{\diamond} B)_m, \Delta$ is provable.

The proofs for sequents of the form $\Gamma \vdash (A \vec{\diamond} B)_{opt_{\mathrm{LCL}}(A) \cdot opt_{\mathrm{LCL}}(B)+k}, \Delta$ and $\Gamma \vdash (A \vec{\diamond} B)_{opt_{\mathrm{LCL}}(A) \cdot opt_{\mathrm{LCL}}(B)+opt_{\mathrm{LCL}}(A)+l}, \Delta$ are analogous. □

As with CCL, the refutation calculus $\mathbf{L}[\mathrm{LCL}]^-$ can be obtained from $\mathbf{L}[\mathrm{LCL}]$ by modifying the $\vec{\diamond}$-rules accordingly. Calculi for preferred model entailment follow immediately.

Multiple Choice Connectives

Lastly, we want to point out that, according to the choice logic framework (cf. Subsection 2.4.1), choice logics can make use of more than one choice connective. Indeed, a combination of QCL and CCL into the so-called QCCL has been suggested in the master thesis of the author (Bernreiter 2020). QCCL is simply the choice logic with choice connectives $\mathcal{C}_{\mathrm{QCCL}} = \{\vec{\times}, \vec{\odot}\}$, with the optionality and satisfaction degree of $\vec{\times}$ (resp. $\vec{\odot}$) defined in the same way as in QCL (resp. CCL). A calculus for QCCL can be obtained simply by adding both the rules for $\vec{\times}$ and $\vec{\odot}$. We demonstrate this with a small example.

Example 5.54. *The following is a proof of a valid sequent in QCCL. We use lexicographic entailment, but one could also use the minmax or inclusion-based approaches instead. Since the formulas $(a \vec{\odot} b)$ and $(b \vec{\times} c)$ are jointly satisfiable to a degree of 1 we can guess the optimal degree-profile $(a \vec{\odot} b)_1, (b \vec{\times} c)_1$. Thus, we only have two branches in the $\mid\!\sim_{lex}$-rule.*

$$
\dfrac{
\dfrac{\dfrac{a, b, b \nvdash \bot}{a, b, (b \vec{\times} c)_1 \nvdash \bot}\;\nvdash \vec{\times} l_1}{(a \vec{\odot} b)_1, (b \vec{\times} c)_1 \nvdash \bot}\;\nvdash \vec{\odot} l_1
\qquad
\dfrac{\dfrac{\dfrac{\vdots}{a, b, b \vdash a \wedge b}}{a, b, (b \vec{\times} c)_1 \vdash a \wedge b}\;\vec{\times} l_1}{(a \vec{\odot} b)_1, (b \vec{\times} c)_1 \vdash a \wedge b}\;\vec{\odot} l_1
}{(a \vec{\odot} b), (b \vec{\times} c) \mid\!\sim^{lex}_{\mathrm{QCCL}} a \wedge b}\;\mid\!\sim_{lex}
$$

5.5 Conclusion

In this chapter, we studied preferred model entailment in choice logics with respect to semantic, computational, and proof-theoretic properties.

Regarding semantic properties, we investigated the principles of cautious monotonicity, cumulative transitivity, and rational monotonicity laid out by Kraus, Lehmann, and Magidor (1990). We showed that preferred model entailment satisfies cautious monotonicity and cumulative transitivity for all choice logics and all preferred model semantics, assuming we are dealing with finite theories. Moreover, under all considered preferred model semantics except for the inclusion-based semantics, rational monotonicity is also satisfied for all choice logics.

As for computational properties, we showed that checking if an interpretation is a preferred model is coNP-complete for choice logics in which more than two satisfaction degrees are obtainable. The complexity of preferred model entailment depends both on the choice logic and on the preferred model semantics considered (cf. Table 5.1). For QCL and CCL, entailment is Θ_2^{P}-complete under minmax semantics, $\Delta_2^{\mathsf{P}}[O(\log^2 n)]$-complete under log-lexicographic semantics, and Δ_2^{P}-complete under lexicographic semantics. For LCL, entailment is Δ_2^{P}-complete for all of these three semantics. Entailment under the inclusion-based semantics is Π_2^{P}-complete for QCL, CCL, and LCL.

Lastly, we introduced a sound and complete sequent calculus for preferred model entailment in QCL. This non-monotonic calculus is built on two calculi: a monotonic labeled sequent calculus and a corresponding refutation calculus. Our systems are modular and can easily be adapted. Calculi for other choice logics can be obtained by introducing suitable rules for the choice connectives of the new logic, as exemplified with our calculi for CCL and LCL. Moreover, non-monotonic calculi for alternative preferred model semantics can be obtained by adapting the inference rule which transitions from preferred model entailment to the labeled calculi (e.g. the $\mid\!\sim_{mm}$, $\mid\!\sim_{lex}$ or $\mid\!\sim_{inc}$-rule).

An interesting avenue for future work is to examine alternative semantics for languages using ordered disjunction or other choice connectives, and see whether our methods can be adapted to those approaches. We now give a brief overview over relevant work in this direction. In Prioritized QCL (PQCL) and QCL+ (Benferhat and Sedki 2008) ordered disjunction is defined in the same way as in QCL, but the classical connectives are given new semantics. As pointed out in previous work (Bernreiter, Maly, and Woltran 2021), both PQCL and QCL+ can be captured by the choice logic framework as fragments by allowing negations only in front of atoms. Another interesting paper is that of Maly and Woltran (2018), in which the concept of satisfaction degrees is abandoned and the semantics rather 'directly' induces a partial order over models. The most recent reinterpretation of QCL that we are aware of is an approach (Freiman and Bernreiter 2023a,b) using game theoretic semantics, with a special focus on providing an alternative negation for the language of QCL. Note that (Freiman and Bernreiter 2023b) also features a proof calculi for this reinterpretation, although these calculi do not allow to decide preferred model entailment. A logic similar to LCL was proposed by Charalambidis et al. (2021). In contrast to LCL, their lexicographic logic uses lists of truth values to rank interpretations rather than satisfaction degrees. In the world of logic programming, recent works (Charalambidis, Nomikos, and Rondogiannis 2022; Charalambidis, Rondogiannis, and Troumpoukis 2021) have suggested a new semantics for logic programs with ordered disjunction (LPODs) (Brewka, Niemelä, and Syrjänen 2004). While the original semantics of LPODs uses satisfaction degrees as in QCL, the new approach uses a four-valued logic.

Specifically regarding our proof calculi, a possibility for future work is to study their proof complexity and how this complexity might depend on which choice logic or preferred model semantics is considered. Moreover, developing a calculus for LPODs might prove to be interesting since they contain two sources of non-monotonicity (logic programming itself as well as ordered disjunction).

Another possible direction for the further study of choice logics is to investigate the connection to conditional knowledge bases (Kraus, Lehmann, and Magidor 1990; Lehmann and Magidor 1992) in more detail. Our complexity results for choice logics (cf. Section 5.3) together with the known complexity of conditional knowledge bases (Eiter and Lukasiewicz 2000) suggest that polynomial time transformations from choice logic theories to conditional knowledge bases and vice versa are possible. How exactly such translations would look like is not immediately clear, however, and should be investigated. Related to this point is the notion of syntax splitting (Parikh 1999), where unrelated parts of a knowledge base can be split up and evaluated independently. Syntax splitting has recently been examined for conditional knowledge bases (Heyninck et al. 2023; Kern-Isberner, Beierle, and Brewka 2020), and studying whether similar ideas can be used for choice logics may provide interesting results.

Lastly, to facilitate practical applications, it will be useful to design and implement efficient algorithms for preferred model entailment. To the best of our knowledge, this has not been done yet. The only implementation of choice logics we are aware of is an encoding in Answer Set Programming (Bernreiter, Maly, and Woltran 2020), which, however, only concerns itself with single choice logic formulas, not with preferred model entailment.

CHAPTER 6

From Choice Logics to Abstract Argumentation

In this chapter, we examine the connection between choice logics and abstract argumentation, and show that they are more closely related than previously known. To this end, we use SETAFs—argumentation frameworks with sets of attacking arguments (Nielsen and Parsons 2006). SETAFs have been in the focus of researchers recently as a more flexible and expressive formalism than standard Dung-style AFs (Dvořák, Fandinno, and Woltran 2019), with intuitive connections to structured argumentation and other related formalisms (König, Rapberger, and Ulbricht 2022), while preserving many desired properties of standard AFs (Dvořák et al. 2024; Flouris and Bikakis 2019).

Despite the differences between choice logics and abstract argumentation, a first connection between them has been established by a translation (Sedki 2015) from PQCL-theories[12] (Benferhat and Sedki 2008) to Value-based AFs (VAFs) (Atkinson and Bench-Capon 2021). While this translation is a valuable first step in connecting choice logics and argumentation, it leaves some issues unaddressed. Firstly, the translation is not syntactic, as each interpretation relevant to a formula is translated into an argument. This implies that the translation is also not polynomial in size. Secondly, only the lexicographic preferred model semantics is considered, while other methods such as the inclusion- or minmax-based approaches are not studied. Thirdly, the translation relies on a redefinition of VAF-semantics which is not commonly used elsewhere.

Contributions. We address the challenges discussed above by providing two purely syntactic and polynomial-size translations from QCL-theories to SETAFs. Depending on the translation, either the inclusion-based or the minmax preferred models of the

[12] PQCL redefines the semantics of the classical connectives, but defines ordered disjunction in the same way as QCL.

original QCL-theory are in direct correspondence to the semi-stable extensions of the constructed SETAF. Moreover, we do not rely on any redefinition of standard SETAF notions. Our work shows that abstract argumentation is well-suited to directly capture formalisms in which hard- and soft-constraints are jointly represented. Moreover, using our translation, preferred model entailment in QCL can be decided using existing solvers for SETAFs (Dvořák, Greßler, and Woltran 2018). Choice logics thus join many other logic-based formalisms that have been studied with respect to their connection to formal argumentation (Bienvenu and Bourgaux 2020; Bochman 2018; Cyras and Toni 2016; Falappa et al. 2011; Modgil and Prakken 2014; Skiba and Thimm 2022; Wyner, Bench-Capon, and Dunne 2013).

Publications. This chapter is based on the paper (Bernreiter and König 2023). Newly added in this version are additional examples and full proofs for all results.

Outline. In Section 6.1 we formally define SETAFs. In Section 6.2 we show how SETAFs can be used to encode single choice logic formulas. We then build upon this encoding in Section 6.3 to capture preferred model entailment under the minmax and inclusion-based preferred model semantics. We conclude in Section 6.4.

Required preliminaries. Before reading this chapter, it is recommended to read Section 2.1 (propositional logic), Section 2.2 (computational complexity), Subsection 2.3.1 (abstract argumentation), Section 2.4 (choice logics), and Section 5.1 (formal definition of preferred model entailment).

6.1 Argumentation Frameworks with Collective Attacks

Nielsen and Parsons (2006) introduced Argumentation Frameworks with *collective attacks* (SETAFs), a generalization of standard AFs (see Section 2.3.1) where arguments cannot only be attacked by a single argument but also by *sets* of arguments.[13]

Definition 6.1 (SETAF)**.** *A SETAF is a pair $SF = (Arg, Att)$ where Arg is a set of* arguments *and $Att \subseteq (2^{Arg} \setminus \{\emptyset\}) \times Arg$ is the* attack relation.

SETAFs $SF = (Arg, Att)$, where for all $(T, h) \in Att$ it holds that $|T| = 1$, amount to (standard Dung) AFs. We usually write (t, h) to denote the set-attack $(\{t\}, h)$. For $SF_1 = (Arg_1, Att_1), SF_2 = (Arg_2, Att_2)$ we define the *union* $SF_1 \cup SF_2$ as $(Arg_1 \cup Arg_2, Att_1 \cup Att_2)$. If there is an attack $(T, h) \in Att$ with $T \subseteq S \subseteq Arg$ and $h \in S' \subseteq Arg$, we write $S \mapsto_{Att} S'$ (or simply $S \mapsto S'$).

[13]Note that, in this chapter, we denote the set of arguments in a framework by Arg instead of A and the set of attacks in a framework by Att instead of R, since upper case letters such as A, B, L, R will be used to denote choice logic formulas.

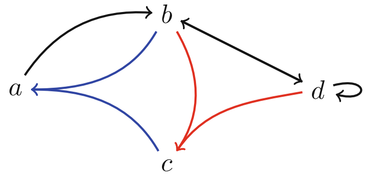

Figure 6.1: Example SETAF. Collective attacks are colored.

SETAF semantics are defined largely analogously to standard AF semantics. In this work, we make use of conflict-free (cf), admissible (adm), stable (stb), and semi-stable (sem) semantics for SETAFs (Flouris and Bikakis 2019).

Definition 6.2 (SETAF Semantics)**.** *Let $SF = (Arg, Att)$ be a SETAF and $E \subseteq Arg$. E is* conflict-free *in SF, written as $E \in cf(SF)$, if $E \not\mapsto E$. An argument $a \in Arg$ is* defended *in SF by a set $S \subseteq Arg$ if $S \mapsto B$ for each $B \subseteq Arg$ such that $B \mapsto \{a\}$. A set $T \subseteq Arg$ is defended in SF by S if each $a \in T$ is defended in SF by S. $E^{\oplus}_{SF} = E \cup \{a \in Arg \mid E \mapsto a\}$ is called the range of E (in SF). Let $S \in cf(SF)$. Then*

- *$S \in adm(SF)$ iff S defends itself in SF;*
- *$S \in stb(SF)$ iff $S \mapsto \{a\}$ for all $a \in Arg \setminus S$;*
- *$S \in sem(SF)$ iff $S \in adm(SF)$ and there is no $T \in adm(SF)$ such that $T^{\oplus} \supset S^{\oplus}$.*

Example 6.3. *Figure 6.1 shows the SETAF $SF = (Arg, Att)$ with*

$$
\begin{aligned}
Arg &= \{a, b, c, d\} \\
Att &= \{(a, b), (b, d), (d, b), (d, d), (\{b, c\}, a), (\{b, d\}, c)\}.
\end{aligned}
$$

It can be confirmed that

$$
\begin{aligned}
cf(SF) &= \{\emptyset, \{a\}, \{b\}, \{c\}, \{a, c\}, \{b, c\}\}, \\
adm(SF) &= \{\emptyset, \{a\}, \{c\}, \{a, c\}, \{b, c\}\}, \\
stb(SF) &= sem(SF) = \{\{b, c\}\}.
\end{aligned}
$$

If SF is clear from the context, we will simply write $E^{\oplus}$ instead of $E^{\oplus}_{SF}$ to denote the range of E in SF.

The reasoning task for SETAFs that is of relevance in this chapter is that of skeptical acceptance for semi-stable semantics, which is defined analogously to the case of standard AFs (see Definition 2.8): $a \in Arg$ is skeptically accepted in $SF = (Arg, Att)$ w.r.t. semi-stable semantics iff $a \in S$ for every $S \in sem(SF)$. Specifically, we are interested in a variant of skeptical acceptance where an argument a is also accepted if $sem(SF) = \{\emptyset\}$.

Moreover, we will make use of the fact that in SETAFs, analogously to standard AFs, $stb(SF) \neq \emptyset$ implies $stb(SF) = sem(SF)$. This is easy to see, since if there is a stable

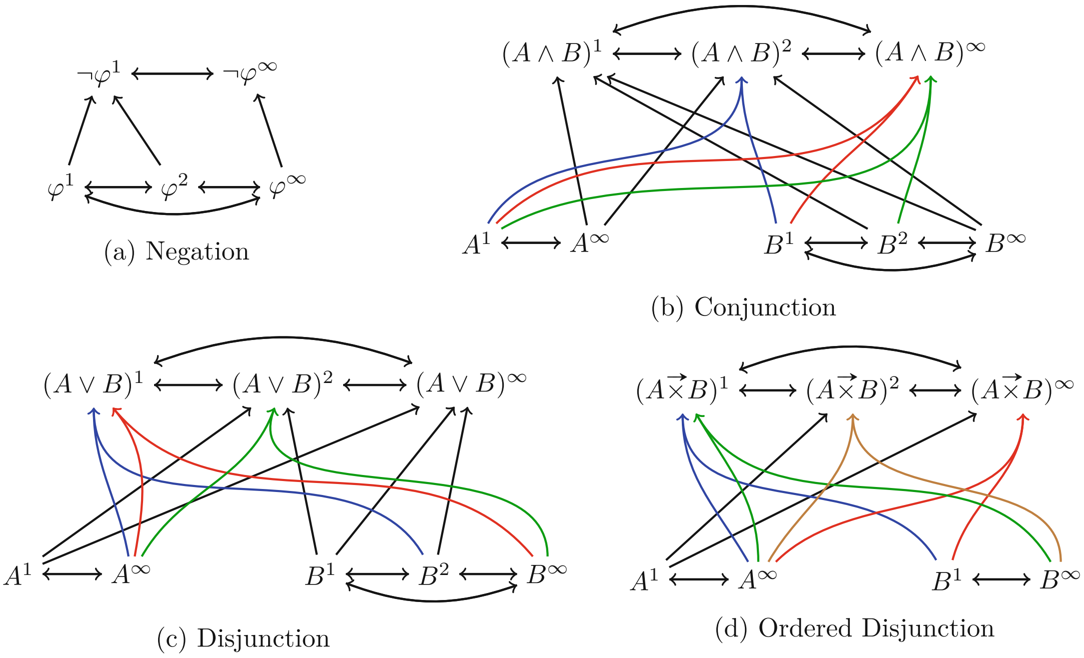

Figure 6.2: QCL-connectives encoded as SETAFs.

extension $E \in stb(SF)$ for $SF = (Arg, Att)$, then $E^{\oplus} = Arg$ and therefore we must have $E'^{\oplus} = Arg$ for all semi-stable extensions $E' \in sem(SF)$.

6.2 Encoding QCL-formulas

We aim to capture QCL-theories via SETAFs such that the preferred models of the initial theory correspond to the extensions of the constructed framework. As a first step, we encode single QCL-formulas to obtain a correspondence between the satisfaction degree ascribed to a formula by an interpretation and the (semi-)stable extensions of the target SETAF. This intermediate step is needed to deal with the monotonic nature of satisfaction degrees, upon which the non-monotonic notion of preferred models is built. Note that a similar intermediate step is utilized in the sequent calculus for QCL presented in Section 5.4, where the calculus for preferred model entailment is built on a labeled monotonic calculus.

The following notation will be used from now on: By $var(\varphi)$ we denote the set of variables occurring in a QCL-formula φ, while $sf(\varphi)$ denotes the set of all subformulas of φ. By $pdeg(\varphi) = \{1, \ldots, opt_{\text{QCL}}(\varphi)\} \cup \{\infty\}$ we denote the set of possible satisfaction degrees that φ may assume. Likewise, for a QCL-theory $T = \{\varphi_1, \ldots, \varphi_t\}$ we let $pdeg(T) = \{1, \ldots, max(opt_{\text{QCL}}(\varphi_1), \ldots, opt_{\text{QCL}}(\varphi_t))\} \cup \{\infty\}$. Note that $opt_{\text{QCL}}(\neg\varphi) = 1$ and thus $pdeg(\neg\varphi) = \{1, \infty\}$ for every QCL-formula φ. This reflects the fact that negation in QCL acts only on truth, but not on preferences (cf. Definition 2.29).

Intuitively, our encoding works as follows: given a QCL-formula φ, we will add arguments ψ^k for each subformula $\psi \in sf(\varphi)$ and each degree $k \in pdeg(\psi)$. Every ψ^k will attack all other ψ^ℓ with $\ell \neq k$ to ensure that only one of $\psi^1, \ldots, \psi^{opt_{\mathrm{QCL}}(\psi)}, \psi^\infty$ can be accepted. Moreover, we add attacks between each ψ^k and the immediate subformulas of ψ according to the degree-semantics of QCL. This will ensure that ψ^k is accepted in a (semi-)stable extension E iff ψ is satisfied to a degree of k in the interpretation $\mathcal{I}$ corresponding to E. For instance, if $\psi = (a \vec{\times} b)$ is satisfied to a degree of 2 by $\mathcal{I}$ (i.e., $\mathcal{I} \models_2^{\mathrm{QCL}} \psi$), then the argument $(a \vec{\times} b)^2$ will be accepted in the corresponding extension E, but $(a \vec{\times} b)^1$ and $(a \vec{\times} b)^\infty$ will be defeated. We now formally specify our translation. Figure 6.2 depicts the encoding for each of the four connectives.

Definition 6.4. *Let φ be a QCL-formula. We define the* corresponding SETAF $SF_\varphi = (Arg_\varphi, Att_\varphi)$ *with arguments*

$$Arg_\varphi = \{\psi^o \mid \psi \in sf(\varphi), o \in pdeg(\varphi)\}$$
$$Att_\varphi = \Big(\bigcup_{\psi \in sf(\varphi)} Att^*_\psi \Big) \cup \{(\psi^o, \psi^p) \mid \psi \in sf(\varphi), o \neq p\}$$

where Att^*_ψ depends on the immediate subformulas of ψ. For $a \in \mathcal{U}$ we have $Att^*_a = \emptyset$. Otherwise, we have

$$Att^*_{\neg L} = \{(L^\infty, \neg L^\infty)\} \cup \{(L^\ell, \neg L^1) \mid \ell \neq \infty\};$$
$$Att^*_{(L \wedge R)} = \{(\{L^\ell, R^r\}, (L \wedge R)^d) \mid d > max(\ell, r)\} \cup$$
$$\{(L^\ell, (L \wedge R)^d) \mid \ell > d\} \cup \{(R^r, (L \wedge R)^d) \mid r > d\};$$
$$Att^*_{(L \vee R)} = \{(\{L^\ell, R^r\}, (L \vee R)^d) \mid d < min(\ell, r)\} \cup$$
$$\{(L^\ell, (L \vee R)^d) \mid \ell < d\} \cup \{(R^r, (L \vee R)^d) \mid r < d\};$$
$$Att^*_{(L \vec{\times} R)} = \{(L^\ell, (L \vec{\times} R)^d) \mid \ell \neq \infty, \ell \neq d\} \cup$$
$$\{(\{L^\infty, R^r\}, (L \vec{\times} R)^d) \mid r \neq \infty, d \neq r + opt_{\mathrm{QCL}}(L)\} \cup$$
$$\{(\{L^\infty, R^\infty\}, (L \vec{\times} R)^d) \mid d \neq \infty\}.$$

Example 6.5. *Let $\varphi = (a \vec{\times} b) \wedge \neg c$. The SETAF SF_φ corresponding to φ is depicted in Figure 6.3. Now consider the interpretation $\mathcal{I} = \{b\}$. Note that, in Figure 6.3, arguments ψ^k corresponding to a subformula $\psi \in sf(\varphi)$ are highlighted iff $deg_{\mathrm{QCL}}(\mathcal{I}, \psi) = k$. Also note that these arguments correspond to a stable extension in SF_φ.*

The construction specified in Definition 6.4 is purely syntactic, since $opt_{\mathrm{QCL}}(\varphi)$, and therefore $pdeg(\varphi)$, can be computed based solely on the structure of φ (cf. Definition 2.28). Moreover, the construction is polynomial, since $opt_{\mathrm{QCL}}(\varphi)$ is bounded by the number of $\vec{\times}$-occurrences in φ. Thus, SF_φ contains $O(opt_{\mathrm{QCL}}(\varphi) \cdot |sf(\varphi)|)$ arguments. Furthermore, attacks in SF_φ never have more than two joint attackers, hence, $|Att_\varphi|$ is polynomial in $|Arg_\varphi|$.

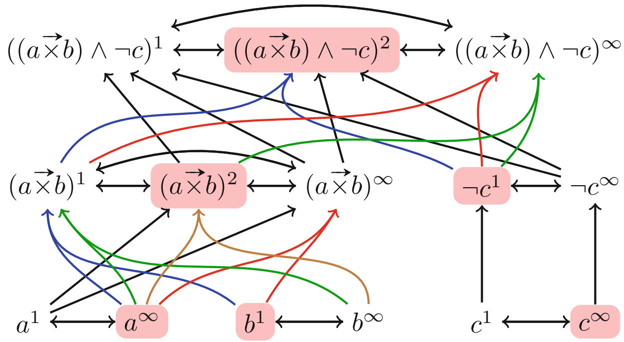

Figure 6.3: The QCL-formula $(a\vec{\times}b) \wedge \neg c$ from Example 6.5 encoded as a SETAF.

We now establish the semantic correspondence between a QCL-formula φ and the SETAF SF_φ. We write $\mathcal{I} \cong E$ for an interpretation $\mathcal{I}$ and an extension E if $\mathcal{I}$ corresponds to the choice of arguments in E i.e., $a^1 \in E$ iff $a \in \mathcal{I}$ and $a^\infty \in E$ iff $a \notin \mathcal{I}$. Likewise, for a set $\mathcal{M} \subseteq 2^{\mathcal{U}}$ of interpretations and a set $\sigma(SF)$ of extensions we write $\mathcal{M} \cong \sigma(SF)$ iff for every $\mathcal{I} \in \mathcal{M}$ there is exactly one $E \in \sigma(SF)$ such tat $\mathcal{I} \cong E$, and for every $E' \in \sigma(SF)$ there is exactly one $\mathcal{I}' \in \mathcal{M}$ such that $\mathcal{I}' \cong E'$.

Lemma 6.6. *Let φ be a QCL-formula and SF_φ its corresponding SETAF. If $\mathcal{I} \cong E$ for $\mathcal{I} \subseteq var(\varphi)$ and $E \in stb(SF_\varphi)$ then for all $\psi \in sf(\varphi)$ we have $\mathcal{I} \models_k^{\mathrm{QCL}} \psi$ iff $\psi^k \in E$.*

Proof. Let φ be a QCL-formula and $SF_\varphi = (Arg_\varphi, Att_\varphi)$ its corresponding SETAF (cf. Definition 6.4). Consider $\mathcal{I} \subseteq var(\varphi)$ and $E \in stb(SF_\varphi)$ such that $\mathcal{I} \cong E$. We proceed by structural induction.

Induction base: by definition, it holds that if $\mathcal{I} \cong E$ then $\mathcal{I} \models_1^{\mathrm{QCL}} a$ iff $a^1 \in E$ and $\mathcal{I} \models_\infty^{\mathrm{QCL}} a$ iff $a^\infty \in E$.

Induction step: as the induction hypothesis (I.H.), assume that for $L \in sf(\varphi)$ we have $\mathcal{I} \models_\ell^{\mathrm{QCL}} L$, $L^\ell \in E$, but $L^k \notin E$ for all $k \in pdeg(L) \setminus \{\ell\}$. Likewise, for $R \in sf(\varphi)$ we have $\mathcal{I} \models_r^{\mathrm{QCL}} R$, $R^r \in E$, but $R^k \notin E$ for all $k \in pdeg(R) \setminus \{r\}$. We consider each connective.

- $\neg L$: there are two possible cases.
  - $\ell < \infty$. Then $\mathcal{I} \models_\infty^{\mathrm{QCL}} \neg L$. By the I.H. we have $L^\ell \in E$. By construction, there is no conflict between L^ℓ and $\neg L^\infty$, but there is one between L^ℓ and $\neg L^1$, i.e., $\neg L^1 \notin E$. Observe that $\neg L^\infty$ is only attacked by L^∞ and $\neg L^1$, both of which are defeated by E. Thus, for E to be stable, it must be that $\neg L^\infty \in E$.
  - $\ell = \infty$. Then $\mathcal{I} \models_1^{\mathrm{QCL}} \neg L$. By the I.H. we have $L^\infty \in E$. By construction, there is no conflict between L^∞ and $\neg L^1$, but there is one between L^∞ and $\neg L^\infty$, i.e, $\neg L^\infty \notin E$. Observe that $\neg L^1$ is only attacked by arguments L^k with $k \neq \infty$ and $\neg L^\infty$, all of which are defeated by E. Thus, for E to be stable, it must be that $\neg L^1 \in E$.

- $(L \wedge R)$: then $\mathcal{I} \models_d^{\text{QCL}} (L \wedge R)$ with $d = max(\ell, r)$. By construction, there is no conflict between the arguments L^ℓ, R^r, and $(L \wedge R)^d$. Now consider $e \in pdeg((L \wedge R))$ such that $e \neq d$. Clearly, $\mathcal{I} \not\models_e^{\text{QCL}} (L \wedge R)$. There are two cases:
  - $e > d = max(\ell, r)$. Then, by construction, $(\{L^\ell, R^r\}, (L \wedge R)^e) \in Att_\varphi$. Since $E \in cf(SF_\varphi)$, it must be that $(L \wedge R)^e \notin E$.
  - $e < d = max(\ell, r)$. Then, either $\ell > e$ or $r > e$. Assume $\ell > e$ (the case that $r > e$ is analogous). Then, by construction, $(L^\ell, (L \wedge R)^e) \in Att_\varphi$, i.e., it must be that $(L \wedge R)^e \notin E$.

  Moreover, $(L \wedge R)^d$ is defended against the attacks from arguments L^x with $x \neq \ell$ and R^y with $y \neq r$, as these arguments are counter-attacked by either L^ℓ or R^r. Thus, for E to be stable, it must be that $(L \wedge R)^d \in E$.

- $(L \vee R)$: then $\mathcal{I} \models_d^{\text{QCL}} (L \wedge R)$ with $d = min(\ell, r)$. By construction, there is no conflict between the arguments L^ℓ, R^r, and $(L \vee R)^d$. Now consider $e \in pdeg((L \vee R))$ such that $e \neq d$. Clearly, $\mathcal{I} \not\models_e^{\text{QCL}} (L \vee R)$. There are two cases:
  - $e < d = min(\ell, r)$. Then, by construction, $(\{L^\ell, R^r\}, (L \vee R)^e) \in Att_\varphi$. Since $E \in cf(SF_\varphi)$, it must be that $(L \vee R)^e \notin E$.
  - $e > d = min(\ell, r)$. Then, either $\ell < e$ or $r < e$. Assume $\ell < e$ (the case that $r < e$ is analogous). Then, by construction, $(L^\ell, (L \vee R)^e) \in Att_\varphi$, i.e., it must be that $(L \vee R)^e \notin E$.

  Moreover, $(L \vee R)^d$ is defended against the attacks from arguments L^x with $x \neq \ell$ and R^y with $y \neq r$, as these arguments are counter-attacked by either L^ℓ or R^r. Thus, for E to be stable, it must be that $(L \vee R)^d \in E$.

- $(L \vec{\times} R)$: we must distinguish the following cases:
  - $\ell < \infty$. Then $\mathcal{I} \models_\ell^{\text{QCL}} (L \vec{\times} R)$. By construction, there is no conflict between the arguments L^ℓ, R^r, and $(L \vec{\times} R)^\ell$. However, for all $e \in pdeg((L \vec{\times} R))$ with $e \neq \ell$ we have $(L^\ell, (L \vec{\times} R)^e) \in Att_\varphi$, i.e., $(L \vec{\times} R)^e \notin E$. Thus, for E to be stable, it must be that $(L \vec{\times} R)^\ell \in E$.
  - $\ell = \infty$ and $r < \infty$. Then $\mathcal{I} \models_d^{\text{QCL}} ((L \vec{\times} R))$ with $d = opt_{\text{QCL}}(L) + r$. By construction, there is no conflict between the arguments L^∞, R^r, and $(L \vec{\times} R)^d$. However, for all $e \in pdeg((L \vec{\times} R))$ with $e \neq d$ we have $(\{L^\infty, R^r\}, (L \vec{\times} R)^e) \in Att_\varphi$, i.e., $(L \vec{\times} R)^e \notin E$. Thus, for E to be stable, it must be that $(L \vec{\times} R)^d \in E$.
  - $\ell = \infty$ and $r = \infty$. Then $\mathcal{I} \models_\infty^{\text{QCL}} (L \vec{\times} R)$. By construction, there is no conflict between the arguments L^∞, R^∞, and $(L \vec{\times} R)^\infty$. However, for all $e \in pdeg((L \vec{\times} R))$ with $e \neq \infty$ we have $(\{L^\infty, R^\infty\}, (L \vec{\times} R)^e) \in Att_\varphi$, i.e., $(L \vec{\times} R)^e \notin E$. Thus, for E to be stable, it must be that $(L \vec{\times} R)^\infty \in E$. □

As a result of the above lemma, each interpretation relevant to a formula φ corresponds to exactly one stable extension in SF_φ, and vice versa.

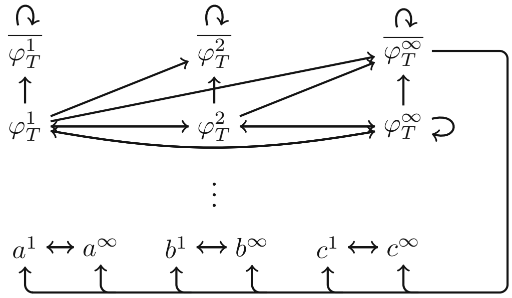

Figure 6.4: The QCL-theory from Example 6.9 encoded as a SETAF (minmax semantics).

Proposition 6.7. $2^{var(\varphi)} \cong stb(SF_\varphi) = sem(SF_\varphi)$.

Proof. Let φ be a QCL-formula and $SF_\varphi = (Arg_\varphi, Att_\varphi)$ its corresponding SETAF, as constructed in Definition 6.4. We first show that $2^{var(\varphi)} \cong stb(SF_\varphi)$. For any $E \in stb(SF_\varphi)$ there is exactly one $\mathcal{I}$ such that $\mathcal{I} \cong E$, since, by construction and the fact that E is stable, for each $a \in var(\varphi)$ we must have either $a^1 \in E$ or $a^\infty \in E$. Now consider any $\mathcal{I} \in 2^{var(\varphi)}$. Let $E = \{\psi^k \mid \psi \in sf(\varphi), \mathcal{I} \models_k^{\text{QCL}} \psi\}$. Clearly, $\mathcal{I} \cong E$. Moreover, by construction we have that $E \in cf(SF_\varphi)$ and that every $\psi^k \in E$ attacks every ψ^ℓ with $\ell \neq k$. This further implies that $E \in stb(SF_\varphi)$. Indeed, if $E' \in stb(SF_\varphi)$ and $\mathcal{I} \cong E'$ then, by Lemma 6.6, $E' = E$. Finally, by the fact that $2^{var(\varphi)} \neq \emptyset$ it follows that $stb(SF_\varphi) \neq \emptyset$, which in turn yields $stb(SF_\varphi) = sem(SF_\varphi)$. □

Note that we can also capture only the *models* of a formula φ, instead of all interpretations relevant for φ, by adding the attack $(\varphi^\infty, \varphi^\infty)$ to SF_φ: if φ is (classically) unsatisfiable, then we will have no stable extensions.

6.3 Capturing Preferred Models

We now extend our construction for QCL-formulas from Section 6.2 to also capture QCL-theories and their preferred models. This then further allows us to decide the problem of preferred model entailment via the constructed framework.

First, we consider preferred models w.r.t. the minmax (mm) semantics (cf. Definition 5.1), where a theory $T = \{\varphi_1, \dots, \varphi_t\}$ is semantically equivalent to the formula $\varphi_T = (\varphi_1 \wedge (\cdots \wedge \varphi_t))$. The key idea is the following: we first construct the SETAF SF_{φ_T} corresponding to φ_T (cf. Definition 6.4). Then, for each argument φ_T^k we introduce a self-attacking argument $\overline{\varphi_T^k}$. Each $\overline{\varphi_T^k}$ is attacked by every φ_T^ℓ such that $\ell \leq k$. As a result, if we consider two admissible sets E, E' such that $\varphi_T^\ell \in E$, $\varphi_T^k \in E'$, and $\ell < k$, then the range $E^\oplus$ of E is a superset of the range $E'^\oplus$ of E' w.r.t. to the arguments $\overline{\varphi_T^m}$, $m \in pdeg(\varphi_T)$. This then means that the semi-stable extensions of the constructed framework correspond to the minmax preferred models of the initial theory. Finally, we add attacks from

$\overline{\varphi_T^\infty}$ to all variable-arguments a^1, a^∞ where $a \in var(\varphi_T)$. This ensures that, if T is not classically satisfiable, the only semi-stable extension of SF_T^{mm} is $\emptyset$. We now provide this construction formally.

Definition 6.8. *Let $T = \{\varphi_1, \ldots, \varphi_t\}$ be a QCL-theory. Let $\varphi_T = (\varphi_1 \wedge (\cdots \wedge \varphi_t))$, and let $SF_{\varphi_T} = (Arg_{\varphi_T}, Att_{\varphi_T})$ be the SETAF corresponding to φ_T. We define $SF_T^{mm} = (Arg_T^{mm}, Att_T^{mm})$ as follows:*

$$
\begin{aligned}
Arg_T^{mm} &= Arg_{\varphi_T} \cup \{\overline{\varphi_T^o} \mid o \in pdeg(\varphi_T)\} \\
Att_T^{mm} &= Att_{\varphi_T} \cup \{(\varphi_T^\infty, \varphi_T^\infty)\} \\
&\cup \{(\overline{\varphi_T^o}, \overline{\varphi_T^o}) \mid o \in pdeg(\varphi_T)\} \\
&\cup \{(\overline{\varphi_T^\infty}, a^1), (\overline{\varphi_T^\infty}, a^\infty) \mid a \in var(\varphi_T)\} \\
&\cup \{(\varphi_T^o, \overline{\varphi_T^p}) \mid o, p \in pdeg(\varphi_i), o \leq p\}.
\end{aligned}
$$

Example 6.9. *Let $T = \{(a \vec{\times} c), (b \vec{\times} c), \neg(a \wedge b)\}$. Then $\varphi_T = ((a \vec{\times} c) \wedge ((b \vec{\times} c) \wedge \neg(a \wedge b)))$ with $pdeg(\varphi_T) = \{1, 2, \infty\}$. SF_T^{mm} is depicted in Figure 6.4. Arguments ψ^k corresponding to non-atomic subformulas ψ of φ_T are not depicted for the sake of succinctness.*

There is a direct semantic correspondence between the initial theory T and the constructed framework SF_T^{mm}, namely, each preferred model of T corresponds to exactly one semi-stable extension of SF_T^{mm}, and vice versa.

Proposition 6.10. *$Prf_{\mathrm{QCL}}^{mm}(T) \cong sem(SF_T^{mm}) \setminus \{\emptyset\}$.*

Proof. Let T be a QCL-theory, and consider $SF_{\varphi_T} = (Arg_{\varphi_T}, Att_{\varphi_T})$ and $SF_T^{mm} = (Arg_T^{mm}, Att_T^{mm})$ constructed according to Definition 6.8.

"$\subseteq$": Let $\mathcal{I} \in Prf_{\mathrm{QCL}}^{mm}(T)$. Let E be the corresponding set of arguments, i.e., $E = \{\psi^d \in Arg_{\varphi_T} \mid \psi \in sf(\varphi_T), \mathcal{I} \models_d^{\mathrm{QCL}} \psi\}$. By Proposition 6.7 we know that $E \in stb(SF_{\varphi_T})$ and thus every argument in Arg_{φ_T} is either in E or attacked by E. Since $\mathcal{I}$ is a classical model of T we have $\varphi_T^k \in E$ for some $k \in pdeg(\varphi_T) \setminus \{\infty\}$. Moreover, $E \in adm(SF_T^{mm})$ since for every $a \in var(\varphi_T)$, the arguments a^1, a^∞ are defended against the attacks from $\overline{\varphi_T^\infty}$ in E. We want to show that there is no other admissible set E' in SF_T^{mm} such that $E'^{\oplus} \supset E^{\oplus}$. Towards a contradiction assume such an E' exists. Then $Arg_{\varphi_T} \subseteq E'^{\oplus}$. In fact, since the only way to achieve this is to have either $a^1 \in E'$ or $a^\infty \in E'$ for each $a \in var(T)$, we also have $E' \cap Arg_{\varphi_T} \in stb(SF_{\varphi_T})$. As a consequence of Proposition 6.7 this means that E' accepts different arguments that correspond to the atoms in T—let $\mathcal{I}'$ be the corresponding model, i.e., $\mathcal{I}' \cong E'$. Clearly $\mathcal{I}'$ is a classical model of T, as otherwise φ_T^∞ is not attacked by E', which contradicts the assumption that $E'^{\oplus} \supset E^{\oplus}$. By construction, the only arguments in Arg_T^{mm} not contained in $E^{\oplus}$ are $\overline{\varphi_T^1}, \ldots, \overline{\varphi_T^{k-1}}$, where k is such that $\mathcal{I} \models_k^{\mathrm{QCL}} \varphi_T$. Thus, $\overline{\varphi_T^\ell} \in E'^{\oplus}$ for some $\ell < k$. Since $\overline{\varphi_T^\ell}$ is self-attacking this means there must be some $m \leq \ell$ such that $\varphi_T^m \in E'$. But then, by Lemma 6.6 applied to SF_{φ_T}, $\mathcal{I}' \models_m^{\text{QCL}} \varphi_T$ which means that $\mathcal{I} \notin \mathit{Prf}_{\text{QCL}}^{mm}(T)$. Contradiction.

"$\supseteq$": Assume $E \in \mathit{sem}(SF_T^{mm}) \setminus \{\emptyset\}$. Let $\mathcal{I}$ be the corresponding model of T, i.e., $\mathcal{I} \cong E$. Towards a contradiction, assume $\mathcal{I} \notin \mathit{Prf}_{\text{QCL}}^{mm}(T)$. Then there are two cases:

1. T is classically satisfiable. Then, since $\mathcal{I} \notin \mathit{Prf}_{\text{QCL}}^{mm}(T)$, there is an interpretation $\mathcal{I}'$ such that $\mathit{deg}_{\text{QCL}}(\mathcal{I}', \varphi_T) < \mathit{deg}_{\text{QCL}}(\mathcal{I}, \varphi_T)$. Let $E' = \{\psi^d \in \mathit{Arg}_T^{mm} \mid \psi \in \mathit{sf}(\varphi_T), \mathcal{I}' \models_d^{\text{QCL}} \psi\}$. By Proposition 6.7 and Lemma 6.6 we know that $E' \in \mathit{stb}(SF_{\varphi_T})$. Indeed, by the same reasoning as in the "$\subseteq$"-case, $E' \in \mathit{adm}(SF_T^{mm})$ and $E'^{\oplus} \supset E^{\oplus}$. Contradiction to $E \in \mathit{sem}(SF_T^{mm})$.

2. T is classically unsatisfiable. We show that then $\mathit{adm}(SF_T^{mm}) = \{\emptyset\}$, which contradicts $E \in \mathit{sem}(SF_T^{mm}) \setminus \{\emptyset\}$. Towards a contradiction, assume that there is some $E' \in \mathit{adm}(SF_T^{mm})$ such that $E' \neq \emptyset$. Thus, $\psi^d \in E'$ for some ψ such that one of the following is true: $\psi = a$ with $a \in \mathit{var}(\varphi)$; $\psi = \neg L$; $\psi = (L \circ R)$ with $\circ \in \{\wedge, \vee, \vec{\times}\}$. If $\psi \neq a$ for $a \in \mathit{var}(\varphi)$ then ψ^d must be defended against the attacks from its immediate subformulas L and R. But L and R must in turn also be defended, and so on. We can conclude that there is at least some $a \in \mathit{var}(\varphi_T)$ such that $a^1 \in E'$ or $a^\infty \in E'$. Note that a^1 and a^∞ are attacked by $\overline{\varphi_T^\infty}$. Since $E' \in \mathit{adm}(SF_T^{mm})$, it must be that $\varphi_T^k \in E'$ for some $k < \infty$. But this means by the same reasoning as above that $\mathcal{I}' \models_k^{\text{QCL}} \varphi_T$ for the interpretation $\mathcal{I}'$ such that $\mathcal{I}' \cong E'$. This contradicts T being classically unsatisfiable. □

We now turn our attention to the inclusion-based (*inc*) preferred model semantics (cf. Definition 5.1). In essence, we can build upon the tools established so far and use the same gadget as in the case of minmax semantics to minimize satisfaction degrees. However, this gadget is now constructed for every $\varphi \in T$, i.e., we add $\overline{\varphi^k}$ for each $\varphi \in T$ and each $k \in \mathit{pdeg}(\varphi)$.

Definition 6.11. *Let $T = \{\varphi_1, \ldots, \varphi_t\}$ be a QCL-theory and $SF_1 = (\mathit{Arg}_1, \mathit{Att}_1), \ldots, SF_t = (\mathit{Arg}_t, \mathit{Att}_t)$ the SETAFs corresponding to $\varphi_1, \ldots, \varphi_t$. Let $SF_T^{inc} = (\mathit{Arg}_T^{inc}, \mathit{Att}_T^{inc})$ s.t.:*

$$
\begin{aligned}
\mathit{Arg}_T^{inc} &= \Big(\bigcup_{1 \leq i \leq t} \mathit{Arg}_i \Big) \cup \{\overline{\varphi_i^o} \mid \varphi_i \in T, o \in \mathit{pdeg}(\varphi_i)\} \\
\mathit{Att}_T^{inc} &= \Big(\bigcup_{1 \leq i \leq t} \mathit{Att}_i \Big) \cup \{(\varphi_i^\infty, \varphi_i^\infty) \mid \varphi_i \in T\} \\
&\quad \cup \{(\overline{\varphi_i^o}, \overline{\varphi_i^o}) \mid \varphi_i \in T, o \in \mathit{pdeg}(\varphi_i)\} \\
&\quad \cup \{(\overline{\varphi_i^\infty}, a^1), (\overline{\varphi_i^\infty}, a^\infty) \mid \varphi_i \in T, a \in \mathit{var}(\varphi_i)\} \\
&\quad \cup \{(\varphi_i^o, \overline{\varphi_i^p}) \mid \varphi_i \in T, o \in \mathit{pdeg}(\varphi_i), o \leq p\}.
\end{aligned}
$$

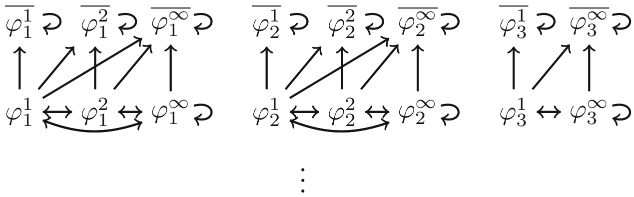

Figure 6.5: The QCL-theory from Example 6.12 encoded as a SETAF (inclusion-based semantics).

Example 6.12. *Let* $T = \{\varphi_1, \varphi_2, \varphi_3\}$ *with* $\varphi_1 = (a \vec{\times} c)$, $\varphi_2 = (b \vec{\times} c)$, *and* $\varphi_3 = \neg(a \wedge b)$. *To obtain* SF_T^{inc} *we construct a minimization gadget for each* $\varphi_i \in T$, *as depicted in Figure 6.5. For succinctness, we omit arguments corresponding to subformulas of each* $\varphi_i \in T$.

Analogously to Proposition 6.10, every preferred model of some QCL-theory T corresponds to exactly one semi-stable extension of SF_T^{inc}, and vice versa.

Proposition 6.13. $\mathit{Prf}_{\mathrm{QCL}}^{inc}(T) \cong \mathit{sem}(SF_T^{inc}) \setminus \{\emptyset\}$.

Proof. Let $T = \{\varphi_1, \ldots, \varphi_t\}$ be a QCL-theory, and consider $SF_1 = (\mathit{Arg}_1, \mathit{Att}_1), \ldots, SF_t = (\mathit{Arg}_t, \mathit{Att}_t)$ and $SF_T^{inc} = (\mathit{Arg}_T^{inc}, \mathit{Att}_T^{inc})$ constructed according to Definition 6.11.

"$\subseteq$": Let $\mathcal{I} \in \mathit{Prf}_{\mathrm{QCL}}^{inc}(T)$. Let E be the corresponding set of arguments, i.e., $E = \{\psi^d \in \mathit{Arg}_T^{inc} \mid \psi \in \mathit{sf}(T), \mathcal{I} \models_d^{\mathrm{QCL}} \psi\}$[14]. By Proposition 6.7 and Lemma 6.6 we know that $E \cap \mathit{Arg}_t \in \mathit{stb}(SF_t)$. Since $\mathcal{I}$ is a classical model of T, for each $\varphi_i \in T$ there is some $k \in \mathit{pdeg}(\varphi_i) \setminus \{\infty\}$ such that $\varphi_i^k \in E$. Thus, for every variable a, the arguments a^1, a^∞ are defended against $\overline{\varphi_i^\infty}$ in E. Thus, $E \in \mathit{adm}(SF_T^{inc})$. We want to show that there is no other admissible set E' in SF_T^{inc} such that $E'^{\oplus} \supset E^{\oplus}$. Towards a contradiction assume such an E' exists. As a consequence of Proposition 6.7 this means that E' accepts different arguments that correspond to the atoms in T—let $\mathcal{I}'$ be the corresponding model, i.e., $\mathcal{I}' \cong E'$. Clearly $\mathcal{I}'$ is a classical model of T, as otherwise one of the self-attacking arguments φ_i^∞ would not be be attacked (contradicting the assumption that $E'^{\oplus} \supset E^{\oplus}$). Let $\overline{\varphi_j^y}$ be an argument that is attacked by E' but not by E. Clearly, we have $\overline{\varphi_j^z} \in E^{\oplus} \cap E'^{\oplus}$ for some $z > y$. Let z be the lowest such number. This means $\mathcal{I} \models_z^{\mathrm{QCL}} \varphi_j$, but $\mathcal{I}' \models_y^{\mathrm{QCL}} \varphi_j$. However, since for all φ_i with $i \neq j$ we also have $\mathcal{I} \models_m^{\mathrm{QCL}} \varphi_i$ and $\mathcal{I}' \models_n^{\mathrm{QCL}} \varphi_i$ with $n \leq m$ by the same line of reasoning, this means that $\mathcal{I} \notin \mathit{Prf}_{\mathrm{QCL}}^{inc}(T)$. Contradiction.

"$\supseteq$": Assume $E \in \mathit{sem}(SF_T^{inc}) \setminus \{\emptyset\}$. Let $\mathcal{I}$ be the corresponding model of T, i.e., $\mathcal{I} \cong E$. Towards a contradiction, assume $\mathcal{I} \notin \mathit{Prf}_{\mathrm{QCL}}^{inc}(T)$. Then there are two cases:

[14]We define for a QCL-theory T the set of subformulas as $\mathit{sf}(T) = \bigcup_{\varphi \in T} \mathit{sf}(\varphi)$.

1. T is classically satisfiable. Then there must be some interpretation $\mathcal{I}'$ that is more preferable than $\mathcal{I}$, i.e., for all φ_i it holds that $\mathcal{I} \models_m^{\text{QCL}} \varphi_i$ and $\mathcal{I}' \models_n^{\text{QCL}} \varphi_i$ with $n \leq m$ and $n < m$ for at least one i. Let $E' = \{\psi^d \in Arg_T^{inc} \mid \psi \in sf(T), \mathcal{I}' \models_d^{\text{QCL}} \psi\}$. By Proposition 6.7 and Lemma 6.6 we know that $E' \in stb(SF_i)$ for each $1 \leq i \leq t$. It follows that $E' \in adm(SF_T^{inc})$. Moreover, E' attacks more arguments than E, because the argument $\overline{\varphi_i^n}$ where $\mathcal{I} \models_m^{\text{QCL}} \varphi_i$ and $\mathcal{J} \models_n^{\text{QCL}} \varphi_i$ with $n < m$ is attacked by E', but not by E. This contradicts $E \in sem(SF_T^{inc}) \setminus \{\emptyset\}$.

2. T is classically unsatisfiable. Note that each $\overline{\varphi_i^\infty}$ attacks the variable-arguments a^1, a^∞. Thus, by the same line of reasoning as in the proof of Proposition 6.10, $adm(SF_T^{inc}) = \{\emptyset\}$. This contradicts $E \in sem(SF_T^{inc}) \setminus \{\emptyset\}$. □

We have established a semantic correspondence between the preferred models of QCL-theories (under both *mm* and *inc* semantics) and the semi-stable extensions of SETAFs. These results can now further be used to decide preferred model entailment $T \mid\!\sim_{\text{QCL}}^{\pi} \varphi$ (cf. Definition 5.5). To this end, we combine the frameworks SF_T^π for the theory T and SF_φ for the entailed (classical) formula φ.

Theorem 6.14. *Let T be a QCL-theory and $\pi \in \{mm, inc\}$. Then $T \mid\!\sim_{\text{QCL}}^{\pi} \varphi$ iff $\varphi^1 \in S$ for all $S \in sem(SF_T^\pi \cup SF_\varphi) \setminus \{\emptyset\}$.*

Proof. Let T be a QCL-theory, φ a classical formula, and let $SF_T^\pi = (Arg_T^\pi, Att_T^\pi)$ and $SF_\varphi = (Arg_\varphi, Att_\varphi)$ be the SETAFs constructed from T (cf. Definition 6.8 if $\pi = mm$, Definition 6.11 if $\pi = inc$) and φ (cf. Definition 6.4) respectively.

Assume $T \mid\!\sim_{\text{QCL}}^{\pi} \varphi$. Let $E \in sem(SF_T^\pi \cup SF_\varphi)$ and let $\mathcal{I}$ be the interpretation such that $\mathcal{I} \cong E$. By Proposition 6.7 and Lemma 6.6, $E \cap Arg_\varphi \in stb(SF_\varphi)$. Moreover, note that $E \cap Arg_T^\pi \in sem(SF_T^\pi)$ and therefore, by Proposition 6.10 (if $\pi = mm$) or Proposition 6.13 (if $\pi = inc$), also $\mathcal{I} \in Prf_{\text{QCL}}^\pi(T)$. Then, since $T \mid\!\sim_{\text{QCL}}^{\pi} \varphi$, we have $\mathcal{I} \models_1^{\text{QCL}} \varphi$. By Lemma 6.6 this further implies $\varphi^1 \in E \cap Arg_\varphi$, i.e., $\varphi^1 \in E$.

Assume $T \not\mid\!\sim_{\text{QCL}}^{\pi} \varphi$. Then there is some $\mathcal{I} \in Prf_{\text{QCL}}^\pi(T)$ such that $\mathcal{I} \not\models_1^{\text{QCL}} \varphi$. Let $E = \{\psi^k \in Arg_T^\pi \cup Arg_\varphi \mid \mathcal{I} \models_k^{\text{QCL}} \psi\}$. Clearly $\mathcal{I} \cong E$. By Lemma 6.6, $\varphi^1 \notin E$. Moreover, by Proposition 6.7 along with Proposition 6.10 (if $\pi = mm$) or Proposition 6.13 (if $\pi = inc$) we can conclude that $E \in sem(SF_T^\pi \cup SF_\varphi)$. □

The above result allows us to apply fast SAT- or ASP-based argumentation solvers to reason on QCL-theories efficiently. See (Dvořák, Greßler, and Woltran 2018) for a SETAF-specific solver. For a more general overview of argumentation solvers, see (Lagniez et al. 2021). Recently, SETAFs have been investigated with focus on efficient algorithms (Dvořák, König, and Woltran 2021, 2022a,b). While QCL has been encoded in ASP (Bernreiter, Maly, and Woltran 2020), to the best of our knowledge there are no implementations for preferred model entailment.

Regarding computational complexity, deciding whether an argument is contained in all semi-stable extensions (as needed in Theorem 6.14) is Π_2^{P}-complete for SETAFs (Dvořák, Greßler, and Woltran 2018). As we know from Section 5.3, deciding $T \mid\!\sim^{\pi}_{\mathrm{QCL}} \varphi$ is Π_2^{P}-complete for $\pi = \mathit{inc}$ and Θ_2^{P}-complete for $\pi = \mathit{mm}$ (cf. Table 5.1). Thus, when capturing $T \mid\!\sim^{\pi}_{\mathrm{QCL}} \varphi$, there is a complexity gap for $\pi = \mathit{mm}$ but not for $\pi = \mathit{inc}$. Note that all discussed problems are on the second level of the polynomial hierarchy.

6.4 Conclusion

We successfully mapped Qualitative Choice Logic (QCL) theories to argumentation frameworks with collective attacks (SETAFs). The preferred models of the initial QCL-theory directly correspond to the semi-stable extensions of the constructed SETAF, which further allows us to decide preferred model entailment. We considered both the inclusion-based and the minmax preferred model semantics (cf. Section 5.1). Unlike the translation (Sedki 2015) from PQCL-theories to Value-based AFs, our construction is purely syntactic and polynomial in size and runtime.

Our results show that the connection between choice logics and argumentation is closer than previously known. Indeed, we find that SETAFs are well-suited for capturing languages such as QCL, where soft and hard constraints are jointly represented. Moreover, we demonstrated that semi-stable semantics are a useful tool that can handle degree-minimization in a straightforward way.

Observe that every SETAF can be translated into an equivalent Dung-style AF with only polynomial overhead (Polberg 2017). However, this requires the introduction of additional arguments. Thus, the usage of SETAFs allows us to capture QCL-formulas more directly, with each argument ψ^k corresponding to a subformula $\psi \in \mathit{sf}(\varphi)$.

Regarding future work, we plan to find a syntactic and polynomial translation from QCL-theories to SETAFs that respects the lexicographically preferred models of the initial theory (Brewka, Benferhat, and Berre 2004). A difficulty here is that this approach relies on *counting* how many formulas are satisfied to a certain degree. From a complexity theoretical standpoint, however, such a translation must exist, since preferred model entailment under the lexicographic approach is Δ_2^{P}-complete (see Section 5.3) while skeptical acceptance of semi-stable semantics is Π_2^{P}-complete.

Finally, our work can be extended to formalisms related to QCL. This of course includes other choice logics featured in this thesis such as Conjunctive Choice Logic (Boudjelida and Benferhat 2016) or Lexicographic Choice Logic (Bernreiter 2020), both of which replace the ordered disjunction of QCL with alternative choice connectives. Note that our construction is in large parts independent of ordered disjunction, i.e., a similar construction may be possible for other choice logics. A more distantly related system is the recently introduced Lexicographic Logic (Charalambidis et al. 2021) which uses lists of truth values rather than satisfaction degrees.

CHAPTER 7

Conclusion

In this thesis, we investigated the notion of integrated preferences in Knowledge Representation and Reasoning (KR), where information about preferences (soft-constraints) and truth (hard-constraints) is represented and/or resolved together instead of separately. Specifically, we examined two different KR-formalisms which belong to the paradigm of integrated preferences, namely choice logics and abstract argumentation with preferences. In choice logics, non-classical choice connectives such as ordered disjunction enable us to jointly represent hard- and soft-constraints. In argumentation, we made use of four preference reductions from the literature that modify the attack relation based on preferences between arguments, i.e., hard- and soft-constraints are jointly resolved.

By studying the syntactic, semantic, and computational properties of choice logics and abstract argumentation with preferences, we gained new insights into the close interplay between hard- and soft-constraints present in formalisms featuring integrated preferences. We found that the impact of preference reductions in abstract argumentation depends greatly on which reduction is used, while the properties of preferred model entailment in choice logics depend on which preferred model semantics is considered. For instance, the introduction of preferences can lead to desirable semantic properties being lost in some cases (e.g. I-maximality in the case of argumentation and rational monotonicity in the case of choice logics), but these properties can still be guaranteed if a suitable preference reduction or preferred model semantics is selected. Similarly, we found that, in both logic and argumentation, integrated preferences can cause a mild increase in complexity (usually by one level in the polynomial hierarchy) depending on how preferences are handled. Therefore, our results show that the exact method of resolving preferences has to be chosen with care, and they allow for informed decisions when designing systems based on choice logics or abstract argumentation with preferences.

7.1 Summary

Let us summarize our contributions in detail, starting with our results on argumentation (Subsection 7.1.1), followed by our study of choice logics (Subsection 7.1.2), and concluded by the translation from choice logics theories to abstract argumentation (Subsection 7.1.3).

7.1.1 Abstract Argumentation with Preferences

The notion of preference reductions has been studied before when applied to Abstract Argumentation Frameworks (AFs) (Dung 1995), resulting in Preference-based AFs (PAFs) (Amgoud and Cayrol 1998; Kaci et al. 2021). We built upon and extended this study by investigating the effect of the four preference reductions in two settings:

- In Chapter 3 we introduced and studied Conditional PAFs (CPAFs), a formalism capable of representing and reasoning with conditional preferences.
- In Chapter 4 we investigated the impact of preferences in Claim-augmented AFs (CAFs) by introducing and examining Preference-based CAFs (PCAFs).

Regarding the syntactic impact of the four preference reductions, we showed that resolving preferences on the crucial class of well-formed CAFs (wfCAFs) results in four new CAF-classes (one for each preference reduction) that lie inbetween wfCAFs and general CAFs. We characterized these classes and studied their relationship to each other (cf. Figure 4.3). Crucially, this syntactic analysis indicated that the impact of preferences on wfCAFs is not arbitrary, and therefore gave rise to the natural question of whether the new CAF classes retain the beneficial semantic and computational properties of wfCAFs or not.

As for semantic properties, we studied the principle of I-maximality for both CPAFs (cf. Section 3.1) and PCAFs (cf. Figure 4.6). For CPAFs we showed that the global naive and global preferred semantics preserve I-maximality under all four preference reductions. This is similar to PCAFs, where I-maximality is preserved under the hybrid naive and hybrid preferred semantics, where subset-maximization is handled on the claim-level rather than the argument level. On the other hand, for CPAFs, stable semantics as well as the local naive semantics preserve I-maximality under Reduction 2–4, but not under Reduction 1. For the local preferred semantics, I-maximality does not hold for any of the four preference reductions. In contrast, for PCAFs, only Reduction 3 preserves I-maximality for any semantics other than hybrid naive and hybrid preferred, namely for the inherited preferred semantics, both variants of semi-stable semantics, and all variants of stable semantics.

Furthermore, we investigated CPAFs with respect to ten semantic principles for preference-based argumentation laid out by Kaci et al. (2021). We showed that most principles that are satisfied by a semantics for PAFs are also satisfied by the corresponding semantics on CPAFs (cf. Table 3.1). In particular, complete and stable semantics under Reduction 3 satisfy most of the ten principles. There are differences to the case of PAFs (cf. Table 2.2), however, with complete semantics under Reduction 3 not satisfying $P5^*$ (extension growth), or grounded semantics no longer satisfying many principles since CPAFs do not always have a unique grounded extension. Moreover, $P8^*$ (preference-based immunity) is no longer satisfied by any semantics except the global preferred semantics.

For PCAFs, in addition to I-maximality, we also examined the relationship between the various semantics (cf. Figure 4.6). Unlike on wfCAFs, the various variants of stable semantics do not coincide under Reduction 1. Under Reductions 2 and 4 however, they do coincide, with the only exception being $stb\text{-}cf^2_{hyb}$ which does not coincide with $stb\text{-}adm^2_{hyb}$ or stb^2_{inh}. Under Reduction 3, on the other hand, the various variants of stable semantics and the two variants of preferred semantics (inherited and hybrid) coincide.

Regarding computational properties, we studied the complexity of CPAFs (cf. Table 3.2) and PCAFs (cf. Table 4.1) under all four preference reductions. For CPAFs, the complexity of credulous (resp. skeptical) acceptance under local naive semantics using Reduction 1 ($naive^1_{cp}$) is NP-complete (resp. coNP-complete), while the corresponding problem for AFs/PAFs is in P (cf. Table 2.1). Under global naive semantics using Reduction 1, verification is now coNP-complete (instead of in P) which causes skeptical acceptance to become Π^{P}_2-complete. For both variants of naive semantics, the complexity rises under Reduction 1, but not under Reduction 2–4. In contrast, for grounded semantics we see a rise in complexity for *all* preference reductions, with credulous and skeptical acceptance now being coNP-complete instead of in P. A similar rise in complexity by one level in the polynomial hierarchy was observed for complete semantics (skeptical acceptance) and local preferred semantics (credulous acceptance). For PCAFs, the complexity of verification rises for Reduction 1 by one level in the polynomial hierarchy when compared to wfCAFs. Under Reductions 2–4, the complexity remains on the same level as for wfCAFs, except for complete semantics where the problem is NP-complete (instead of in P) under Reductions 2 and 4. Similarly to global naive semantics for CPAFs, under Reduction 1, skeptical acceptance of hybrid naive semantics for PCAFs is Π^{P}_2-complete.

Overall, we can conclude that Reduction 3, the most conservative of the four preference reductions, preserves semantic properties in more cases and features a lower complexity when compared to the other preference reductions. Reductions 2 and 4 constitute a middle ground, preserving some semantic properties and featuring a lower complexity at least on PCAFs. Reduction 1 on the other hand preserves only few semantic properties and leads to a higher complexity in more cases than the other reductions.

Lastly, we compared CPAFs to important related formalisms in detail (cf. Section 3.4). We found that CPAFs can naturally capture Value-based AFs (Atkinson and Bench-Capon 2021), and that CPAFs exhibit crucial differences to Extended AFs (Modgil 2009) when it comes to how preferences are interpreted, with our CPAFs being designed specifically with conditional preferences in mind. Moreover, we discussed the recently introduced lifting-based CPAFs of Alfano et al. (2023), which are similar to our reduction-based CPAFs, but belong to the paradigm of separated preferences rather than integrated preferences, and exhibit significant differences when it comes to semantic and computational properties.

7.1.2 Choice Logics

In Chapter 5 we comprehensively studied preferred model entailment in choice logics with regards to logical, computational, and proof-theoretic properties. We considered large classes of choice logics, and moreover put a special focus on Qualitative Choice Logic (QCL) (Brewka, Benferhat, and Berre 2004), Conjunctive Choice Logic (CCL) (Boudjelida and Benferhat 2016), and Lexicographic Choice Logic (LCL) (Bernreiter 2020). Moreover, we investigated several preferred model semantics, namely the inclusion-based and lexicographic approaches from the literature (Brewka, Benferhat, and Berre 2004) as well as the newly introduced minmax and log-lexicographic approaches.

Regarding logical properties, we first showed that preferred model entailment is non-monotonic for all considered preferred model semantics and all choice logics in which more than two satisfaction degrees are obtainable (cf. Proposition 5.7). We then investigated preferred model entailment with regards to key principles for non-monotonic entailment laid out by Kraus, Lehmann, and Magidor (1990), namely cautious monotonicity, cumulative transitivity, and rational monotonicity (cf. Section 5.2). Our results show that cautious monotonicity and cumulative transitivity are satisfied for all choice logics and all considered preferred model semantics. Rational monotonicity is satisfied for all choice logics and all considered preferred model semantics, except for the inclusion-based approach.

As for computational properties, we found that both the choice logic and the preferred model semantics have an impact on the complexity of preferred model entailment (cf. Table 5.1). Specifically, for QCL and CCL, preferred model entailment is Θ_2^{P}-complete under the minmax semantics, $\Delta_2^{\mathsf{P}}[O(\log^2 n)]$-complete under the log-lexicographic semantics, Δ_2^{P}-complete under the lexicographic semantics, and Π_2^{P}-complete under the inclusion-based semantics. For LCL, however, preferred model entailment is Δ_2^{P}-complete under the minmax, log-lexicographic, and lexicographic semantics, and Π_2^{P}-complete under the inclusion-based semantics. Moreover, we investigated the complexity of checking whether a given interpretation is a preferred model of a choice logic theory, and found that this problem is coNP-complete for all considered preferred model semantics and all choice logics in which more than two satisfaction degrees are obtainable.

Finally, we introduced sequent calculi for preferred model entailment and proved that they are sound and complete (cf. Section 5.4). Specifically, we described how calculi for QCL, CCL, and LCL under the minmax, lexicographic, and inclusion-based semantics can be obtained. Every one of these calculi is in turn based on a labeled monotonic calculus, a labeled refutation calculus, and a rule for preferred model entailment that makes use of the two labeled calculi in its premises. In addition, while the labeled calculi are sound and complete without a cut-rule, we show that *cut* (resp. *cut*2) is admissible in the labeled monotonic calculus (resp. the labeled refutation calculus).

In summary, we found that preferred model entailment satisfies many important properties typically associated with non-monotonic entailment: firstly, crucial logical properties hold for most preferred model semantics; secondly, while we do see a rise in complexity compared to classical entailment, all problems remain on the second level of the polynomial hierarchy and are not harder than entailment in other prominent non-monotonic logics (Eiter and Gottlob 1993; Eiter and Lukasiewicz 2000); lastly, preferred model entailment in choice logics can be decided syntactically using our proof systems.

7.1.3 From Choice Logics to Abstract Argumentation

In Chapter 6 we studied the connection between choice logics and abstract argumentation by translating QCL-theories to Argumentation Frameworks with *collective attacks* (SETAFs) (Nielsen and Parsons 2006). As a first step, we encoded single QCL-formulas into SETAFs and showed that there is a correspondence between the interpretations relevant to the initial formula and the stable (and semi-stable) extensions of the constructed SETAF. We built upon this encoding for single QCL-formulas to translate QCL-theories, both under the minmax and inclusion-based preferred model semantics, to SETAFs. We then showed that the initial QCL-theory is in semantic correspondence with the constructed SETAF, and that we can decide preferred model entailment in QCL via our translation by deciding skeptical acceptance for semi-stable semantics on the constructed SETAF.

Note that the encoding of single QCL-formulas as a stepping stone to an encoding of QCL-theories is similar to the approach we used in our sequent calculus for choice logics, where single formulas are encoded via the labeled calculi upon which the calculi for preferred model entailment are built.

7.2 Future Work

In this section, we highlight some challenges and avenues for future work when it comes to argumentation with preferences, choice logics, and integrated preferences in KR.[15]

In abstract argumentation, a possibility for future work is to apply the notion of preference reductions to other generalizations of standard AFs. For instance, Preference-based SETAFs have been considered before as a tool to capture prioritized knowledge bases (Bienvenu and Bourgaux 2020), but, to the best of our knowledge, they have not been studied in detail yet. A further interesting option to this end are Bipolar AFs (Amgoud et al. 2008) where both attack and support relations are present. Another possibility for future work in argumentation with preferences lies in the connection between abstract and structured formalisms. Concerning conditional preferences, it would be valuable to explore the connection between our CPAFs and the work by Dung, Thang, and Son (2019), in which conditional preferences are considered in structured argumentation. Regarding claim-centric argumentation, we plan to investigate in more detail how PCAFs are related to structured formalisms such as ABA+ (Cyras and Toni 2016) or ASPIC (Modgil and Prakken 2013) where preference reductions are also used.

[15]Note that future work specific to each chapter is contained at the end of the respective chapter. Specifically, we refer to Section 3.5, Section 4.5, Section 5.5, and Section 6.4.

In choice logics, an interesting avenue for future work is to design and implement efficient algorithms for preferred model entailment to facilitate practical applications. Another promising possibility for future work is to further study related logics that are, however, not part of the choice logic framework as defined in (Bernreiter 2020). Two such formalisms are Prioritized QCL (PQCL) and QCL+ (Benferhat and Sedki 2008), although they both can be captured by the choice logic framework as fragments (Bernreiter, Maly, and Woltran 2021). We also mention the alternative semantics for the language of QCL described by Maly and Woltran (2018), the game-theoretic approach to ordered disjunction of Freiman and Bernreiter (2023a,b), and the lexicographic logic of Charalambidis et al. (2021).

Regarding the connection between choice logics and abstract argumentation, an immediate open question is how QCL-theories under the lexicographic preferred model semantics can be translated to SETAFs or other argumentation formalisms. Moreover, translations for CCL and LCL have not been defined yet, although adapting the existing translation for QCL should be straightforward.

From a broader perspective, there are also possibilities for future work when it comes to integrated preferences in KR in general. Indeed, the term integrated preferences has not been established before and was coined in this thesis. One avenue for future work in this context is to explicitly investigate the different subparadigms of integrated preferences we have seen: in choice logics, hard- and soft-constraints are jointly represented, whereas in argumentation using preference reductions hard- and soft-constraints are only jointly resolved. It would be valuable to investigate the relationship between these different kinds of integrated preferences more deeply by, for example, finding a connection between choice logics and argumentation that makes use of preference reductions in a natural way. Recall that for our translation from QCL-theories to SETAFs we used no preferences on the argumentation side. Indeed, it is not clear to us how to do so in this case in a meaningful way. More general possibilities for future work are to identify other existing formalisms that feature integrated preferences, to define entirely new systems that make use of integrated preferences in a natural way, and to methodically compare formalisms using integrated preferences with formalisms using separated preferences.

Lastly, there are also possibilities for future work when it comes to applying formalisms with integrated preferences to other fields of study. For instance, the fact that our CPAFs are capable of capturing Value-based AFs suggests that they can be used for normative reasoning, and a deeper investigation of this connection may prove fruitful. Moreover, formal argumentation can be used to provide explanations in various settings (Cyras et al. 2021; Vassiliades, Bassiliades, and Patkos 2021), and studying how choice logics and abstract argumentation with preferences can be used as tools for explainable AI may provide interesting results.

List of Figures

List of Tables

Bibliography

Aaronson, S. (2017). „P=?NP". In: *Electron. Colloquium Comput. Complex.* TR17-004.

Alfano, G., Greco, S., Parisi, F., and Trubitsyna, I. (July 2022). „On Preferences and Priority Rules in Abstract Argumentation". In: *Proc. IJCAI'22*. IJCAI Organization, pp. 2517–2524.

Alfano, G., Greco, S., Parisi, F., Trubitsyna, I., et al. (2023). „Abstract Argumentation Framework with Conditional Preferences". In: *Proc. AAAI'23 (to appear)*.

Amgoud, L. and Cayrol, C. (1998). „On the Acceptability of Arguments in Preference-based Argumentation". In: *Proc. UAI'98*. Morgan Kaufmann, pp. 1–7.

— (2002). „A Reasoning Model Based on the Production of Acceptable Arguments". In: *Ann. Math. Artif. Intell.* 34.1-3, pp. 197–215.

Amgoud, L., Cayrol, C., Lagasquie-Schiex, M., and Livet, P. (2008). „On bipolarity in argumentation frameworks". In: *Int. J. Intell. Syst.* 23.10, pp. 1062–1093.

Amgoud, L. and Vesic, S. (2014). „Rich preference-based argumentation frameworks". In: *Int. J. Approx. Reason.* 55.2, pp. 585–606.

Arora, S. and Barak, B. (2009). *Computational Complexity: A Modern Approach*. Cambridge University Press.

Atkinson, K. and Bench-Capon, T. J. M. (2021). „Value-based Argumentation". In: *Handbook of Formal Argumentation, Volume 2*. Also appears in IfCoLog Journal of Logics and their Applications 8(6):1543–1588. College Publications, pp. 299–354.

Avron, A. (1996). „The method of hypersequents in the proof theory of propositional nonclassical logics". In: *Logic: from foundations to applications*. Ed. by W. Hodges, M. Hyland, C. Steinhorn, and J. Truss. Oxford: Oxford Science Publications, pp. 1–32.

Baader, F., Horrocks, I., Lutz, C., and Sattler, U. (2017). *An Introduction to Description Logic*. Cambridge University Press.

Baaz, M., Lahav, O., and Zamansky, A. (2013). „Finite-valued Semantics for Canonical Labelled Calculi". In: *J. Autom. Reason.* 51.4, pp. 401–430.

Baroni, P., Caminada, M., and Giacomin, M. (2018). „Abstract Argumentation Frameworks and Their Semantics". In: *Handbook of Formal Argumentation*. College Publications. Chap. 4, pp. 159–236.

Baroni, P., Cerutti, F., Giacomin, M., and Guida, G. (2009). „Encompassing Attacks to Attacks in Abstract Argumentation Frameworks". In: *Proc. ECSQARU'09*. Vol. 5590. LNCS. Springer, pp. 83–94.

Baroni, P., Cerutti, F., Giacomin, M., and Guida, G. (2011). „AFRA: Argumentation framework with recursive attacks“. In: *Int. J. Approx. Reason.* 52.1, pp. 19–37.

Baroni, P., Gabbay, D., Giacomin, M., and van der Torre, L. (2018). *Handbook of formal argumentation*. College Publications.

Baroni, P. and Giacomin, M. (2007). „On principle-based evaluation of extension-based argumentation semantics“. In: *Artif. Intell.* 171.10-15, pp. 675–700.

Baroni, P., Governatori, G., and Riveret, R. (2016). „On Labelling Statements in Multi-Labelling Argumentation“. In: *Proc. ECAI'16*. Vol. 285. Frontiers in Artificial Intelligence and Applications. IOS Press, pp. 489–497.

Baroni, P. and Riveret, R. (2019). „Enhancing Statement Evaluation in Argumentation via Multi-labelling Systems“. In: *J. Artif. Intell. Res.* 66, pp. 793–860.

Baumeister, D., Järvisalo, M., Neugebauer, D., Niskanen, A., and Rothe, J. (2021). „Acceptance in incomplete argumentation frameworks“. In: *Artif. Intell.* 295, p. 103470.

Bench-Capon, T. J. M., Doutre, S., and Dunne, P. E. (2007). „Audiences in argumentation frameworks“. In: *Artif. Intell.* 171.1, pp. 42–71.

Benferhat, S. and Sedki, K. (2008). „Two alternatives for handling preferences in qualitative choice logic“. In: *Fuzzy Sets Syst.* 159.15, pp. 1889–1912.

— (2010). „An alert correlation approach based on security operator's knowledge and preferences“. In: *J. Appl. Non Class. Logics* 20.1-2, pp. 7–37.

Bernreiter, M. (2020). „A General Framework for Choice Logics“. Master's thesis. Vienna, Austria: TU Wien.

Bernreiter, M., Dvořák, W., Rapberger, A., and Woltran, S. (2022a). „The Effect of Preferences in Abstract Argumentation Under a Claim-Centric View“. In: *Proc. NMR'22*. Vol. 3197. CEUR Workshop Proceedings. CEUR-WS.org, pp. 27–38.

— (2023). „The Effect of Preferences in Abstract Argumentation under a Claim-Centric View“. In: *Proc. AAAI'23*. AAAI Press, pp. 6253–6261.

Bernreiter, M., Dvořák, W., and Woltran, S. (2022). „Abstract Argumentation with Conditional Preferences“. In: *Proc. COMMA'22*. Vol. 353. Frontiers in Artificial Intelligence and Applications. IOS Press, pp. 92–103.

— (2023). „Abstract Argumentation with Conditional Preferences“. In: *Argument & Computation* pre-press, pp. 1–29.

Bernreiter, M. and König, M. (2023). „From Qualitative Choice Logic to Abstract Argumentation“. In: *Proc. KR'23*, pp. 737–741.

Bernreiter, M., Lolic, A., Maly, J., and Woltran, S. (2022b). „Sequent Calculi for Choice Logics“. In: *Proc. IJCAR'22*. Springer, pp. 331–349.

— (2024a). „Sequent Calculi for Choice Logics“. In: *J. Autom. Reason. (to appear)*.

Bernreiter, M., Maly, J., Nardi, O., and Woltran, S. (2024b). „Combining Voting and Abstract Argumentation to Understand Online Discussions“. In: *Proc. AAMAS'24 (to appear)*.

Bernreiter, M., Maly, J., and Woltran, S. (2020). „Encoding Choice Logics in ASP“. In: *Proc. ASPOCP@ICLP'20*. Vol. 2678. CEUR Workshop Proceedings. CEUR-WS.org.

— (2021). „Choice Logics and Their Computational Properties“. In: *Proc. IJCAI'21*. IJCAI Organization, pp. 1794–1800.

— (2022). „Choice logics and Their Computational Properties“. In: *Artif. Intell.* 311, p. 103755.

Bhattarai, B., Granmo, O.-C., and Jiao, L. (2023). „An Interpretable Knowledge Representation Framework for Natural Language Processing with Cross-Domain Application“. In: *European Conference on Information Retrieval*. Springer, pp. 167–181.

Bienvenu, M. and Bourgaux, C. (2020). „Querying and Repairing Inconsistent Prioritized Knowledge Bases: Complexity Analysis and Links with Abstract Argumentation“. In: *Proc. KR'20*, pp. 141–151.

Bienvenu, M., Lang, J., and Wilson, N. (2010). „From Preference Logics to Preference Languages, and Back“. In: *Proc. KR'10*. AAAI Press, pp. 414–424.

Biere, A., Heule, M., Maaren, H. van, and Walsh, T., eds. (2021). *Handbook of Satisfiability - Second Edition*. Vol. 336. Frontiers in Artificial Intelligence and Applications. IOS Press.

Bistarelli, S. and Santini, F. (2021). „Weighted Argumentation“. In: *Handbook of Formal Argumentation, Volume 2*. Also appears in IfCoLog Journal of Logics and their Applications 8(6):1589–1622. College Publications, pp. 355–395.

Bobadilla, J., Ortega, F., Hernando, A., and Gutiérrez, A. (2013). „Recommender systems survey“. In: *Knowl. Based Syst.* 46, pp. 109–132.

Bochman, A. (2018). „Argumentation, nonmonotonic reasoning and logic“. In: *Handbook of Formal Argumentation*. Also appears in IfCoLog Journal of Logics and their Applications 4(8):2889–2926. College Publications. Chap. 19, pp. 981–1016.

Bonatti, P. A. and Olivetti, N. (2002). „Sequent calculi for propositional nonmonotonic logics“. In: *ACM Trans. Comput. Log.* 3.2, pp. 226–278.

Bonatti, P. A. (1993). „A Gentzen system for non-theorems“. In: *Technisce Universität Wien, Institut für Informationssysteme, Technical Report CD-TR* 93, p. 52.

Boudjelida, A. and Benferhat, S. (2016). „Conjunctive Choice Logic“. In: *Proc. ISAIM'16*.

Boutilier, C., Brafman, R. I., Domshlak, C., Hoos, H. H., and Poole, D. (2004). „CP-nets: A Tool for Representing and Reasoning with Conditional Ceteris Paribus Preference Statements“. In: *J. Artif. Intell. Res.* 21, pp. 135–191.

Brandt, F., Conitzer, V., Endriss, U., Lang, J., and Procaccia, A. D. (2016). *Handbook of computational social choice*. Cambridge University Press.

Brewka, G., Benferhat, S., and Berre, D. L. (2004). „Qualitative choice logic“. In: *Artif. Intell.* 157.1-2, pp. 203–237.

Brewka, G., Delgrande, J. P., Romero, J., and Schaub, T. (2015). „asprin: Customizing Answer Set Preferences without a Headache“. In: *Proc. AAAI'15*. AAAI Press, pp. 1467–1474.

Brewka, G., Niemelä, I., and Syrjänen, T. (2004). „Logic Programs with Ordered Disjunction“. In: *Comput. Intell.* 20.2, pp. 335–357.

Brewka, G., Niemelä, I., and Truszczynski, M. (2003). „Answer set optimization“. In: *Proc. IJCAI'03*. Vol. 3. Citeseer, pp. 867–872.

Brewka, G., Niemelä, I., and Truszczynski, M. (2008). „Preferences and Nonmonotonic Reasoning“. In: *AI Mag.* 29.4, pp. 69–78.

Brewka, G., Truszczynski, M., and Woltran, S. (2010). „Representing Preferences Among Sets“. In: *Proc. AAAI’10*. AAAI Press, pp. 273–278.
Caminada, M. (2018). „Argumentation semantics as formal discussion“. In: *Handbook of Formal Argumentation*. College Publications. Chap. 10, pp. 487–518.
Carnielli, W. A. (1987). „Systematization of Finite Many-Valued Logics Through the Method of Tableaux“. In: *J. Symb. Log.* 52.2, pp. 473–493.
Castro, J. and Seara, C. (1992). „Characterizations of Some Complexity Classes Between Θ_2^P and Δ_2^P“. In: *Proc. STACS’92*. Vol. 577. LNCS. Springer, pp. 305–317.
Charalambidis, A., Nomikos, C., and Rondogiannis, P. (2022). „Strong Equivalence of Logic Programs with Ordered Disjunction: A Logical Perspective“. In: *Theory Pract. Log. Program.* 22.5, pp. 708–722.
Charalambidis, A., Papadimitriou, G., Rondogiannis, P., and Troumpoukis, A. (Nov. 2021). „A Many-valued Logic for Lexicographic Preference Representation“. In: *Proc. KR’21*, pp. 646–650.
Charalambidis, A., Rondogiannis, P., and Troumpoukis, A. (2021). „A Logical Characterization of the Preferred Models of Logic Programs with Ordered Disjunction“. In: *Theory Pract. Log. Program.* 21.5, pp. 629–645.
Creignou, N., Pichler, R., and Woltran, S. (2018). „Do Hard SAT-Related Reasoning Tasks Become Easier in the Krom Fragment?“ In: *Log. Methods Comput. Sci.* 14.4.
Cyras, K., Rago, A., Albini, E., Baroni, P., and Toni, F. (2021). „Argumentative XAI: A Survey“. In: *Proc. IJCAI’21*. ijcai.org, pp. 4392–4399.
Cyras, K. and Toni, F. (2016). „ABA+: Assumption-Based Argumentation with Preferences“. In: *Proc. KR’16*. AAAI Press, pp. 553–556.
Delgrande, J. P., Schaub, T., and Tompits, H. (2003). „A Framework for Compiling Preferences in Logic Programs“. In: *Theory Pract. Log. Program.* 3.2, pp. 129–187.
Domshlak, C., Hüllermeier, E., Kaci, S., and Prade, H. (2011). „Preferences in AI: An overview“. In: *Artif. Intell.* 175.7-8, pp. 1037–1052.
Dung, P. M. (1995). „On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-Person Games“. In: *Artif. Intell.* 77.2, pp. 321–358.
Dung, P. M., Kowalski, R. A., and Toni, F. (2009). „Assumption-Based Argumentation“. In: *Argumentation in Artificial Intelligence*. Springer, pp. 199–218.
Dung, P. M., Thang, P. M., and Son, T. C. (2019). „On Structured Argumentation with Conditional Preferences“. In: *Proc. AAAI’19*. AAAI Press, pp. 2792–2800.
Dunne, P. E., Dvořák, W., Linsbichler, T., and Woltran, S. (2015). „Characteristics of multiple viewpoints in abstract argumentation“. In: *Artif. Intell.* 228, pp. 153–178.
Dunne, P. E., Modgil, S., and Bench-Capon, T. J. M. (2010). „Computation in Extended Argumentation Frameworks“. In: *Proc. ECAI’10*. Vol. 215. FAIA. IOS Press, pp. 119–124.
Dvořák, W. and Dunne, P. E. (2018). „Computational Problems in Formal Argumentation and their Complexity“. In: *Handbook of Formal Argumentation*. Also appears in IfCoLog Journal of Logics and their Applications 4(8):2557–2622. College Publications. Chap. 14, pp. 631–687.

Dvořák, W., Fandinno, J., and Woltran, S. (2019). „On the expressive power of collective attacks“. In: *Argument Comput.* 10.2, pp. 191–230.

Dvořák, W., Greßler, A., Rapberger, A., and Woltran, S. (2023). „The complexity landscape of claim-augmented argumentation frameworks“. In: *Artif. Intell.* 317, p. 103873.

Dvořák, W., Greßler, A., and Woltran, S. (2018). „Evaluating SETAFs via Answer-Set Programming“. In: *Proc. SAFA@COMMA'18*. Vol. 2171. CEUR Workshop Proceedings. CEUR-WS.org, pp. 10–21.

Dvořák, W., König, M., Ulbricht, M., and Woltran, S. (2024). „Principles and their Computational Consequences for Argumentation Frameworks with Collective Attacks“. In: *J. Artif. Intell. Res.* 79, pp. 69–136.

Dvořák, W., König, M., and Woltran, S. (2021). „On the Complexity of Preferred Semantics in Argumentation Frameworks with Bounded Cycle Length“. In: *Proc. KR'21*, pp. 671–675.

— (2022a). „Deletion-Backdoors for Argumentation Frameworks with Collective Attacks“. In: *Proc. SAFA@COMMA'22*. Vol. 3236. CEUR Workshop Proceedings. CEUR-WS.org, pp. 98–110.

— (2022b). „Treewidth for Argumentation Frameworks with Collective Attacks“. In: *Proc. COMMA'22*, pp. 140–151.

Dvořák, W., Rapberger, A., and Woltran, S. (2023). „A claim-centric perspective on abstract argumentation semantics: Claim-defeat, principles, and expressiveness“. In: *Artif. Intell.* 324, p. 104011.

Dvořák, W. and Woltran, S. (2020). „Complexity of abstract argumentation under a claim-centric view“. In: *Artif. Intell.* 285, p. 103290.

Eiter, T. and Gottlob, G. (1993). „Propositional Circumscription and Extended Closed-World Reasoning are Π_2^P-Complete“. In: *Theor. Comput. Sci.* 114.2, pp. 231–245.

Eiter, T. and Lukasiewicz, T. (2000). „Default reasoning from conditional knowledge bases: Complexity and tractable cases“. In: *Artificial Intelligence* 124.2, pp. 169–241.

Faber, W., Truszczyński, M., and Woltran, S. (2013). „Strong equivalence of qualitative optimization problems“. In: *Journal of Artificial Intelligence Research* 47, pp. 351–391.

Falappa, M. A., García, A. J., Kern-Isberner, G., and Simari, G. R. (2011). „On the evolving relation between belief revision and argumentation“. In: *The Knowledge Engineering Review* 26.1, pp. 35–43.

Fazzinga, B., Flesca, S., and Furfaro, F. (2020). „Revisiting the Notion of Extension over Incomplete Abstract Argumentation Frameworks.“ In: *Proc. IJCAI'20*, pp. 1712–1718.

Flasiński, M. (2016). *Introduction to artificial intelligence*. Springer.

Flouris, G. and Bikakis, A. (2019). „A comprehensive study of argumentation frameworks with sets of attacking arguments“. In: *Int. J. Approx. Reason.* 109, pp. 55–86.

Freiman, R. and Bernreiter, M. (2023a). „Truth and Preferences - A Game Approach for Qualitative Choice Logic“. In: *Proc. JELIA'23*. Vol. 14281. LNCS. Springer, pp. 547–560.

— (2023b). „Validity in Choice Logics - A Game-Theoretic Investigation“. In: *Proc. WoLLIC'23*. Vol. 13923. LNCS. Springer, pp. 211–226.

Geibinger, T. and Tompits, H. (2020). „Sequent-Type Calculi for Systems of Nonmonotonic Paraconsistent Logics“. In: *ICLP Technical Communications*. Vol. 325. EPTCS, pp. 178–191.

Gençay, E., Schüller, P., and Erdem, E. (2019). „Applications of non-monotonic reasoning to automotive product configuration using answer set programming“. In: *Journal of Intelligent Manufacturing* 30.3, pp. 1407–1422.

Governatori, G. and Rotolo, A. (2006). „Logic of violations: a Gentzen system for reasoning with contrary-to-duty obligations“. In: *The Australasian Journal of Logic* 4, pp. 193–215.

Hausman, D. M. (2011). *Preference, value, choice, and welfare*. Cambridge University Press.

Heyninck, J., Kern-Isberner, G., Meyer, T. A., Haldimann, J. P., and Beierle, C. (2023). „Conditional Syntax Splitting for Non-monotonic Inference Operators“. In: *Proc. AAAI'23*. AAAI Press, pp. 6416–6424.

Hodel, R. E. (2013). *An introduction to mathematical logic*. Courier Corporation.

Horty, J. F. (2002). „Skepticism and floating conclusions“. In: *Artif. Intell.* 135.1-2, pp. 55–72.

Kaci, S., van der Torre, L. W. N., Vesic, S., and Villata, S. (2021). „Preference in Abstract Argumentation“. In: *Handbook of Formal Argumentation, Volume 2*. College Publications, pp. 211–248.

Kaci, S., van der Torre, L. W. N., and Villata, S. (2018). „Preference in Abstract Argumentation“. In: *Proc. COMMA'18*. Vol. 305. FAIA. IOS Press, pp. 405–412.

Kaminski, M. and Francez, N. (2021). „Calculi for Many-Valued Logics“. In: *Logica Universalis* 15.2, pp. 193–226.

Kern-Isberner, G., Beierle, C., and Brewka, G. (2020). „Syntax Splitting = Relevance + Independence: New Postulates for Nonmonotonic Reasoning From Conditional Belief Bases“. In: *Proc. KR'20*, pp. 560–571.

König, M. (2020). „Graph-classes of argumentation frameworks with collective attacks“. Master's thesis. Vienna, Austria: TU Wien.

König, M., Rapberger, A., and Ulbricht, M. (2022). „Just a Matter of Perspective – Intertranslating Expressive Argumentation Formalisms“. In: *Proc. COMMA'22*, pp. 212–223.

Kraus, S., Lehmann, D., and Magidor, M. (1990). „Nonmonotonic Reasoning, Preferential Models and Cumulative Logics“. In: *Artif. Intell.* 44.1-2, pp. 167–207.

Krentel, M. W. (1988). „The Complexity of Optimization Problems“. In: *J. Comput. Syst. Sci.* 36.3, pp. 490–509.

Lagniez, J., Lonca, E., Mailly, J., and Rossit, J. (2021). „Design and results of ICCMA 2021“. In: *arXiv preprint arXiv:2109.08884*.

Lehmann, D. and Magidor, M. (1992). „What does a Conditional Knowledge Base Entail?“ In: *Artif. Intell.* 55.1, pp. 1–60.

Liétard, L., Hadjali, A., and Rocacher, D. (2014). „Towards a Gradual QCL Model for Database Querying“. In: *IPMU (3)*. Vol. 444. Communications in Computer and Information Science. Springer, pp. 130–139.

Lifschitz, V. (2019). *Answer set programming.* Springer Heidelberg.

Liu, F. (2010). „Von Wright's *"The Logic of Preference"* revisited". In: *Synth.* 175.1, pp. 69–88.

Łukasiewicz, J. (1951). *Aristotle's Syllogistic from the Standpoint of Modern Formal Logic.* Oxford: Clarendon Press.

Maly, J. and Woltran, S. (2018). „A New Logic for Jointly Representing Hard and Soft Constraints". In: *PRUV@IJCAR.* Vol. 2157. CEUR Workshop Proceedings. CEUR-WS.org.

McCarthy, J. (1980). „Circumscription - A Form of Non-Monotonic Reasoning". In: *Artif. Intell.* 13.1-2, pp. 27–39.

Modgil, S. (2009). „Reasoning about Preferences in Argumentation Frameworks". In: *Artif. Intell.* 173.9-10, pp. 901–934.

Modgil, S. and Bench-Capon, T. J. M. (2011). „Metalevel argumentation". In: *J. Log. Comput.* 21.6, pp. 959–1003.

Modgil, S. and Prakken, H. (2010). „Reasoning about Preferences in Structured Extended Argumentation Frameworks". In: *Proc. COMMA'10.* Vol. 216. FAIA. IOS Press, pp. 347–358.

— (2013). „A general account of argumentation with preferences". In: *Artif. Intell.* 195, pp. 361–397.

— (2014). „The $ASPIC^+$ framework for structured argumentation: a tutorial". In: *Argument Comput.* 5.1, pp. 31–62.

— (2018). „Abstract Rule-Based Argumentation". In: *Handbook of Formal Argumentation.* Also appears in IfCoLog Journal of Logics and their Applications 4(8):2319–2406. College Publications. Chap. 6, pp. 287–364.

Moore, R. C. (1985). „Semantical Considerations on Nonmonotonic Logic". In: *Artif. Intell.* 25.1, pp. 75–94.

Nielsen, S. H. and Parsons, S. (2006). „A Generalization of Dung's Abstract Framework for Argumentation: Arguing with Sets of Attacking Arguments". In: *ArgMAS, Revised Selected and Invited Papers.* Vol. 4766. LNCS. Springer, pp. 54–73.

Papadimitriou, C. H. and Yannakakis, M. (1982). „The complexity of facets (and some facets of complexity)". In: *Proceedings of the fourteenth annual ACM symposium on Theory of computing*, pp. 255–260.

Parikh, R. (1999). „Beliefs, belief revision, and splitting languages". In: *Logic, language and computation* 2.96, pp. 266–268.

Paulius, D. and Sun, Y. (2019). „A survey of knowledge representation in service robotics". In: *Robotics and Autonomous Systems* 118, pp. 13–30.

Pigozzi, G., Tsoukiàs, A., and Viappiani, P. (2016). „Preferences in artificial intelligence". In: *Ann. Math. Artif. Intell.* 77.3-4, pp. 361–401.

Polberg, S. (2017). „Developing the Abstract Dialectical Framework". PhD thesis. TU Wien.

Prakken, H. and Vreeswijk, G. A. (2002). „Logics for defeasible argumentation". In: *Handbook of Philosophical Logic.* Vol. 4. Springer.

Reiter, R. (1980). „A Logic for Default Reasoning". In: *Artif. Intell.* 13.1-2, pp. 81–132.

Rocha, V. H. N. and Cozman, F. G. (2022a). „A Credal Least Undefined Stable Semantics for Probabilistic Logic Programs and Probabilistic Argumentation“. In: *Proc. KR'22*. IJCAI Organization, pp. 309–319.

— (2022b). „Bipolar Argumentation Frameworks with Explicit Conclusions: Connecting Argumentation and Logic Programming“. In: *Proc. NMR'22@FLoC2022*. Vol. 3197. CEUR Workshop Proceedings. CEUR-WS.org, pp. 49–60.

Russell, S. and Norvig, P. (2021). *Artificial Intelligence: A Modern Approach (4th Edition)*. Pearson.

Sakama, C. and Inoue, K. (2000). „Prioritized logic programming and its application to commonsense reasoning“. In: *Artif. Intell.* 123.1-2, pp. 185–222.

Schnoebelen, P. (2003). „Oracle Circuits for Branching-Time Model Checking“. In: *Proc. ICALP'03*. Vol. 2719. LNCS. Springer, pp. 790–801.

Schumm, G. F. (1987). „Transitivity, preference and indifference“. In: *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 52.3, pp. 435–437.

Sedki, K. (2015). „Value-based argumentation framework built from prioritized qualitative choice logic“. In: *Int. J. Approx. Reason.* 64, pp. 75–94.

Sedki, K., Lamy, J., and Tsopra, R. (2020). „Learning Preferences in Prioritized Qualitative Choice Logic“. In: *Proc. ICTAI'20*. IEEE, pp. 368–375.

— (2022). „Qualitative Choice Logic For Modeling Experts Recommendations of Antibiotics“. In: *Proc. FLAIRS'22*.

Segoufin, L. and ten Cate, B. (2013). „Unary negation“. In: *Log. Methods Comput. Sci.* 9.3.

Sheth, A., Roy, K., and Gaur, M. (2023). „Neurosymbolic Artificial Intelligence (Why, What, and How)“. In: *IEEE Intelligent Systems* 38.3, pp. 56–62.

Shoham, Y. (1987). „Nonmonotonic Logics: Meaning and Utility“. In: *Proc. IJCAI'87*. Morgan Kaufmann, pp. 388–393.

Skiba, K. and Thimm, M. (2022). „Ordinal Conditional Functions for Abstract Argumentation“. In: *Proc. COMMA'22*. Vol. 353. Frontiers in Artificial Intelligence and Applications. IOS Press, pp. 308–319.

Słupecki, J., Bryll, G., and Wybraniec-Skardowska, U. (1971). „Theory of rejected propositions. I“. In: *Studia Logica: An International Journal for Symbolic Logic* 29, pp. 75–123.

Tiomkin, M. L. (1988). „Proving unprovability“. In: *Proc. LICS'88*. IEEE Computer Society, pp. 22–26.

Troelstra, A. S. and Schwichtenberg, H. (2000). *Basic Proof Theory*. Cambridge: Cambridge University Press.

van der Torre, L. and Vesic, S. (2018). „The Principle-Based Approach to Abstract Argumentation Semantics“. In: *Handbook of Formal Argumentation*. Also appears in IfCoLog Journal of Logics and their Applications 4(8):2737–2780. College Publications. Chap. 16, pp. 797–838.

Vassiliades, A., Bassiliades, N., and Patkos, T. (2021). „Argumentation and explainable artificial intelligence: a survey“. In: *Knowl. Eng. Rev.* 36, e5.

von Wright, G. H. (1963). *The Logic of Preference*. Edinburgh University Press.
Wagner, K. W. (1990). „Bounded Query Classes". In: *SIAM J. Comput.* 19.5, pp. 833–846.
Wakaki, T. (2020). „Consistency in Assumption-Based Argumentation". In: *Proc. COMMA'20*. Vol. 326. FAIA. IOS Press, pp. 371–382.
Wyner, A. Z., Bench-Capon, T. J. M., and Dunne, P. E. (2013). „On the Instantiation of Knowledge Bases in Abstract Argumentation Frameworks". In: *Proc. CLIMA'14*. Vol. 8143. LNCS. Springer, pp. 34–50.

Appendix

Additional Proofs for Section 4.2

Lemma A.1. *Let $\mathcal{F} = (A, R, cl)$ be a CAF. $\mathcal{F} \in \mathcal{R}_2$-**CAF** iff there are no arguments a, a', b, b' in $\mathcal{F}$ with $cl(a) = cl(a')$ and $cl(b) = cl(b')$ such that $(a, b) \in wfp(\mathcal{F})$, $(b, a) \notin R$, $(a', b) \in R$, and either $(b, a') \in R$ or $((a', b') \notin R$ and $(b', a') \notin R)$.*

Proof. " $\Longrightarrow$ ": By contrapositive. Suppose that there are $a, a', b, b' \in A$ with $cl(a') = cl(a)$ and $cl(b') = cl(b)$ such that $(a, b) \in wfp(\mathcal{F})$, $(b, a) \notin R$, $(a', b) \in R$, and either $(b, a') \in R$ or $((a', b') \notin R$ and $(b', a') \notin R)$. Towards a contradiction, assume that $\mathcal{F} \in \mathcal{R}_2$-**CAF**. Then there must be a PCAF $\mathcal{P} = (A, R', cl, \succ)$ such that $\mathcal{R}_2(\mathcal{P}) = \mathcal{F}$. Reduction 2 cannot completely remove conflicts between arguments. Since there is no conflict between a and b insemantic-properties $\mathcal{F}$ there can be no conflict in $\mathcal{P}$ either, i.e., $(a, b) \notin R'$ and $(b, a) \notin R'$. Therefore, since the underlying CAF (A, R', cl) of $\mathcal{P}$ must be well-formed, $(a', b) \notin R'$. Since $(a', b) \in R$ it must be that $(b, a') \in R'$ and $a' \succ b$. Then $(b, a') \notin \mathcal{R}_2(\mathcal{P})$. Furthermore, by the well-formedness of (A, R', cl), we have that $(b', a') \in R'$ and therefore either $(a', b') \in \mathcal{R}_2(\mathcal{P})$ or $(b', a') \in \mathcal{R}_2(\mathcal{P})$. Contradiction to $\mathcal{R}_2(\mathcal{P}) = \mathcal{F}$.

" $\Longleftarrow$ ": Our underlying assumption is that there are no arguments a, a', b, b' in $\mathcal{F}$ with $cl(a) = cl(a')$ and $cl(b) = cl(b')$ such that $(a, b) \in wfp(\mathcal{F})$, $(b, a) \notin R$, $(a', b) \in R$, and either $(b, a') \in R$ or $((a', b') \notin R$ and $(b', a') \notin R)$. We will construct a PCAF $\mathcal{P} = (A, R'', cl, \succ)$ such that $\mathcal{R}_2(\mathcal{P}) = \mathcal{F}$.

But first, as an intermediate step, we construct the CAF $\mathcal{F}' = (A, R', cl)$. We say that (b, a) is forced in $\mathcal{F}$ if $(a, b) \in R$ and if there is an argument a' with $cl(a') = cl(a)$ such that $(a', b) \notin R$ and $(b, a') \notin R$. Observe that if (b, a) is forced in $\mathcal{F}$, then (a, b) cannot be forced in $\mathcal{F}$ by our underlying assumption. Furthermore, if (b, a) is forced in $\mathcal{F}$, then $(b, a) \notin R$, again by our underlying assumption. We construct $R' = (R \cup \{(b, a) \mid (b, a) \text{ is forced in } \mathcal{F}\}) \setminus \{(a, b) \mid (b, a) \text{ is forced in } \mathcal{F}\}$. Note that $(a, b) \in wfp(\mathcal{F}')$ implies $(b, a) \in R'$ for all arguments a, b: towards a contradiction, assume otherwise. Then there is some $(a, b) \in wfp(\mathcal{F}')$ such that $(b, a) \notin R'$. Then $(a, b) \notin R$ and $(b, a) \notin R$ by construction of R'. Furthermore, since $(a, b) \in wfp(\mathcal{F}')$, there must be some a' with $cl(a') = cl(a)$ and $(a', b) \in R'$. It cannot be that $(a', b) \in R$, otherwise (b, a') would be forced in $\mathcal{F}$ and $(a', b) \notin R'$. Thus, $(b, a') \in R$ and (a', b)

was added to R' because it is forced in $\mathcal{F}$. But this is only possible if there is some b' with $cl(b') = cl(b)$ and $(a', b') \notin R$ and $(b', a') \notin R$. This contradicts our underlying assumption: $(b', a') \in wfp(\mathcal{F})$, $(a', b') \notin R$, $(b, a') \in R$, $(a, b) \notin R$, and $(b, a) \notin R$.

Now we construct $R'' = R' \cup \{(a, b) \mid (a, b) \in wfp(\mathcal{F}')\}$. Furthermore, $b \succ a \iff (a, b) \in R'' \setminus R$. This gives us $\mathcal{P} = (A, R'', cl, \succ)$. The underlying CAF of $\mathcal{P}$ is well-formed since $wfp((A, R'', cl)) = \emptyset$ by construction. Moreover, $\succ$ is asymmetric since if $(a, b) \in R''$ and $(b, a) \in R''$ then, by construction of R' and R'', either $(a, b) \in R$ or $(b, a) \in R$. Lastly, we show that $\mathcal{R}_2(\mathcal{P}) = \mathcal{F}$: if $(a, b) \in R'' \setminus R$, then we defined $b \succ a$ and thus $(a, b) \notin \mathcal{R}_2(\mathcal{P})$. If $(a, b) \in R \setminus R''$, then (b, a) was forced in $\mathcal{F}$, i.e., $(b, a) \notin R$ but $(b, a) \in R'$ and therefore also $(b, a) \in R''$. Thus, we define $a \succ b$ which means that $(a, b) \in \mathcal{R}_2(\mathcal{P})$. □

Lemma A.2. *Let $\mathcal{F} = (A, R, cl)$ be a CAF. $\mathcal{F} \in \mathcal{R}_3$-**CAF** iff $(a, b) \in wfp(\mathcal{F})$ implies $(b, a) \in R$.*

Proof. "$\Longrightarrow$": By contrapositive. Suppose there is $(a, b) \in wfp(\mathcal{F})$ such that $(b, a) \notin R$. Towards a contradiction, assume $\mathcal{F} \in \mathcal{R}_3$-**CAF**. Then there is a PCAF $\mathcal{P} = (A, R', cl, \succ)$ such that $\mathcal{R}_3(\mathcal{P}) = \mathcal{F}$. Since Reduction 3 can only delete but not introduce attacks, and since (A, R', cl) must be well-formed, $(a, b) \in R'$. However, Reduction 3 cannot completely remove conflicts between arguments, i.e., either $(a, b) \in \mathcal{R}_3(\mathcal{P})$ or $(b, a) \in \mathcal{R}_3(\mathcal{P})$. Contradiction.

"$\Longleftarrow$": Suppose $(a, b) \in wfp(\mathcal{F})$ implies $(b, a) \in R$. Then $\mathcal{R}_3(\mathcal{P}) = \mathcal{F}$ for the PCAF $\mathcal{P} = (A, R', cl, \succ)$ with $R' = R \cup \{(a, b) \mid (a, b) \in wfp(\mathcal{F})\}$ as well as $a \succ b \iff (b, a) \in R' \setminus R$. (A, R', cl) is well-formed since $wfp((A, R', cl)) = \emptyset$. Furthermore, $\succ$ is asymmetric by construction. □

Lemma A.3. *Let $\mathcal{F} = (A, R, cl)$ be a CAF. $\mathcal{F} \in \mathcal{R}_4$-**CAF** iff there are no arguments a, a', b, b' in $\mathcal{F}$ with $cl(a) = cl(a')$ and $cl(b) = cl(b')$ such that $(a, b) \in wfp(\mathcal{F})$, $(b, a) \notin R$, $(a', b) \in R$, and either $(b, a') \notin R$ or $((a', b') \notin R$ and $(b', a') \notin R)$.*

Proof. Similar to the proof of Lemma A.1:

"$\Longrightarrow$": By contrapositive. Suppose that there are $a, a', b, b' \in A$ with $cl(a') = cl(a)$ and $cl(b') = cl(b)$ such that $(a, b) \in wfp(\mathcal{F})$, $(b, a) \notin R$, $(a', b) \in R$, and either $(b, a') \notin R$ or $((a', b') \notin R$ and $(b', a') \notin R)$. Towards a contradiction, assume that $\mathcal{F} \in \mathcal{R}_4$-**CAF**. Then there must be a PCAF $\mathcal{P} = (A, R', cl, \succ)$ such that $\mathcal{R}_4(\mathcal{P}) = \mathcal{F}$. Reduction 4 cannot completely remove conflicts between arguments. Since there is no conflict between a and b in $\mathcal{F}$ there can be no conflict in $\mathcal{P}$ either, i.e., $(a, b) \notin R'$ and $(b, a) \notin R'$. Therefore, since the underlying CAF of $\mathcal{P}$ must be well-formed, $(a', b) \notin R'$. The only way to obtain $(a', b) \in R$ from $(a', b) \notin R'$ via Reduction 4 is to have $(b, a') \in R'$ and $a' \succ b$. Then $(b, a') \in \mathcal{R}_4(\mathcal{P})$. Furthermore, by the well-formedness of (A, R', cl), we have that $(b', a') \in R'$ and therefore either $(a', b') \in \mathcal{R}_4(\mathcal{P})$ or $(b', a') \in \mathcal{R}_4(\mathcal{P})$. Contradiction to $\mathcal{R}_4(\mathcal{P}) = \mathcal{F}$.

“ $\Longleftarrow$ ”: Our underlying assumption is that there are no arguments a, a', b, b' in $\mathcal{F}$ with $cl(a) = cl(a')$ and $cl(b) = cl(b')$ such that $(a, b) \in wfp(\mathcal{F})$, $(b, a) \notin R$, $(a', b) \in R$, and either $(b, a') \notin R$ or $((a', b') \notin R$ and $(b', a') \notin R)$. We will construct a PCAF $\mathcal{P} = (A, R'', cl, \succ)$ such that $\mathcal{R}_4(\mathcal{P}) = \mathcal{F}$.

But first, as an intermediate step, we construct the CAF $\mathcal{F}' = (A, R', cl)$. We say that (b, a) is forced in $\mathcal{F}$ if $(a, b) \in R$, $(b, a) \in R$, and if there is an argument a' with $cl(a') = cl(a)$ such that $(a', b) \notin R$ and $(b, a') \notin R$. Observe that if (b, a) is forced in $\mathcal{F}$, then (a, b) cannot be forced in $\mathcal{F}$ by our underlying assumption. We construct $R' = R \setminus \{(a, b) \mid (b, a) \text{ is forced in } \mathcal{F}\}$. Note that $(a, b) \in wfp(\mathcal{F}')$ implies $(b, a) \in R'$ for all arguments a, b: towards a contradiction, assume otherwise. Then there is some $(a, b) \in wfp(\mathcal{F}')$ such that $(b, a) \notin R'$. Then $(a, b) \notin R$ and $(b, a) \notin R$ by construction of R'. Furthermore, there must be some a' with $cl(a') = cl(a)$ and $(a', b) \in R'$. It cannot be that $(a', b) \in R$ and $(b, a') \in R$, otherwise (b, a') would be forced in $\mathcal{F}$ and $(a', b) \notin R'$. Thus, $(a', b) \in R$ and $(b, a') \notin R$ by construction of $\mathcal{F}'$. But this contradicts our underlying assumption: $(a, b) \in wfp(\mathcal{F})$, $(b, a) \notin R$, $(a', b) \in R$, and $(b, a') \notin R$.

Now we construct $R'' = R' \cup \{(a, b) \mid (a, b) \in wfp(\mathcal{F}')\}$. Furthermore, $b \succ a \iff (a, b) \in R'' \setminus R$ or $(b, a) \in R \setminus R''$. This gives us, $\mathcal{P} = (A, R'', cl, \succ)$. The underlying CAF of $\mathcal{P}$ is well-formed since $wfp((A, R'', cl)) = \emptyset$ by construction. Moreover, $\succ$ is asymmetric: if $b \succ a$, there are two cases.

1. $(a, b) \in R'' \setminus R$. Clearly, $(a, b) \notin R \setminus R''$. Moreover, $(a, b) \in R'' \setminus R$ implies $(b, a) \in R$ since we did not add attacks to R'' if there was no conflict between these attacks in R. Thus, $(b, a) \notin R'' \setminus R$. We can conclude $a \not\succ b$.

2. $(b, a) \in R \setminus R''$. Clearly, $(b, a) \notin R'' \setminus R$. Moreover, $(b, a) \in R \setminus R''$ implies $(a, b) \in R''$, since we never completely removed conflicts when constructing R'' from R. Thus, $(a, b) \notin R \setminus R''$. We can conclude $a \not\succ b$.

Lastly, we show that $\mathcal{R}_4(\mathcal{P}) = \mathcal{F}$: if $(a, b) \in R'' \setminus R$, then we defined $b \succ a$. As above, $(a, b) \in R'' \setminus R$ implies $(b, a) \in R$. The only possible reason for why we added (a, b) to R'' is because $(a, b) \in wfp(\mathcal{F}')$. As previously discussed, this means that $(b, a) \in R'$ and therefore also $(b, a) \in R''$. Thus, $(a, b) \notin \mathcal{R}_4(\mathcal{P})$. If $(a, b) \in R \setminus R''$, then $a \succ b$. As above, this implies $(b, a) \in R''$, and therefore $(a, b) \in \mathcal{R}_4(\mathcal{P})$. □

Additional Proofs for Section 4.4

Lemma A.4. *$Ver^{PCAF}_{\sigma^i_\mu}$ is Σ^{P}_2-hard for $\sigma^i_\mu \in \{stg^1_{inh}, stg^1_{hyb}\}$, even if we restrict ourselves to PCAFs with transitive preference relations.*

Proof. We provide a reduction from $\text{QBF}^2_\forall$ to the complementary problem. Let $\Phi = \forall Y \exists Z \varphi$ be an instance of $\text{QBF}^2_\forall$, where φ is given by a set Ω of clauses over atoms $X = Y \cup Z$. We construct the CAF $\mathcal{F} = (A, Att, cl)$ with underlying AF $F = (A, R)$ and a set of claims C:

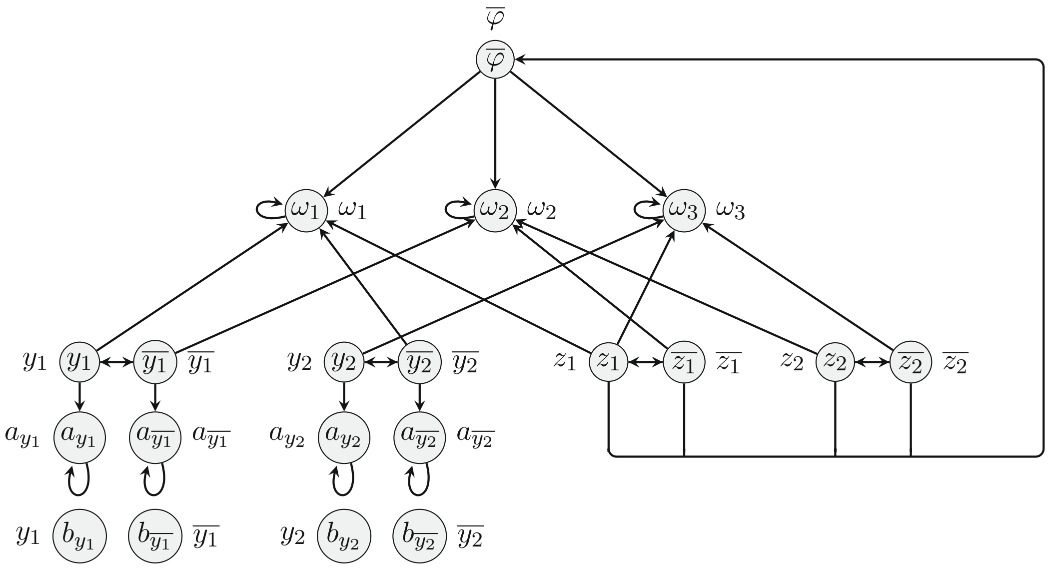

Figure A.1: Reduction of the $\text{QBF}_\forall^2$ instance $\Phi = \forall y_1, y_2 \exists z_1, z_2 \varphi$ with φ given by clauses $\omega_1 = \{y_1, \neg y_2, z_1\}$, $\omega_2 = \{\neg y_1, \neg z_1, z_2\}$, $\omega_3 = \{y_2, z_1, \neg z_2\}$ to an instance of $Ver^{PCAF}_{stg^1_{inh}}$.

- $A = \{\overline{\varphi}\} \cup \Omega \cup X \cup \overline{X} \cup Y_a \cup \overline{Y}_a \cup Y_b \cup \overline{Y}_b$, where $\overline{X} = \{\overline{x} \mid x \in X\}$, $Y_a = \{a_y \mid y \in Y\}$, $\overline{Y}_a = \{a_{\overline{y}} \mid y \in Y\}$, $Y_b = \{b_y \mid y \in Y\}$, $\overline{Y}_b = \{b_{\overline{y}} \mid y \in Y\}$;
- $Att = \{(x, \overline{x}), (\overline{x}, x) \mid x \in X\} \cup \{(\omega, \omega), (\overline{\varphi}, \omega) \mid \omega \in \Omega\} \cup$
  $\{(x, \omega) \mid x \in \omega, \omega \in \Omega\} \cup \{(\overline{x}, \omega) \mid \neg x \in \omega, \omega \in \Omega\} \cup$
  $\{(a_v, a_v), (v, a_v) \mid v \in Y \cup \overline{Y}\} \cup \{(z, \overline{\varphi}), (\overline{z}, \overline{\varphi}) \mid z \in Z\}$;
- $cl(b_v) = v$ for $b_v \in Y_b \cup \overline{Y}_b$ and $cl(v) = v$ else;
- $C = Y \cup \overline{Y} \cup \{\overline{\varphi}\}$.

Figure A.1 illustrates the above construction. Note that $\mathcal{F} \in \mathcal{R}_1$-$\mathbf{CAF}_{tr}$ since all paths in $wfp(\mathcal{F}) = \{(b_v, v) \mid v \in Y \cup \overline{Y}\}$ are of length 1 (only arguments in $Y_b \cup \overline{Y}_b$ have outgoing edges in $wfp(\mathcal{F})$). It remains to verify the correctness of the reduction, i.e., we will show that Φ is valid iff $C \notin \sigma_\mu(\mathcal{F})$. The proof proceeds similar as the proof of Proposition 4.33.

"$\implies$": Assume Φ is valid. Consider any $S \subseteq A$ such that $S \in cf(F)$ and $cl(S) = C$. Then $S \subseteq Y \cup \overline{Y} \cup Y_b \cup \overline{Y}_b \cup \{\overline{\varphi}\}$. Let $Y' = S \cap Y$. Since Φ is valid, there is $Z' \subseteq Z$ such that $M = Y' \cup Z'$ is a model of φ. Let $T = M \cup \{\overline{x} \mid x \in X \setminus M\} \cup Y_b \cup \overline{Y}_b$. Note that $T \in cf(F)$ by construction. Moreover, $S \setminus \{\overline{\varphi}\} \subseteq T$. Since for each $z \in Z$ we have either $z \in T$ or $\overline{z} \in T$, and since $(z, \overline{\varphi}), (\overline{z}, \overline{\varphi}) \in R$, we have $\overline{\varphi} \in T_F^+$ (resp. $\overline{\varphi} \in T_F^*$). Since $M \models \varphi$, all clause-arguments $\omega \in \Omega$ are attacked by T and we have $\{\overline{\varphi}\}_F^+ = \Omega \subseteq T_F^+$ (resp. $\{\overline{\varphi}\}_\mathcal{F}^* = \Omega \subseteq T_F^*$). We can conclude that $S \cup S_F^+ \subset T \cup T_F^+$ (resp. $cl(S) \cup S_\mathcal{F}^* \subset cl(T) \cup T_\mathcal{F}^*$), i.e., $C \notin stg_{inh}(\mathcal{F})$ (resp. $C \notin stg_{hyb}(\mathcal{F})$).

“ $\Longleftarrow$ ”: Assume $C \notin stg_{inh}(\mathcal{F})$ (resp. $C \notin stg_{hyb}(\mathcal{F})$) and consider an arbitrary subset $Y' \subseteq Y$. We must show that there is $Z' \subseteq Z$ such that $Y' \cup Z' \models \varphi$. Let $S = Y' \cup \{\overline{y} \mid y \in Y \setminus Y'\} \cup Y^* \cup \overline{Y}^* \cup \{\overline{\varphi}\}$. Observe that $cl(S) = C$ and that $S \in cf(F)$. By $C \notin stg_{inh}(\mathcal{F})$ (resp. $C \notin stg_{hyb}(\mathcal{F})$) there is some $T \in cf(F)$ with $S \cup S_F^+ \subset T \cup T_F^+$ (resp. $cl(S) \cup S^*_{\mathcal{F}} \subset cl(T) \cup T^*_{\mathcal{F}}$).

In particular, we have $Y' \cup \{\overline{y} \mid y \in Y \setminus Y'\} \subseteq T$ since each $a_v \in S_F^+$ (resp. $a_v \in S^*_{\mathcal{F}}$) with $v \in Y \cup \overline{Y}$ has precisely one non-self-attacking attacker (namely the argument v). Moreover, we can assume that T contains each argument $v \in Y_b \cup \overline{Y}_b$ since each such v is unattacked and does not attack any other argument. Thus, $T \supseteq S \setminus \{\overline{\varphi}\}$.

Furthermore, $\overline{\varphi} \notin T$ since $S \in naive(F)$ (note that $\overline{\varphi}$ attacks each $z, \overline{z}$ with $z \in Z$ as well as every clause-argument $\omega \in \Omega$ and thus cannot be extended any further). Therefore, it must be that $\overline{\varphi} \in T_F^+$ (resp. $\overline{\varphi} \in T^*_{\mathcal{F}}$). Also, we have that T attacks each clause-argument $\omega \in \Omega$ since $\Omega \subseteq S_F^+$ (resp. $\Omega \subseteq S^*_{\mathcal{F}}$), and since each clause-argument $\omega \in \Omega$ is self-attacking.

Now, let $Z' = Z \cap T$. We show that $M = Y' \cup Z'$ is a model of φ. Consider some arbitrary clause $\omega \in \Omega$. Then there is some argument $v \in T$ such that $(v, \omega) \in Att$. As outlined above, $v \neq \overline{\varphi}$ since $\overline{\varphi}$ is not contained in T. Consequently, we have $v \in X \cup \overline{X}$. In case $v \in X$ we have $v \in M \cap \omega$, in case $v \in \overline{X}$ we have $\neg v \in \omega$ and $v \notin M$ by definition of Att. In every case, the clause ω is satisfied by M. As ω was chosen arbitrary it follows that $M \models \varphi$. We can conclude that Φ is valid. □

Lemma A.5. *$Ver^{PCAF}_{\sigma^i_\mu}$ is* DP*-hard for $\sigma^i_\mu = naive^1_{hyb}$, even if we restrict ourselves to PCAFs with transitive preference relations.*

Proof. Before showing DP-hardness, we show NP- and coNP-hardness separately:

Let $(\mathcal{P}, C)$ be an instance of $Ver^{PCAF}_{cf^1_{inh}}$, i.e., $\mathcal{P} = (A, R, cl, \succ)$ is a PCAF and $C \subseteq cl(A)$ is the claim-set to be verified for conflict-freeness. Recall that $Ver^{PCAF}_{cf^1_{inh}}$ is NP-complete, even when restricted to transitive preferences (see Proposition 4.31).

- First, we construct a PCAF $\mathcal{P}' = (A', R', cl', \succ')$ with $A' = \{x \mid x \in A, cl(x) \in C\}$ as well as $R' = \{(x, y) \mid x, y \in A', (x, y) \in R\}$, $cl'(x) = cl(x)$ for all $x \in A'$, and $x \succ' y$ iff $x \succ y$ and $x, y \in A'$. Observe that (A', R', cl') is still well-formed. Furthermore, if $\succ$ is transitive, then so is $\succ'$. It is easy to see that $C \in cf_{inh}(\mathcal{R}_1(\mathcal{P}))$ iff $C \in cf_{inh}(\mathcal{R}_1(\mathcal{P}'))$. Since $C = cl(A')$, $C \in cf_{inh}(\mathcal{R}_1(\mathcal{P}))$ iff $C \in naive_{hyb}(\mathcal{R}_1(\mathcal{P}'))$.
- Second, we construct another PCAF $\mathcal{P}'' = (A'', R'', cl'', \succ'')$. Without loss of generality, we can assume $C \neq \emptyset$. We fix an arbitrary claim $c \in C$ and for each claim $d \in C \setminus \{c\}$ introduce a fresh argument z_d. Let Z be the set of those fresh arguments. Then $A'' = A' \cup Z$, $R'' = R' \cup \{(x, z_d) \mid cl(x) = c, z_d \in Z\} \cup \{(z_d, y) \mid z_d \in Z, y \in A'$, there exists $x \in A'$ with $cl(x) = d$ such that $(x, y) \in R'\}$, $cl''(x) = cl'(x)$ for all $x \in A'$, $cl''(z_d) = d$ for all $z_d \in Z$, and $x \succ'' y$ iff $x \succ' y$. (A'', R'', cl'') is

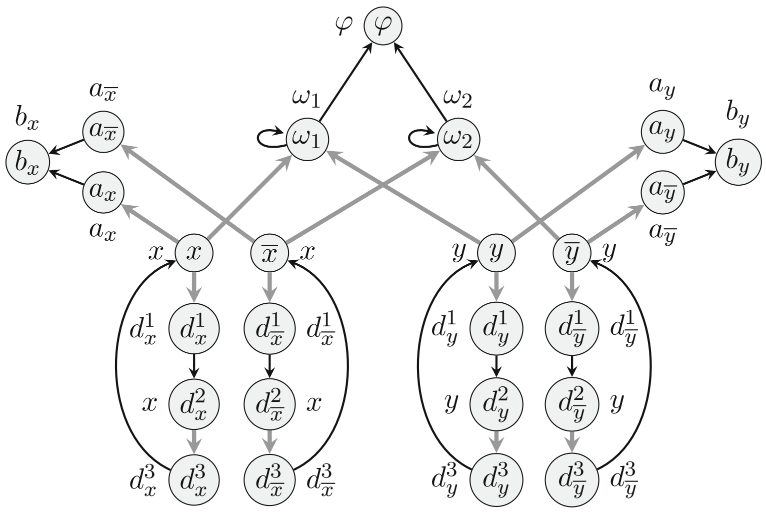

Figure A.2: $\mathcal{R}_2(\mathcal{P})$ from the proof of Lemma A.6, with φ given by clauses $\omega_1 = \{x, y\}$, $\omega_2 = \{\neg x, \neg y\}$. Gray/thick attacks have been reversed by Reduction 2.

well-formed by construction, and $\succ''$ can still be assumed to be transitive. Now we show that $C \notin cf_{inh}(\mathcal{R}_1(\mathcal{P}))$ iff $C \setminus \{c\} \in naive_{hyb}(\mathcal{R}_1(\mathcal{P}''))$: (1) assume $C \in cf_{inh}(\mathcal{R}_1(\mathcal{P}))$. Then also $C \in cf_{inh}(\mathcal{R}_1(\mathcal{P}''))$ and thus $C \setminus \{c\} \notin naive_{hyb}(\mathcal{R}_1(\mathcal{P}''))$. (2) assume $C \notin cf_{inh}(\mathcal{R}_1(\mathcal{P}))$. Then also $C \notin cf_{inh}(\mathcal{R}_1(\mathcal{P}''))$ since all arguments x with $cl(x) = c$ are in conflict with the fresh arguments z_d. But because the fresh arguments z_d do not attack each other, $C \setminus \{c\} \in cf_{inh}(\mathcal{R}_1(\mathcal{P}''))$. Since $C = cl(A'')$, $C \setminus \{c\} \in naive_{hyb}(\mathcal{R}_1(\mathcal{P}''))$.

The construction of $\mathcal{P}'$ shows NP-hardness, and the construction of $\mathcal{P}''$ shows coNP-hardness. Now we show DP-hardness: let (φ_1, φ_2) be an arbitrary instance of SAT-UNSAT, with φ_1 and φ_2 sharing no variables. We can construct instances $(\mathcal{P}_1 = (A_1, R_1, cl_1, \succ_1), C_1)$ and $(\mathcal{P}_2 = (A_2, R_2, cl_2, \succ_2), C_2)$ of $Ver^{PCAF}_{naive^1_{hyb}}$ such that φ_1 is satisfiable iff $C_1 \in naive_{hyb}(\mathcal{R}_1(\mathcal{P}_1))$ and φ_2 is unsatisfiable iff $C_2 \in naive_{hyb}(\mathcal{R}_1(\mathcal{P}_2))$. Note that we can assume $\mathcal{P}_1$ and $\mathcal{P}_2$ to be disjoint, i.e., they share no arguments and claims. Let $\mathcal{P} = (A_1 \cup A_2, R_1 \cup R_2, cl_1 \cup cl_2, \succ_1 \cup \succ_2)$ be the combination of $\mathcal{P}_1$ and $\mathcal{P}_2$. Observe that $C_1 \cup C_2 \in naive_{hyb}(\mathcal{R}_1(\mathcal{P}))$ iff $C_1 \in naive_{hyb}(\mathcal{R}_1(\mathcal{P}_1))$ and $C_2 \in naive_{hyb}(\mathcal{R}_1(\mathcal{P}_2))$. Thus, (φ_1, φ_2) is a yes-instance of SAT-UNSAT iff $C_1 \cup C_2 \in naive_{hyb}(\mathcal{R}_1(\mathcal{P}))$. □

Lemma A.6. *$Ver^{PCAF}_{\sigma^i_\mu}$ is* NP*-hard for $\sigma^i_\mu = com^2_{inh}$, even if we restrict ourselves to PCAFs with transitive preference relations.*

Proof. Let φ be an arbitrary instance of 3-SAT given as a set Ω of clauses over variables X and let $\overline{X} = \{\overline{x} \mid x \in X\}$. We construct a PCAF $\mathcal{P} = (A, Att, cl, \succ)$ as well as a set of claims C:

- $A = \{\varphi\} \cup \Omega \cup X \cup \overline{X} \cup \{a_x \mid x \in X \cup \overline{X}\} \cup \{b_x \mid x \in X\} \cup \{d_x^j \mid x \in X \cup \overline{X}, 1 \leq j \leq 3\}$;
- $Att = \{(\omega, \varphi) \mid \omega \in \Omega\} \cup \{(\omega, \omega) \mid \omega \in \Omega\} \cup$
  $\{(\omega, x) \mid x \in \omega, \omega \in \Omega\} \cup \{(\omega, \overline{x}) \mid \neg x \in \omega, \omega \in \Omega\} \cup$
  $\{(d_x^1, x), (d_x^1, d_x^2), (d_x^3, d_x^2), (d_x^3, x), (a_x, x) \mid x \in X \cup \overline{X}\} \cup$
  $\{(a_x, b_x), (a_{\overline{x}}, b_x) \mid x \in X\}$;
- $cl(x) = cl(\overline{x}) = cl(d_x^2) = cl(d_{\overline{x}}^2) = x$ for $x \in X$, $cl(v) = v$ else;
- $x \succ \omega$, $x \succ d_x^1$, $x \succ a_x$, $d_x^2 \succ d_x^3$ for all $x \in X \cup \overline{X}$ and all $\omega \in \Omega$;
- $C = X \cup \{\varphi\}$.

Figure A.2 illustrates the above construction. It remains to show that φ is satisfiable if and only if $C \in com_{inh}(\mathcal{R}_2(\mathcal{P}))$.

Assume φ is satisfiable. Then there is an interpretation I such that $I \models \varphi$. Let $S = \{x, d_x^2 \mid x \in X, x \in I\} \cup \{\overline{x}, d_{\overline{x}}^2 \mid x \in X, x \notin I\} \cup \{\varphi\}$. Clearly, $cl(S) = C$. Furthermore, S defends φ in $\mathcal{R}_2(\mathcal{P})$ since each clause is satisfied by I, and thus each clause argument ω_j is attacked by some x (or $\overline{x}$) in S. For each variable x, if $x \in S$, then x defends d_x^2 and d_x^2 defends x. Moreover, if $x \in S$, then $\overline{x} \notin S$ and none of $\overline{x}$, $a_{\overline{x}}$, b_x, or $d_{\overline{x}}^j$ with $1 \leq j \leq 3$ is defended by S. Analogously for the case that $\overline{x} \in S$. Thus, S is admissible, and contains all arguments it defends, i.e., $S \in com(\mathcal{R}_2(\mathcal{P}))$.

Assume $C \in com_{inh}(\mathcal{R}_2(\mathcal{P}))$. Then there is $S \subseteq A$ such that $cl(S) = C$ and $S \in com(\mathcal{R}_2(\mathcal{P}))$. For each $x \in X$, at least one of $x, \overline{x}, d_x^2, d_{\overline{x}}^2$ must be contained in S. In fact, if $x \in S$, then also $d_x^2 \in S$ and vice versa. Analogous for $\overline{x}$ and $d_{\overline{x}}^2$. However, it cannot be that $x \in S$ and $\overline{x} \in S$, otherwise b_x would be defended by S and we would have $cl(S) \neq C$. Thus, for each $x \in X$, there is either $x \in S$ or $\overline{x} \in S$, but not both. Furthermore, S defends φ, i.e., S attacks all clause arguments ω_j. Therefore, $I \models \varphi$ for $I = X \cap S$. □