Reasoning About Legal Concepts with Propositional Dependence Logic

MSc Thesis *(Afstudeerscriptie)*

written by

Lorenz Hornung

under the supervision of **Dr. Maria Aloni**, **Dr. Aleksi Anttila** and **Dr. Giovanni Sileno**, and submitted to the Examinations Board in partial fulfillment of the requirements for the degree of

MSc in Logic

at the *Universiteit van Amsterdam.*

Date of the public defense:
July 9, 2025

Members of the Thesis Committee:
Dr. Balder Ten Cate (chair)
Prof. Dr. Davide Grossi
Dr. Aybüke Özgün
Dr. Maria Aloni
Dr. Aleksi Anttila
Dr. Giovanni Sileno

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

Abstract

This thesis investigates reasoning about the applicability of legal concepts, a core element in the application of law, using a variant of propositional dependence logic. The contribution of this work is two-fold.

The first contribution is the development of a new framework. This framework is based on a team-based logic called *propositional dependence logic with the might-operator* $\mathbf{PL}(=(\cdot), \blacklozenge)$. With this new framework, we demonstrate how $\mathbf{PL}(=(\cdot), \blacklozenge)$ can be applied as a formal model of legal reasoning, thereby highlighting the utility of team-based logics in a domain to which they have not previously been applied.

The second contribution concerns the framework's capabilities, addressing several tasks related to the applicability of legal concepts. The framework can (i) express legal information regarding the applicability of such concepts. It introduces a new mechanism capable of (ii) determining whether a concept is applicable to a case. This thesis examines this mechanism and explores how much information is needed to determine the applicability of a concept. Further, the framework enables (iii) reasoning about the stability of predictions. That is, it allows us to analyze whether acquiring more information about a case would lead to a different judgment regarding the applicability of the concept. Additionally, the framework provides the means to (iv) investigate which properties determine the applicability of a concept. Lastly, the framework is capable of (v) predicting the applicability of a concept to a sequence of cases.

Acknowledgements

I would like to express my sincere gratitude to my supervisors — Maria Aloni, Aleksi Anttila, and Giovanni Sileno — whose guidance and support were essential to the development of this thesis. Their commitment to the project and the many insightful discussions we shared have been a steady source of encouragement and inspiration.

Maria, thank you for sharing your passion for how formal systems can illuminate complex phenomena, for your thoughtful and continuous guidance in shaping the formalism developed here, and for the many conceptual conversations that helped clarify the foundation of this work.

I am also deeply appreciative of Aleksi, whose teaching introduced me to dependence logic and whose insightful engagement with the conceptual and logical dimensions of my ideas greatly shaped my thinking. Your eye for detail has been invaluable throughout.

Finally, I am grateful to Giovanni for encouraging me to look beyond the technical aspects of the developed formalism and to consider its broader context and conceptual underpinnings. Your perspective helped me frame my work more clearly within a larger narrative.

To those who fill my days with warmth and joy: thank you for reminding me of the beauty found in ordinary moments. Your support always helps me see what truly matters.

Contents

1 Introduction and Background **2**

1.1 Legal Theory on Legal Concept Applicability 3

1.1.1 Legal Concepts and Meaning Postulates 3

1.1.2 Legal Concepts in Legal Norms 5

1.1.3 Legal Concepts in Legal Reasoning 6

1.2 Towards Modeling Legal Concept Applicability 7

1.3 Tasks of Modeling Legal Concept Applicability 8

1.4 Approaches to Modeling Legal Concept Applicability 9

2 An Intuition for Modeling Legal Concept Applicability **15**

2.1 Legal Concepts, Cases, Conditions and Legal Knowledge 15

2.2 Concept Applicability Functions 16

3 A Logic for Legal Concept Applicability **21**

3.1 The Logics $\mathbf{PL}(=(\cdot))$ and $\mathbf{PL}(=(\cdot), \blacklozenge)$ 21

3.2 $\mathbf{PL}(=(\cdot), \blacklozenge)$ for Legal Concept Applicability 25

4 Derivation of Legal Concept Applicability **33**

4.1 An Intuition about the Derivation of Legal Concept Applicability 33

4.2 Direct Derivation of Legal Concept Applicability 35

4.3 Heuristic Derivation of Legal Concept Applicability 36

4.4 Criteria of Heuristic Derivation of Legal Concept Applicability 41

5 Tasks in the Formal Analysis of Legal Concept Applicability **47**

5.1 Stability of Derived Legal Concept Applicability 47

5.2 Sequential Derivations of Legal Concept Applicability 51

5.3 Relevant Conditions of Legal Concept Applicability 55

6 Conclusion and Future Work **60**

Bibliography **64**

A Appendix **71**

A.1 Proofs of Chapter 3 71

A.2 Proofs of Chapter 4 73

A.3 Proofs of Chapter 5 78

Chapter 1

Introduction and Background

Anyone who travels by train in the Netherlands will have seen people with bicycles on the platforms and on the trains. Depending on the bike, a special ticket needs to be purchased in order to travel with the bike. Let us call these bikes "ticketed bikes", a concept that refers to bikes allowed on trains and requiring a ticket. If a train conductor has to decide whether a passenger needs to possess a bike-ticket, they have to decide whether the passenger's bike qualifies as a "ticketed bike".

In this work, we want to investigate legal reasoning of this nature. To be precise, we want to investigate legal reasoning regarding legal concepts—such as "ticketed bike", "theft", "person" and "accident at work"—and what determines the applicability of such concepts.

The investigation of such reasoning is the subject of extensive research at the interface between law and AI.[1] This research has led to different formal frameworks, such as the reason model by Horty [Hor11] and the ANGELIC methodology by Al-Abdulkarim et al. [AAB16]. The present thesis proposes a novel framework for modeling legal reasoning about concept applicability using a variant of propositional dependence logic.

Propositional dependence logic is a team-based logic. Team-based logics are used in many fields. For instance, Hannula and Kontinen [HK16] demonstrate their use in database theory, Ciardelli et al. [CGR18] and Degano [Deg24] illustrate their significance for formal semantics, and Pacuit and Yang [PY16] show their applicability to the study of computational social choice. This thesis thus extends the possible applications of these logics by applying them to a new field: the formal analysis of legal reasoning.

The thesis is structured as follows. In the remainder of this chapter, we will outline the relevant legal theory on legal concepts (Section 1.1), motivate the simplifications needed to model the applicability of legal concepts (Section 1.2), and characterize the tasks related to modeling the applicability of legal concepts (Section 1.3). In Section 1.4, we will provide an overview of relevant approaches to modeling the applicability of legal concepts and indicate how our framework differs from these frameworks.

In Chapter 2, we will formally introduce the foundational concepts used to model the applicability of legal concepts. We will present the notions of *concepts*, *conditions*, *cases*, *associated knowledge of cases*, and *legal possibilities* in Section 2.1. In Section 2.2, we introduce *concept applicability functions*, which characterize our understanding of how the applicability of legal concepts is specified.

[1] For an overview of the research on law and AI, see Prakken and Sartor [PS15], Bench-Capon [Ben22], and Villata et al. [Vil+22].

Our main contributions are presented in Chapters 3 to 5. In Chapter 3, we demonstrate how propositional dependence logic with the might-operator $\mathbf{PL}(=(\cdot), \blacklozenge)$, a variant of propositional dependence logic $\mathbf{PL}(=(\cdot))$, can be used to model the applicability of legal concepts. To this end, we introduce $\mathbf{PL}(=(\cdot), \blacklozenge)$ in Section 3.1. In Section 3.2, we introduce *legal teams*, which represent concept applicability functions as models of $\mathbf{PL}(=(\cdot), \blacklozenge)$. We further characterize the formulas of $\mathbf{PL}(=(\cdot), \blacklozenge)$ that a team must satisfy in order to be a legal team.

In Chapter 4, we investigate how the applicability of a legal concept to a legal case can be derived from a legal team. A first intuition about the derivation of legal concept applicability is given in Section 4.1. To derive the applicability of legal concepts, we differentiate between two cases. First, the legal team possesses explicit information about the applicability of the concept to the case, allowing the applicability to be directly derived. This is studied in Section 4.2. Second, the legal team lacks explicit information about the concept's applicability to the case, requiring a heuristic derivation instead. Heuristic derivations aim to predict the applicability of legal concepts based on properties satisfied by the legal team. We formalize this approach in Section 4.3. Furthermore, we characterize the conditions necessary for a legal team to support such predictions in Section 4.4.

We examine additional challenges associated with the derivation of the applicability of legal concepts in Chapter 5. Given a prediction about a concept's applicability to a case, one might ask whether additional information about the case might alter the prediction. We address this question in Section 5.1. Additionally, in Section 5.2, we examine sequential predictions of a concept's applicability across a sequence of cases. In Section 5.3, we analyze how to determine the conditions that directly affect the applicability of legal concepts.

Lastly, in Chapter 6, we discuss possible future work.

1.1 Legal Theory on Legal Concept Applicability

To justify the formal study of reasoning about the applicability of legal concepts and to emphasize the importance of legal concepts, we will now outline the functions of legal concepts in the legal domain.

1.1.1 Legal Concepts and Meaning Postulates

As discussed by Frändberg [Frä87] and Searle [Sea11], legal concepts are used to express the content of legal norms, and they function as classifications. This is illustrated by the following example.

Example 1.1. Let us investigate the concept of "ticketed bikes" that refer to bikes that are allowed to be brought onto trains and need a ticket. If we look at the regulations of Nederlandse Spoorwegen, the Dutch railway company, we will notice that the conditions under which an object is considered a "ticketed bike" are not straightforward. Besides many other rules, Nederlandse Spoorwegen articulates for instance that a folded folding bike is not a "ticketed bike", while a folding bike that is not folded is a "ticketed bike". The terms and conditions contain statements as the following:

(1) If a folding bike is not folded, it is a "ticketed bike".

(2) If a bike is a "ticketed bike", one ought to possess a bike-ticket to travel with it on trains.

Statement (1) expresses what is classified as a "ticketed bike". Statement (2) states the legal consequences of the concept "ticketed bike". △

Following Pigozzi and van der Torre [PT17], statements like (1) are an example of what we will call a *meaning postulate* of "ticketed bike". We understand *meaning postulates* of a legal concept as statements that specify what falls under that concept. In other words, a meaning postulate defines or clarifies the conditions for the concept's applicability or non-applicability. (We will elaborate on this below.)

Furthermore, Example 1.1 indicates that the concept "ticketed bike" instantiates a link between statements (1) and (2), connecting the factual descriptions given in (1) with the consequences illustrated in (2) [Ros57].[2] This exemplifies that legal concepts bridge the gap between legal consequences and factual descriptions given by their meaning postulates [Lie87; HPT07; BT12].

Linking factual descriptions to legal consequences might require considering multiple legal concepts. Rather than being isolated, legal concepts are organized hierarchically within legal systems [PS15]. Thus, meaning postulates of one legal concept might refer to other legal concepts and factual descriptions. As a result, the connection between legal consequences and factual descriptions might not be established by a single legal concept, but by multiple concepts and their associated meaning postulates.

Example 1.2. Consider the following characterization of the concept "theft", discussed by Pigozzi and van der Torre [PT17]:

(3) An act is *theft* if the person exercising the act has taken a movable object from the possession of another person into their own possession without the consent of the other person or any other legal authorization.

According to the authors, this meaning postulate provides a definition of the concept "theft" by referring to more foundational legal concepts. These more foundational legal concepts are "person", "movable object", and "possession of movable object by a person". These concepts are specified by the following meaning postulates:

(3a) A *person* is a born human.

(3b) A *movable object* is any physical entity that is not a person or a piece of land.

(3c) A *person possesses a movable object* if the person controls the uses and the location of the movable object.

Note that the concepts "movable object" and "possession of movable object by a person" refer to further foundational concepts, whereas the concept "person" is specified solely by referring to a factual description. △

Statements (1), (3), and (3a) to (3c) specify conditions that are sufficient for the applicability of the relevant condition. However, meaning postulates can, but do not have to, take this form—we now examine the various forms they may assume.

First, a meaning postulate for a concept can be given as an exhaustive list of entities that fall under

[2] This analysis refers to Ross's famous example "tû-tû" introduced in Ross [Ros57]. This example is based on the following norms (1) and (2). The first norm states that (1) if someone eats the food of the chief, then they are "tû-tû". This norm articulates a specification of "tû-tû" by stating a sufficient condition for its application. The second norm states that (2) if someone is "tû-tû", then they ought to undergo a purification ceremony. Thus, this norm articulates an associated normative consequence. As with the concept of a "ticketed bike", the concept "tû-tû" functions as an intermediate link. According to Ross, the normative content of (1) and (2) can be expressed even if this concept is omitted by the following norm: If someone eats the food of the chief, then they ought to undergo a purification ceremony. Based on this analysis, Ross argues that legal concepts are meaningless and can be omitted. Sartor [Sar09] and Gizbert-Studnicki and Klinowski [GK12] discuss this claim, as they still have an inferential function (see Sections 1.1.2 and 1.1.3).

the classification of the concept [Lin04].

Example 1.3. Assume that there are exactly three bikes in the Netherlands: bike_1, bike_2 and bike_3. Further, consider the following statement (4):

(4) The following bikes are "ticketed bikes": bike_1 and bike_2.

The concept "ticketed bike" is applicable to a given object if this object is one of the listed bikes. Thus, applicability is assessed by examining whether the situation in question is exactly one of the instances named in the exhaustive list defining the concept. △

Second, meaning postulates are presented as definitions that specify a legal concept through a set of conditions. The satisfaction of the conditions implies the applicability or non-applicability of the concept and hence is sufficient for the applicability or non-applicability [Lie87]. The statements (1), (3), and (3a) to (3c) are examples of such meaning postulates.

Third, a meaning postulate is constituted by a set of conditions without defining an implicative relation between these conditions and the applicability of the concept. In this case, the applicability of the concept is determined by the conditions collectively, without any one of them being decisive on its own [AB05a; AB05b; Wyn08].

Example 1.4. Consider the following meaning postulate that refers to a list of factual descriptions.

(5) The following conditions are further relevant to determine whether a bike is a "ticketed bike": (a) whether the bike is a cargo bike, (b) whether it is a support bike (c) whether it is suitable for transporting children. If all of these conditions are satisfied, then it is a "ticketed bike". Otherwise, it is not a "ticketed bike".

Unlike the previous meaning postulate, this characterization does not refer to sufficient conditions for the applicability of a concept. (5) provides a list of factual descriptions, and collectively, these descriptions determine whether or not the concept is applicable. △

1.1.2 Legal Concepts in Legal Norms

As discussed by Ross [Ros57] and Lindahl [Lin04], legal concepts simplify the representation of legal norms. Using the authors' argument, let us assume that $b_1, \ldots, b_n$ are individually sufficient conditions of the concept "ticketed bike" expressed by some meaning postulates. Further, $f_1, \ldots, f_m$ are legal consequences that need to be faced if a bike is considered a "ticketed bike". If the connection between the conditions and legal consequences is provided directly without referring to "ticketed bike", this results in the following presentation:

$$\begin{array}{ccccc}
b_1 \to f_1 & b_2 \to f_1 & b_3 \to f_1 & \ldots & b_n \to f_1 \\
b_1 \to f_2 & b_2 \to f_2 & b_3 \to f_2 & \ldots & b_n \to f_2 \\
\vdots & \vdots & \vdots & & \vdots \\
b_1 \to f_m & b_2 \to f_m & b_3 \to f_m & \ldots & b_n \to f_m
\end{array}$$

This representation requires to explicitly provide the link between any sufficient condition of "ticketed bike" and any of its legal consequences. Using the concept "ticketed bike", these legal norms can be expressed in a simpler fashion that reduces the number of statements needed to present legal norms. The first representation without the concept requires $n \times m$ many statements, whereas the following representation based on the concept necessitates only $n + m$ many statements.

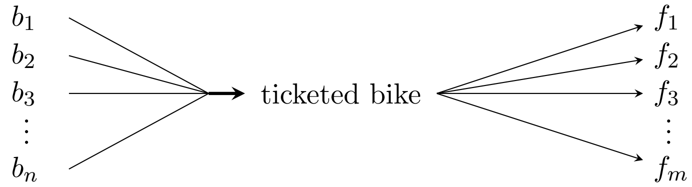

Legal concepts can not only be used to simplify certain norms. They also facilitate a simplification of entire legal systems. For this purpose, they can be used to systematize legal systems, where legal norms can be compared by examining whether they refer to the same legal concepts. This method of comparison can be used in general to compare any structure involving legal concepts [Wyn08].

1.1.3 Legal Concepts in Legal Reasoning

Legal concepts not only structure legal norms, but also guide legal reasoning. When legal reasoning deals with a situation and has to decide which legal consequence this situation entails, the legal concepts separate this step by first asking which legal concepts are applicable to the case and which legal consequences they entail [AB05a; AB05b]. The first step involves the question of whether the considered conditions meet the requirements to be classified by a legal concept.[3] For instance, before a conductor investigates whether a passenger needs a bike-ticket, they investigate whether the passenger's bike is a "ticketed bike". Based on this insight, the conductor then reasons towards the legal consequence that the passenger might need a bike-ticket.

The conductor's reasoning whether the passenger's bike is a "ticketed bike", is based on the information the conductor acquires. This means that reasoning about the applicability of concepts gives rise to an epistemic perspective. A factual situation is subjected to an inquiry that leads to an assessment of what is known about the situation. In accordance with this knowledge, the concept is determined to be applicable or not applicable [Kel17].

Besides the acquired information, the determination of the applicability of a concept depends on the meaning postulates specifying the concept. Although legal concepts might be specified by meaning postulates, it cannot be guaranteed that these postulates define the concept exhaustively. That is, a legal system cannot explicitly determine whether a concept is applicable to any imaginable and potentially unforeseeable circumstance [HM18; Fis91].

Example 1.5. Currently, Nederlandse Spoorwegen specifies that cargo bikes used for transporting goods are not allowed on trains and are therefore not considered a "ticketed bike". It also states that, in general, an unfolded folding bike is a "ticketed bike". Given the existence of folding *cargo* bikes, this raises the question of whether an unfolded folding *cargo* bike qualifies as a "ticketed bike". This represents a potential case in which Nederlandse Spoorwegen does not provide explicit information on the applicability of the concept. △

[3] This pattern of argumentation corresponds to what is called a legal syllogism. According to the legal syllogism, the application of a norm, which assigns a legal consequence to a legal classification, to a situation is based on the subsumption of that situation under the classification. For a characterization of the legal syllogism, see Wróblewski [Wró71; Wró74] and Joerden and Hilgendorf [JH21]. Further, see Duarte d'Almeida's [Dua19; Dua21] discussion of legal syllogism.

Consequently, the determination of the applicability of a legal concept to a situation might require a heuristic derivation. This means that the determination of the applicability of a concept to a case might be solely based on the meaning postulates, but not entailed by them. This highlights a significant property of legal concepts. They provide some flexibility so that the legal system can reason beyond what is explicitly defined.

1.2 Towards Modeling Legal Concept Applicability

The analysis of the classification using legal concepts conducted in Section 1.1 provides the foundation for modeling the determination of the applicability of legal concepts. According to this analysis, the applicability of a legal concept is determined based on the meaning postulates specifying the concept and further concepts mentioned in these meaning postulates themselves. As a result, the applicability of a legal concept can only be determined with respect to the applicability of another legal concept due to the hierarchical structure of legal concepts within a legal system.

However, for the sake of simplicity, we will focus exclusively on the most fundamental legal concepts—namely, those legal concepts whose applicability is independent of the applicability of other legal concepts. Consequently, all conditions of these concepts are descriptions of factual information. The focus on these concepts serves to better illustrate the formalism developed later by making a clear distinction between legal concepts and conditions that determine their applicability. However, this is not a limitation of the proposed formalism. In Chapter 6, we will outline how legal concepts whose applicability depends on other concepts can be formally modeled.

Furthermore, we will simplify the nature of legal concepts, conditions, and cases. We will treat conditions and concepts as non-negated or negated atomic propositions. In addition, we generalize the numerous entities to which concepts are applicable and which can satisfy conditions. We will conceptualize these entities as cases. Cases can be bicycles, individuals, or events, depending on whether we are considering the applicability of the concepts “ticketed bike”, “person” or “theft”, respectively.

Recall that legal reasoning is based on an epistemic perspective. Therefore, the determination of a concept’s applicability to a particular case is guided by the conditions that are known to be satisfied by the case, rather than the facts that the case actually satisfies. This epistemic position differentiates between the facts that are true about a case and the acquired knowledge about a case. This means that knowing that a case satisfies a condition does not imply that the case actually satisfies the condition. This allows for modeling legal reasoning realistically. Thus, our model accurately depicts real world legal scenarios, where the system possesses either incomplete or erroneous information, resulting in a discrepancy between the acquired information and the factual situation.

In summary, our subject of investigation—legal reasoning regarding the applicability of a legal concept to a case—can be characterized as follows: the determination of the applicability of a legal concept is based on both the legal information specifying the conditions for its applicability and the information obtained about the case. The legal information specifies the applicability or non-applicability of the concept in relation to certain conditions, which are factual descriptions. The information about the case consists of knowledge about which of these conditions are fulfilled.

1.3 Tasks of Modeling Legal Concept Applicability

Since we aim to formalize legal reasoning about the applicability of legal concepts, one needs to provide a formalism that is capable of deriving the applicability or non-applicability based on the meaning postulates that specify the concept and the obtained knowledge about the case. This objective leads to several independent tasks. To illustrate these tasks, let us consider the following specification of the concept of "ticketed bike".

Example 1.6. According to Nederlandse Spoorwegen, city bikes, including electric ones, are "ticketed bikes". Folding bikes, like folding city bikes, are not "ticketed bikes". However, if a folding bike is not folded or folded larger than 45 x 86 x 80 cm, then it falls under the category "ticketed bike". △

This specification of the concept "ticketed bike" provides multiple meaning postulates. The first task, we are investigating, is to formally represent such legal information.

> *Task 1: Formally specifying legal information on the applicability of concepts*
> Meaning postulates of legal concepts indicate the applicability of legal concepts with respect to conditions. To formally reason about the applicability of legal concepts, we want to represent this legal information. Importantly, such legal information must be capable of expressing the dependence on knowledge about satisfied conditions without necessitating a list of individually decisive conditions.

The legal information conveyed by the meaning postulates of the concept "ticketed bike", given by Example 1.6, provides a foundation to reason about whether unknown cases can be subsumed under this classification. As discussed in Section 1.1, determining the applicability of a legal concept to an unknown case might require a heuristic derivation of the concept's applicability. For instance, the meaning postulates in Example 1.6 do not explicitly determine whether "ticketed bike" is applicable to an electric bike that, when folded, is larger than 45 x 86 x 80 cm.

> *Task 2: Direct and heuristic derivation of the applicability of a concept to a case*
> In order to determine the applicability of a legal concept to a particular case, a comparison must be made between the information regarding the case and the legal information specifying the concept's applicability. This comparison gives rise to two possibilities. The concept's applicability to such a case is either explicitly stated or not. If the former is the case, the formalized meaning postulates can be used to derive directly the applicability or non-applicability of the concept to the case. Conversely, in the absence of such explicit information, the applicability can only be derived heuristically. Consequently, it can only be based on the formalized meaning postulates, and not derived from them. Thus, we want a formalism that is capable of directly or heuristically determining the applicability of a concept.

The determination of the applicability of a concept to a case is based on the information about the case. Additional information might result in a revised judgment regarding the applicability of the concept.

> *Task 3: Determining the stability of a judgment about the applicability of a concept to a case*
> In determining the applicability of a concept to a case, the question emerges whether the

applicability of the concept is unstable or stable. In other words, whether additional information about the case might result in a different judgment about the applicability of the concept. Consequently, the formalism must be able to reflect on the potential acquisition of knowledge.

For instance, if an object is known to be a city bike, then the concept of a "ticketed bike" is applicable to it. However, if the information that the object is a folding (city) bike is added, then the concept is not applicable anymore.

Determining the applicability of a concept is rarely limited to a single case. Often, the applicability of a legal concept must be assessed across multiple cases. This is addressed in the following task.

Task 4: Derivation of the applicability of a concept to a sequence of cases
To assess the applicability of legal concepts across multiple legal cases, the mechanism that determines a concept's applicability must be extended. We want to provide such an extension that is capable of assessing the applicability of a concept across a sequence of cases. It is important to note that the decision on the applicability of a legal concept in one case may influence its applicability in cases decided later. Therefore, this extension must take these influences into account when deciding subsequent cases.

Lastly, not every condition considered by meaning postulates strictly guides the applicability of a concept. This is because only some conditions distinguish whether the concept is applicable or not.

Task 5: Determining the conditions guiding the application of the concept
The applicability of legal concepts depends solely on conditions that distinguish applicability from non-applicability. This highlights the importance of identifying these conditions. Therefore, it is essential to assess which conditions are relevant to the applicability of a concept. We want to define and investigate the relevancy of such conditions.

For instance, in Example 1.6, whether the conditions "folding bike" and "folded larger than $45 \times 86 \times 80$ cm" are met is the only information needed to determine if a bike qualifies as a "ticketed bike". In contrast, whether a bike is a "city bike" does not affect this determination.

1.4 Approaches to Modeling Legal Concept Applicability

In this section, we will characterize some influential approaches to model the applicability of concepts to tasks. To embed our approach in the established literature, we outline our approach and indicate at the end of this section how it differs from the existing approaches.

The first two approaches present frameworks for case-based reasoning, a subcategory of non-monotonic reasoning. The third approach examines classificatory rules, while the fourth and fifth present frameworks for abstract reasoning and machine learning induction, respectively. All these have in common that they investigate legal classifications.

Approach 1: Case-based reasoning using HYPO and CATO

HYPO and CATO are representations of the conditions influencing a legal concept and of the arguments to reason about the applicability of the concept to a situation. They are used to investigate the concept of trade secret misappropriation [RAB05; Ash17].

HYPO and CATO characterize the factors that determine whether a case constitutes trade secret misappropriation. HYPO employs a dimensional representation of these factors, where each dimension favors either the defendant or the plaintiff to varying degrees. HYPO utilizes examples and explanations of analogous decided cases to argue for the plaintiff or defendant based on a description of a factual situation. In contrast, CATO does not employ dimensionality to represent factors. It utilizes a binary representation, indicating either the presence or absence of factors which again influences the presence or absence of more abstract factors. Thus, these factors are embedded in a hierarchical structure that illustrates which factors favor the plaintiff or the defendant [Ash17].

Approach 2: Reason and result model

The reason and result models are influential formal frameworks to formalize judicial precedents and precedential constraints. Doing so, they can be used to represent formal reasoning about the applicability of concepts and are therefore relevant approaches for our research question.

Horty [Hor11] introduces the reason model of precedent to formalize legal reasoning based on constraints imposed by precedent cases. The foundation of his analysis is a set of factors F that represent relevant facts or fact patterns. According to this model, a legal decision expressed in a case favors either the defendant δ or the plaintiff π. Each factor $f \in F$ supports either the defendant or the plaintiff. A case is represented as a tuple $\langle X, r, s\rangle$, where $X \subseteq F$ is the fact situation, r is a rule, and s is the outcome. The rule r takes the form $Y \to s$, where $Y \subseteq X$ and $s \in \{\delta, \pi\}$. That is, the rule maps a subset of the fact situation to the outcome, indicating whether the defendant or the plaintiff prevails. Note that a rule $r : Y \to s$ is restricted by the fact that any $f \in Y$ needs to favor the outcome s.

Based on these cases, a preference order is introduced. To characterize this preference, let us investigate Horty's example. Let $X = \{f_1^\delta, f_2^\delta, f_3^\delta, f_4^\delta, f_5^\pi, f_6^\pi\}$ be a fact situation where the superscript indicates whether the factor supports the plaintiff or defendant. Further, let the case be given by the following tuple $\langle X, r, \pi\rangle$ where $r : \{f_5^\pi, f_6^\pi\} \to \pi$. Due to this characterization, one can conclude that the legal decision prefers the reasons $\{f_5^\pi, f_6^\pi\}$ over the reasons $\{f_1^\delta, f_2^\delta, f_3^\delta, f_4^\delta\}$ resulting in a preference order on sets of factors. Given a set of cases, called case base, the expressed preference orders can be inconsistent. That is, there is a case such that $\{f_5^\pi, f_6^\pi\}$ is preferred over the reasons $\{f_1^\delta, f_2^\delta, f_3^\delta, f_4^\delta\}$ while another case prefers the latter over the former.

A consistent case base Γ can be used to constrain the judgment on undecided cases. Given a fact set X, the precedential constraints demand that the decision s is based on some rule r such that $\Gamma \cup \{\langle X, r, s\rangle\}$ is consistent. Since the condition only requires that it does not conflict with the preference orders of the case base, the rule r does not have to be used in any other case. In fact, the rule r can be any rule that does not impose a conflicting preference.

The result model, introduced in Horty [Hor04], proposes a strength ordering that compares the strength of sets of facts relative to their support of an outcome s. Informally, this strength ordering is defined as follows: A set Y of factors is at least as strong as a set Z of factors for outcome s if and only if Y contains all the facts supporting s contained by Z, while Z contains all the facts supporting the opposite of s contained by Y. Given this ordering, Horty defines *a fortiori* constraints. Such a constraint requires that a fact set X has to be decided in favor of outcome s if there exists a case such that X is at least as strong as the facts of the case, and the case is decided in favor of s.

The reason model and the result model spark active research. The following provides an overview of current research conducted on the model to offer insight into the ongoing debates.

While Canavotto [Can25] investigates how precedential constraints can be defined when the case base is inconsistent with respect to preferences over sets of reasons, Horty [Hor19] illustrates how dimensions can be integrated into the reason and result model. Dimensions refer to an ordered set of values where the order represents the outcome supported by the specific value. This extension is discussed also by Prakken [Pra21].

Canavotto and Horty [CH23a] generalize the reason model to include hierarchies of factors. This generalization contains a generalization of precedential constraints to hierarchical precedential constraints and a flattening technique to reduce hierarchical reason models to standard reason models. Based on this generalization, the same authors [CH23b] demonstrate the importance of legal concepts.

Furthermore, van Woerkom et al. [Van+23b] provides a generalized result model that integrates hierarchies of factors. Similar to the generalized reason model by Canavotto and Horty [CH23a], this generalization allows for reasoning in multiple steps through intermediate steps. Further, van Woerkom et al. [Van+23a] present a generalized result model that integrates factor hierarchies and dimensions.

These extensions, which can incorporate hierarchies of factors, have been discussed by Bench-Capon [Ben23; Ben24]. He argues against the use of hierarchies of factors to determine precedential constraints, favoring the standard models instead.

Approach 3: Formal analysis of counts-as conditionals

Various formalisms have been proposed for modeling counts-as conditionals that are statements of the form "X counts as Y in context C". Meaning postulates of legal concepts can be seen as what counts as a legal concept in a legal system. Therefore, the study of counts-as conditionals provides a framework to examine this classificatory aspect. To provide a brief introduction to this area of research, two influential logical formalizations and one algebraic formalization will be outlined. An exhaustive overview of relevant proposals is given by Grossi and Jones [GJ13].

Jones and Sergot [JS96] aim to formally represent institutional power by examining institutional classifications. An institutional classification refers to the idea that within an institution, there are usually rules according to which certain states of one type are considered states of another type. They introduce a new connective $\varphi \Rightarrow_s \psi$, which is interpreted as "φ counts as ψ in normative system s". To characterize the connective $\Rightarrow_s$, they introduce the following validities and rules within a propositional logic extended by this connective, in addition to the validities of standard propositional logic:

(V1) $((\varphi \Rightarrow_s \psi) \wedge (\varphi \Rightarrow_s \chi)) \rightarrow (\varphi \Rightarrow_s (\psi \wedge \chi))$

(V2) $((\psi \Rightarrow_s \varphi) \wedge (\chi \Rightarrow_s \varphi)) \rightarrow ((\psi \vee \chi) \Rightarrow_s \varphi)$

(V3) $((\varphi \Rightarrow_s \psi) \rightarrow ((\psi \Rightarrow_s \chi) \rightarrow (\varphi \Rightarrow_s \chi)))$

(R1) $\varphi \leftrightarrow \psi \vdash (\chi \Rightarrow_s \varphi) \leftrightarrow (\chi \Rightarrow_s \psi)$

(R2) $\varphi \leftrightarrow \psi \vdash (\varphi \Rightarrow_s \chi) \leftrightarrow (\psi \Rightarrow_s \chi)$

Validity (V1) expresses the conjunction of the consequents of the counts-as conditionals, while (V2) expresses the disjunction of their antecedents. Further, validity (V3) encodes the transitivity of counts-as conditionals. Finally, rules (R1) and (R2) ensure that both the antecedent and the consequent of the connective $\Rightarrow_s$ are closed under logical equivalence.

In contrast, Grossi et al. [GMD05; GMD08] provide an analysis of counts-as conditionals using modal logic. They distinguish between different types of classifications, including counts-as statements as constitutive rules that reflect legal classifications. According to their analysis, constitutive classifications occur as part of a set of rules Γ. In addition to this set Γ, the formalization of counts-as statements involves modal operators that express different contexts.

A context c is characterized as a subset of worlds $W_c \subseteq W$ and corresponds to a modal operator $[c]$, where $[c]\varphi$ is satisfied in a model if and only if φ is true in every world of W_c. Additionally, the modal operator $[-c]$ allows reasoning about worlds that are not part of W_c. That is, $[-c]\varphi$ is true if and only if φ is true in every world that is not contained in W_c. Given a set of formulas Γ such that $\varphi \to \psi \in \Gamma$, the constitutive classification of φ as ψ in context c is defined as follows:

$$\varphi \Rightarrow_{c,T}^{con} \psi ::= [c]\Gamma \wedge [-c]\neg\Gamma \wedge \neg[u](\varphi \to \psi)$$

where $[u]$ denotes the global context containing all worlds. This definition expresses that a constitutive classification of φ as ψ is valid if $\psi \to \psi$ is an element of Γ, Γ defines the context c and the classification of φ as ψ is not universally valid. Due to this definition, constitutive classifications are those that are explicitly articulated by the normative system expressed by Γ.

The algebraic analysis of the classifications of Lindahl and Odelstad [LO06; LO08; LO13] is embedded in a formalism that captures the intermediary role of legal concepts. For this purpose, they introduce joining systems. A joining system is a triple $\langle \mathcal{A}_1, \mathcal{A}_2, J\rangle$, where each $\mathcal{A}_i = \langle A_i, R_i\rangle$ is a quasi-ordering, meaning that R_i is a reflexive and transitive relation on A_i. The relation J relates elements of A_1 with elements of A_2, so that $J \subseteq A_1 \times A_2$. While Lindahl and Odelstad interpret the relations R_1 and R_2 as implicative relations on the sets A_1 and A_2, they interpret J as a relation that declares a correspondence between elements of A_1 and A_2. Using these joining systems, they specify the intermediary function of legal classifications. Given three joining systems $\mathcal{S}_1 = \langle \mathcal{A}_1, \mathcal{A}_2, J_{1,2}\rangle$, $\mathcal{S}_2 = \langle \mathcal{A}_2, \mathcal{A}_3, J_{2,3}\rangle$, and $\mathcal{S}_3 = \langle \mathcal{A}_1, \mathcal{A}_3, J_{1,3}\rangle$, the elements of A_2 can be interpreted as intermediaries. Therefore, for each correspondence $\langle a_1, a_3\rangle \in J_{1,3}$, there exists a link through A_2. That is, there exists $a_2 \in A_2$ such that $\langle a_1, a_2\rangle \in J_{1,2}$ and $\langle a_2, a_3\rangle \in J_{2,3}$.

Such a system represents the intermediate function of legal concepts discussed in Section 1.1. The elements of A_1 are interpreted as facts, while A_2 and A_3 are interpreted as legal concepts and consequences, respectively. A classificatory statement is modeled by the relation $J_{1,2}$, which connects facts with legal classifications, while the relation $J_{2,3}$ connects classifications with legal consequences. Based on this intuition, Lindahl and Odelstad explore properties of such intermediary elements by introducing further constraints on the joining systems.

Approach 4: ANGELIC and Abstract Dialectical Frameworks

Abstract Dialectical Frameworks (ADFs) were introduced by Brewka and Woltran [BW10] and Brewka et al. [Bre+13], and they extend Dung's [Dun95] argumentation framework. These frameworks have recently been used to provide computational implementations to assist legal reasoning about the applicability of certain concepts. An ADF is a tuple $(S, L, \{C_s\}_{s\in S})$, where S is a set of nodes and L is a binary relation on S that induces links connecting arguments. If $(a, b) \in L$, then a is a child of b. Furthermore, $V = \{0, 1\}$ is interpreted as the set of truth values. The set $\{C_s\}_{s\in S}$ is a collection of acceptance conditions given as functions of the form $(\text{children}(S) \to V) \to V$. That is, each function assigns a truth value to a node based on the truth values of the children of S.

Hierarchies of factors can be modeled using ADFs. For this purpose, let the set of nodes S be a set of factors. It is easy to see that the hierarchical structure can be stipulated by the relation L. This property is exploited by the so-called ANGELIC methodology, specified by Al-Abdulkarim et al. [Al-+18] based on Al-Abdulkarim et al. [AAB16], which uses ADFs to model factor hierarchies in CATO. Furthermore, they provide a computational implementation of the ADF that is capable of generating a judgment on the top-level factors if the truth values for the base-level factors are provided.

To exemplify their findings, Al-Abdulkarim et al. [AAB16] provide an analysis of U.S. trade secrets law, wild animal cases and the Automobile Exception to the Fourth Amendment domain. Collenette et al. [CAB20; CAB23] extend these findings by demonstrating that the same methodology can be applied to the legal domain of European Court of Human Rights cases.

Approach 5: Logic for binary classifier

Since the decision on the applicability of a concept is a binary classification, the formal investigation of these classifiers is of interest. Liu and Lorini [LL21] define a modal-logical framework to model binary classifiers and investigate the explainability of classifiers. This logic is based on classifier models, which are tuples $\langle S, f\rangle$. The set S is a set of states modeled as sets of propositions characterizing which features are satisfied at this state. The function $f : S \to \mathrm{Val}$ is a function that assigns values to states and models the classification. Since Liu and Lorini aim to develop a framework capable of formally reasoning about the explainability of binary classifiers, they introduce a modal operator $[X]$ to compare states. Informally, $[X]\varphi$ is satisfied by a state if φ is satisfied by any state that is indistinguishable with respect to the propositions in X. Using this modal operator, they characterize whether a classification of a state is biased, in the sense that there exists another state to which a different value is assigned, and the two states differ only with respect to protected features that do not permit different treatment.

Furthermore, Liu et al. [Liu+22] show that this logic for binary classifiers is capable of expressing case-based legal reasoning. They demonstrate that both the reason and the result model are expressible in this logic. Among other things, they define consistent case bases and constraints using modal logic formulas based on the operator $[X]$. Based on these results, modal logic techniques can be applied to study the base and result model. For example, Di Florio et al. [Di +23] investigate factors for which it is not yet determined which outcome they favor.

Our approach: Propositional dependence logic with the might-operator

In this work, we will propose a new approach to formalize legal reasoning about concept application. This approach is based on *propositional dependence logic with the might-operator* $\mathbf{PL}(=(\cdot), \blacklozenge)$, which is introduced in Section 3.1. The intuition of this model is given in Sections 2.1 and 2.2 and is logically characterized in Section 3.2.

Similarly to Approaches 2, 4 and 5, we will characterize concept applicability as a function from some set of satisfied conditions to the determination of whether a concept is applicable or not. These sets of satisfied conditions are referred to as factors in Approach 2, as nodes in Approach 4, and as features in Approach 5. In contrast to the reason and result model formalized in Approaches 2 and 5, we do not assume that conditions either favor applicability or disfavor applicability.

The representation of the functions specifying applicability will be based on meaning postulates. This foundation in meaning postulates relates to Approach 3. The presented works of this approach represent these meaning postulates in the sense that they formalize the connection between conditions and legal concepts for different classifications. However, our approach does not aim to formally explore different kinds of classifications, but rather to determine the applicability of a concept to cases.

In this sense, our approach pursues the same objective as Approaches 1, 2, 4 and 5. However, we will provide a new mechanism to determine the applicability of a concept to individual cases and to sequences of cases. This mechanism will be based on properties of sets of satisfied conditions with respect to concept applicability. This distinguishes our approach from Approaches 2 and 5. Furthermore, unlike in Approaches 1 and 4, this mechanism can reason beyond the information explicitly represented in the legal system. That is, it can determine the applicability of a concept to a case even when the system lacks explicit knowledge about whether the concept is applicable or not.

In addition, our formalism enables the direct handling of further tasks. First, given a judgment about the applicability of a concept to a case, our approach can determine whether acquiring additional knowledge about the case would alter the conclusion. Second, it provides sufficient means to identify which conditions are relevant for the applicability of a concept. These tasks are not directly addressed by the previously discussed proposals and therefore require extensions of the underlying models. Our formalism thus offers a unified perspective on these tasks that other formalisms do not directly provide.

Summarizing these aspects, we position our approach relative to the established literature and the tasks outlined in Section 1.3, as shown in Table 1.1. The tick (✓) indicates that the work in question explicitly deals with the task, while the cross (✗) means that this is not the case.

	[Hor04] [Hor11]	[CH23a] [CH23b]	[Van+23b] [Van+23a]	[JS96]	[GMD05] [GMD08]	[LO06] [LO08] [LO13]	[AAB16] [Al-+18] [CAB23]	[Liu+22]	New Proposal
Representation of Meaning Postulates	✓	✓	✓	✓	✓	✓	✓	✓	✓
Derivation of Applicability	✓	✓	✓	✗	✗	✗	✓	✓	✓
Heuristic Derivation of Applicability	✓	✓	✓	✗	✗	✗	✗	✓	✓
Stability of Prediction	✗	✗	✗	✗	✗	✗	✗	✗	✓
Derivation for Sequence of Cases	✓	✓	✓	✗	✗	✗	✗	✓	✓
Conditions Relevant for Applicability	✗	✗	✗	✗	✗	✗	✗	✗	✓

Table 1.1: Comparison of existing approaches and new proposal by task coverage

Chapter 2

An Intuition for Modeling Legal Concept Applicability

This chapter aims to provide intuitions regarding how the applicability of legal concepts can be logically determined. Legal concepts, conditions, cases, and legal knowledge about cases are characterized in Section 2.1. Meaning postulates articulating the applicability of legal concepts with respect to conditions are formalized as concept applicability functions, specified in Section 2.2.

2.1 Legal Concepts, Cases, Conditions and Legal Knowledge

Following Section 1.2, we consider legal reasoning on the applicability and non-applicability of legal concepts to be concerned with three types of elements: legal concepts, cases and conditions. The following symbols denote these elements in a formal context:

- $\mathcal{C} = \{c_1, \ldots, c_n\}$ is a finite and non-empty set of legal concepts
- $\mathcal{B} = \{b_1, \ldots, b_m\}$ is a finite and non-empty set of conditions
- $\mathcal{A} = \{a_1, \ldots, a_k\}$ is a finite and non-empty set of cases
- for any $a \in \mathcal{A}$, $K_a \subseteq \mathcal{B}$ is the associated set of conditions known to be satisfied by case a

For any $i, j \in \mathbb{N}$, B_i and A_j denote sets containing some conditions or cases, respectively. That is, $B_i \subseteq \mathcal{B}$ and $A_j \subseteq \mathcal{A}$. Recall that the determination of the legal applicability of a concept to a case is based on the acquired knowledge about the case. This is why any case $a \in \mathcal{A}$ is associated with a set of conditions K_a known to be satisfied by case a.

Conditions and concepts represent non-negated or negated atomic propositions. For instance, let the conditions b_i and b_j refer to "city bike" and "not folded", respectively. The statement that $b_i \in K_a$ expresses that it is known that "The a is a city bike". This does not imply that it is actually the case that a is a city bike, only that the legal system has concluded that it knows this. Similarly, $b_j \in K_a$ does not imply that a is actually not folded. Thus, the epistemic perspective does not allow any conclusions to be drawn about the real world facts of a case, and merely indicates what the legal system knows about the case under consideration, without this constituting a guarantee for the correctness of what is known.

Some conditions can conflict with one another in the sense that the legal system deems it impossible that each holds simultaneously for a specific case. Examples of conflicting pairs of conditions include "folded" and "not folded"; and "inside" and "outside". The current formalism lacks the capacity to reason about conflicting conditions. To model conflicting conditions formally, we fix for each $b_i \in \mathcal{B}$ a set of conditions $\mathcal{B}_i^{\perp} \subseteq \mathcal{B}$ such that $b_i \notin \mathcal{B}_i^{\perp}$ and for any $b_i, b_j \in \mathcal{B}$, $b_j \in \mathcal{B}_i^{\perp}$ if and only if $b_i \in \mathcal{B}_j^{\perp}$. The set $\mathcal{B}_i^{\perp}$ denotes the *set of opposing conditions* of b_i.

The first property of opposing conditions encodes that opposing conditions are irreflexive, while the later states that they are symmetric. This implies that if b_i opposes b_j, then b_j opposes b_i as well. Moreover, irreflexivity expresses that no condition can oppose itself. Note that $\mathcal{B}_i^{\perp}$ can be empty if $\mathcal{B}$ contains no conditions that oppose b_i.

Given a non-empty set of conditions $B \subseteq \mathcal{B}$, the *set of opposing conditions of* b_i *relative to* B is the set of conditions $B_i^{\perp}$ defined as $B_i^{\perp} := \mathcal{B}_i^{\perp} \cap B$. Thus, $B_i^{\perp}$ provides a restriction of the opposing conditions of a condition b_i to a set of conditions B.

Example 2.1. From now on, let shirt_1 refer to a specific shirt. Let b_1 denote "wearing shirt_1", b_2 denote "not wearing shirt_1", b_3 denote "storing shirt_1 in the wardrobe", and b_4 denote "not storing shirt_1 in the wardrobe". The legal system has determined that $\mathcal{B}_1^{\perp} = \{b_2, b_3\}$. This means that, within the context of legal system knowledge, it is assumed that b_1 and b_2 cannot simultaneously hold for a single case (for instance, a human). That is, a human cannot both wear shirt_1 and not wear shirt_1 at the same time. Further, it means that b_1 and b_3 cannot simultaneously hold for a human, and thus a human cannot simultaneously wear shirt_1 and store shirt_1 in the wardrobe. Lastly, this articulates that b_4 does not oppose b_1. Thus, the legal system renders it possible that a human wears and does not store shirt_1. $\triangle$

Legal reasoning is concerned with sets of conditions, such as the set of conditions K_a that a is known to satisfy. Since the objective is not to model legal reasoning involving conflicting or non-existent information, it is necessary to restrict the set of possible conditions constituting the set of conditions known to be satisfied by a case. For this purpose, let us introduce the notion of legal possibilities.

Definition 2.2 (Conflict-free, Informative, and Legal Possibility)**.** Let $B \subseteq \mathcal{B}$ be a set of conditions.

- B is *conflict-free* if and only if there does not exist $b_i, b_j \in B$ such that $b_j \in \mathcal{B}_i^{\perp}$.
- B is *informative* if and only if $B \neq \emptyset$.
- B is a *legal possibility* if and only if B is conflict-free and informative.

Let LP denote the set of legal possibilities. Given a case $a \in \mathcal{A}$, K_a is *legal knowledge* if and only if K_a is a legal possibility. Thus, legal knowledge about a case satisfies two constraints: the legal system knows at least something about the case, and the acquired legal knowledge about a case is consistent. Further, a case $a \in \mathcal{A}$ is *a legal case* in our framework if and only if K_a is legal knowledge. This means that legal cases are these cases where the associated knowledge is legal knowledge.

2.2 Concept Applicability Functions

Approaches 2 and 5 illustrate that legal decisions are, broadly speaking, a function that maps some sets of facts to a legal outcome. Following this functional understanding of representing legal decisions, we will characterize the applicability of legal concepts as functions.

Within our setting, meaning postulates express legal decisions. That is, meaning postulates state whether a concept is applicable or not applicable to certain sets of conditions. Formally, the meaning postulates of a concept are collectively represented as a concept applicability function that assigns Boolean values $\mathbf{1}$ and $\mathbf{0}$ to a set of legal possibilities. The Boolean values $\mathbf{1}$ and $\mathbf{0}$ are interpreted as "applicable" and "inapplicable", respectively.

Definition 2.3 (Concept Applicability Function)**.** Let $\mathcal{D} \subseteq (\mathcal{P}(LP) \setminus \emptyset)$ be a non-empty set of legal possibilities. A *concept applicability function* is a function $f : \mathcal{D} \to \{\mathbf{1}, \mathbf{0}\}$.

$\mathcal{P}(LP) \setminus \emptyset$ denotes the powerset of the set of legal possibilities, excluding the empty set. Intuitively, this demands that any legal decision about the applicability of some concept is based on some non-conflicting knowledge.

Further, the set $\bigcup \mathcal{D}$ is the set of *ground conditions of* f, and every condition $b \in \bigcup \mathcal{D}$ is called a ground condition of f. For any concept applicability function, the set of all ground conditions might be a legal possibility mapped to applicability or non-applicability by the function. However, this is not necessarily the case and the set of legal possibilities might only include several proper subsets of ground conditions.

Since the meaning postulates of a concept collectively determine its applicability relative to conditions, they are represented by a concept applicability function. For each legal system containing a concept, there exists a concept applicability function that formalizes its meaning postulates. Therefore, a legal system is modeled as a function that assigns a concept applicability function to each legal concept.

Definition 2.4 (Legal System)**.** Let F be the set of concept applicability functions. A *legal system* is a partial function $l : \mathcal{C} \to \mathcal{F}$.

The concept applicability function assigned to a concept c by a legal system l is called *the concept applicability function for c relative to legal system l.* Since we are only concerned with the applicability of legal concepts with respect to a single legal system, the reference "relative to legal system l" is omitted, and for each concept $c \in \mathcal{C}$, there exists exactly one concept applicability function, denoted by $f^c : \mathcal{D}^c \to \{\mathbf{1}, \mathbf{0}\}$. For any legal possibility $X \in \mathcal{D}^c$, we say that the *concept c is applicable to X* if $f^c(X) = \mathbf{1}$. Otherwise, we say that *concept c is not applicable to X.*

The domain $\mathcal{D}^c$ of any concept applicability function f^c is called *the set of decisive legal possibilities of c*, and the elements of $\mathcal{D}^c$ are called *decisive legal possibilities of concept c.* Since we are only concerned with one legal system, so that for each concept there exists exactly one concept applicability function, there exists exactly one set of decisive legal possibilities as well. Decisive legal possibilities are interpreted as sets that contain sufficient information so that the applicability or non-applicability of a concept is determined.

Example 2.5. Let us consider the concept of "possession of shirt$_1$" denoted by c_1. Note that "possession of shirt$_1$" in the legal sense is not identical to the concept of "ownership of shirt$_1$" and, roughly speaking, expresses that you have the shirt$_1$ at your disposal. Further, let us consider the following conditions: "wearing shirt$_1$", "not wearing shirt$_1$", "storing shirt$_1$ in the wardrobe", "not storing shirt$_1$ in the wardrobe". They are denoted by b_1, b_2, b_3 and b_4, respectively. Further, consider the following concept applicability function f^{c_1} specifying the applicability of the concept c:

$$f^{c_1}(X) = \begin{cases} \mathbf{1} \text{ if } X = \{b_1\}, X = \{b_3\}, X = \{b_1, b_4\}, \text{ or } X = \{b_2, b_3\} \\ \mathbf{0} \text{ if } X = \{b_2, b_4\} \end{cases}$$

Accordingly, the set of decisive legal possibilities of c_1 is $\mathcal{D}^{c_1} = \{\{b_1\}, \{b_3\}, \{b_1, b_4\}, \{b_2, b_3\}, \{b_2, b_4\}\}$. Since $\{b_1\}$ and $\{b_3\}$ are decisive legal possibilities, knowing only that a human (and hence a case) wears shirt_1 or stores shirt_1 in the wardrobe provides enough information to determine whether the legal system can classify this human as possessing shirt_1. Similarly, the fact that $\{b_2, b_3\}$, $\{b_1, b_4\}$, and $\{b_2, b_4\}$ are decisive legal possibilities indicates that if a legal system encounters a human about whom it knows that they wear shirt_1 and do not store it in the wardrobe, store shirt_1 in the wardrobe and do not wear it, or neither wear nor store it, then the legal system can infer whether this human qualifies as a possessor of shirt_1.

Also taking into account the Boolean values assigned to the decisive legal possibilities, the concept applicability function f^{c_1} expresses the following meaning postulates. If someone is known to be wearing shirt_1 or storing it in their wardrobe, then the concept "possession of shirt_1" is applicable to them. Also, if they are known to be wearing it but not storing it in their wardrobe, or storing it but not wearing it, then the concept "possession of shirt_1" is applicable to them as well. Conversely, if it is known that they neither wear it nor store it in their wardrobe, then the concept "possession of shirt_1" is not applicable to them. △

The relation between legal possibilities and judgment of applicability and non-applicability characterized by a concept applicability function provides a framework to reason about properties of legal possibilities with respect to this function. The following definition specifies the properties of sufficiency, exception, positive necessity, negative necessity and contingency relative to a concept applicability function.

Definition 2.6 (Properties of Legal Possibilities)**.** Let $f : \mathcal{D} \to \{\mathbf{1}, \mathbf{0}\}$ be a concept applicability function. We say that a legal possibility $X \in \mathcal{D}$ is

(1) *sufficient* relative to f if and only if for every $Y \in \mathcal{D}$ with $X \subseteq Y$ it holds that $f(Y) = \mathbf{1}$.

(2) *an exception* relative to f if and only if for every $Y \in \mathcal{D}$ with $X \subseteq Y$ it holds that $f(Y) = \mathbf{0}$.

(3) *positively necessary* relative to f if and only if for every $Y \in \mathcal{D}$ with $f(Y) = \mathbf{1}$ it holds that $X \subseteq Y$.

(4) *negatively necessary* relative to f if and only if for every $Y \in \mathcal{D}$ with $f(Y) = \mathbf{0}$ it holds that $X \subseteq Y$.

(5) *contingent* relative to f if and only if there exist $Y, Z \in \mathcal{D}$ such that $f(Y) = \mathbf{1}$, $f(Z) = \mathbf{0}$ and $X \subseteq Y, Z$.

If the concept applicability function is the concept applicability function f^c for concept c, then these propensities characterize properties of decisive legal possibilities with respect to the applicability of a concept. For instance, the sufficiency of decisive legal possibility relative to f^c characterizes the property of sufficiency for concept c, the negative necessity of decisive legal possibility relative to f^c characterizes the negative necessity for concept c and so on.

Example 2.7. Let c_1 denote the concept "possession of shirt_1", b_1 refer to the condition "wearing shirt_1", b_2 denote the condition "not wearing shirt_1", b_3 refer to the condition "storing shirt_1 in the wardrobe" and b_4 refer to the condition "not storing shirt_1 in the wardrobe". Furthermore, let the decisive legal possibilities and the concept applicability function f^{c_1} be defined as in Example 2.5. Then, the set of decisive legal possibilities for c is $\mathcal{D}^{c_1} = \{\{b_1\}, \{b_3\}, \{b_1, b_4\}, \{b_2, b_3\}, \{b_2, b_4\}\}$.

The decisive legal possibilities $\{b_1\}$ and $\{b_3\}$ are sufficient for c_1 because for all $Y \in \mathcal{D}^{c_1}$ with $\{b_1\} \subseteq Y$, or $\{b_3\} \subseteq Y$, it is the case that $f^{c_1}(Y) = \mathbf{1}$. This means that the concept c_1 is applicable to any decisive legal possibility of conditions that include either b_1 or b_3. However, neither $\{b_1\}$ nor $\{b_3\}$ are necessary for c_1 because $\{b_1\} \not\subseteq \{b_3\}$ and $\{b_3\} \not\subseteq \{b_1\}$. Since for all $Y \in \mathcal{D}^{c_1}$ with $\{b_2, b_4\} \subseteq Y$ it is the case that $f^{c_1}(Y) = \mathbf{0}$, it follows that the decisive legal possibility $\{b_2, b_4\}$ is an exception for c_1. Additionally, the decisive legal possibility $\{b_2, b_4\}$ is negatively necessary c_1 because for any $Y \in \mathcal{D}^{c_1}$ with $f^{c_1}(Y) = \mathbf{0}$ it holds that $\{b_2, b_4\} \subseteq Y$. Therefore, the concept c_1 is not applicable to a decisive legal possibility if and only if the decisive legal possibility contains $\{b_2, b_4\}$. △

Example 2.8. Let c_2 be the concept "ticketed bike" which defines the concept of bikes that can be taken on Dutch trains and that require a ticket. Further, let b_1 refer to the condition "city bike", b_2 refer to the condition "electric bike" and b_3 refer to the condition "folding bike". Consider the following concept applicability function f^{c_2} and the set of decisive legal possibilities $\mathcal{D}^{c_2} := \mathcal{P}(\{b_1, b_2, b_3\}) \setminus \{\emptyset\}$. That is, every non-empty set $X \subseteq \{b_1, b_2, b_3\}$ is a decisive legal possibility.

$$f^{c_2}(X) = \begin{cases} \mathbf{1} \text{ if } X = \{b_1\}, X = \{b_2\}, \text{ or } X = \{b_1, b_2\} \\ \mathbf{0} \text{ otherwise} \end{cases}$$

This function illustrates that whenever it is only known that the bike is a city bike or an electric, then it is considered a "ticketed bike". However, if it is known that the bike is a folding bike, then the bike is never a "ticketed bike". Therefore, the concept c_2 does not apply to any decisive legal possibility that contains the subset $\{b_3\}$. This illustrates why the decisive legal possibility $\{b_3\}$ is an exception c_2: For any decisive legal possibility X such that $\{b_3\}$ is a subset of it, it is the case that c_2 is not applicable to this decisive legal possibility. Furthermore, the decisive legal possibilities $\{b_1\}$, $\{b_2\}$ and $\{b_1, b_2\}$ are contingent decisive legal possibilities. To see this, first note that c_2 is applicable to these decisive legal possibilities. Second, consider the decisive legal possibility $\{b_1, b_2, b_3\}$. It is the case that the decisive legal possibilities $\{b_1\}$, $\{b_2\}$ and $\{b_1, b_2\}$ are subsets of $\{b_1, b_2, b_3\}$ and c_2 is not applicable to $\{b_1, b_2, b_3\}$. This means that the decisive legal possibilities $\{b_1\}$, $\{b_2\}$ and $\{b_1, b_2\}$ are subsets of decisive legal possibilities to which the concept is applicable and subsets of decision sets to which the concept is not applicable. △

It is important to note that, in general, any decisive legal possibility that is a superset of an exception is also an exception. The following lists this and other similar facts.

Fact 2.9. Let $f : \mathcal{D} \to \{\mathbf{1}, \mathbf{0}\}$ be a concept applicability function.

(1) If $X \in \mathcal{D}$ is sufficient (an exception) relative to f, then any $Y \in \mathcal{D}$ with $X \subseteq Y$ is sufficient (an exception) relative to f.

(2) If $X \in \mathcal{D}$ is positively (negatively) necessary relative to f, then any $Y \in \mathcal{D}$ with $Y \subseteq X$ is is positively (negatively) necessary relative to f.

(3) $X \in \mathcal{D}$ is contingent relative to f if and only if X is neither sufficient nor an exception relative to f.

Proof. Follows immediately from Definition 2.6. □

Let us briefly summarize the core notions and their function within our formalism. Conditions, denoted by $b_1, \dots, b_m$, are negated or non-negated atomic propositions. Sets of conditions that are informative and conflict-free are legal possibilities.

Cases, denoted by $a_1, \ldots, a_k$, are associated with sets of conditions. These sets are denoted by $K_{a_1}, \ldots, K_{a_k}$ and represent the acquired knowledge of the legal system about a case—that is, information about which conditions the case satisfies. If the associated knowledge of a case is a legal possibility, the the associated knowledge is legal knowledge and the case is a legal case.

Legal concepts are denoted by $c_1, \ldots, c_n$. Like conditions, they are negated or non-negated atomic propositions, and they are applicable to cases, as will be examined in Chapter 4.

The applicability of legal concepts is modeled using concept applicability functions, which assign Boolean values to sets of legal possibilities. Given a legal system, there exists one concept applicability function for each concept in the legal system, specifying the applicability of the concept according to that system. We will provide a logical characterization of concept applicability functions in Chapter 3.

Chapter 3

A Logic for Legal Concept Applicability

In this chapter, we will address Task 1. We will present a logical framework for modeling legal information regarding the applicability of concepts based on the previously formalized notions of concepts, cases, decisive legal possibilities, and concept applicability functions. Legal information is modeled using a variant of propositional dependence logic. We describe the logic in Section 3.1 and the modeling of legal reasoning in Section 3.2.

3.1 The Logics $\mathbf{PL}(=(\cdot))$ and $\mathbf{PL}(=(\cdot), \blacklozenge)$

In our daily lives, functional dependencies are a regular feature: the cost of a train ticket depends on the age of the passenger or the gas mileage of a car depends on the speed. The fact that the feature "gas mileage of a car" functionally depends on another feature such as "speed" means that the value of the second feature determines the value of the first feature. To formally represent these dependencies, classical logics, like classical propositional and classical first-order logic, are not sufficient. As noted by Väänänen [Vää07a] and Anttila [Ant25], dependencies demand the investigation of multiple data points, rather than a single data point. This is the case because if only one data point is considered, any dependency trivially holds. To illustrate this, let us consider the following data about restaurants.

Restaurant	Pizza	Salad	Pasta
restaurant_1	1	1	1
restaurant_2	1	1	0
restaurant_3	1	1	0
restaurant_4	0	0	1

Table 3.1: Data about restaurants

Each restaurant specifies a data point that specifies values for the features "pizza", "salad" and "pasta" where 1 indicates that the restaurant offers such a dish and 0 indicates that the restaurant does not offer such a dish. To say that the feature "pizza" functionally depends on the feature "pasta" means that there is a function such that it determines whether a restaurant offers pizza solely based on whether that restaurant offers pasta.

If only one restaurant and thus a single data point is examined, then each of these values depends on each of the other values. For example, the feature "pizza" depends on the feature "pasta" and vice versa. However, when each restaurant is considered, resulting in multiple data points, neither of the features “pizza” and “pasta” depends on the other. This is because restaurant_1 and restaurant_4 offer pasta, but they don't both offer pizza, nor do they both not offer pizza. Similarly, restaurant_1 and restaurant_2 offer pizza, but they don't both offer pasta, nor do they both not offer pasta.

In addition to these insights, Table 3.1 allows us to examine several other dependencies. Although the fact that a restaurant offers salad does not depend on whether it offers pasta, it does depend on whether it offers pizza. This is because every restaurant that offers pizza also offers salad, so there exists the following function determining whether a restaurant offers salad based only on whether it offers pizza. This function articulates that a restaurant offers salad if it offers pizza. Analogously, the fact that a restaurant offers pizza depends on whether it offers salad.

A logic that can express functional dependencies is *dependence logic*, introduced by Väänänen [Vää07a; Vää07b]. In comparison to classical propositional logic, dependence logic is based on *team semantics* and includes *dependence atoms*. Further, we will enrich dependence logic with the *might-operator*. Let us introduce these components intuitively before defining them formally.

The foundation of the semantics of classical propositional logic is a valuation function. This function assigns a Boolean value to each atomic proposition. Based on this assignment, the truth value of more complex propositions is obtained. Importantly, to interpret a formula only one valuation is considered in classical propositional logic. A *team* generalizes this idea so that a team is a set of valuation functions. Within *team semantics* the truth value of propositions are then discussed with respect to teams of valuations. This generalization provides a framework to reason about the relationships between individual valuations and the information they encode. For instance, in Table 3.1 each restaurant instantiates a valuation assigning values to the discussed features. Based on these valuations, patterns of dependencies between the features are analyzable.

These patterns of functional dependency are expressible using dependence atoms. Given that $x_1, \ldots, x_n$ and y are features, the dependence atom $=(x_1, \ldots, x_n, y)$ expresses that the feature y functionally depends on the features $x_1, \ldots, x_n$ which means that there exists a function from the values of the features $x_1, \ldots, x_n$ to the value of the feature y. In the previously discussed example, $=(\text{pizza}, \text{salad})$ holds because the value for “pizza” determines the value for “salad”.

To indicate that there exists some restaurant offering pasta, the might-operator, denoted by $\blacklozenge$, can be used. In Table 3.1, one can state that there exists a non-empty set of restaurants offering pasta using the might-operator. Therefore, the operator expresses that something might be the case in the sense that, given that you are in a $\text{restaurant}_\text{i}$ but do not know which one, it cannot be ruled out that the restaurant offers pasta. Thus, the $\text{restaurant}_\text{i}$ might offer pasta. In general, the might-operator allows one to state the existence of a non-empty set of valuations that satisfy certain constraints.[4]

To formally introduce the dependence atom and the might-operator, we characterize *propositional dependence logic*, denoted by $\mathbf{PL}(=(\cdot))$, and *propositional dependence logic with the might-operator*, denoted by $\mathbf{PL}(=(\cdot), \blacklozenge)$. Further, we will specify the logic $\mathbf{PL}$, which is team-based classical propositional logic. Propositional dependence logic $\mathbf{PL}(=(\cdot))$ was introduced by Yang and Väänänen [YV16]. To our knowledge, the logic $\mathbf{PL}(=(\cdot), \blacklozenge)$ has not been studied in the literature, but it is similar to a logic discussed by Anttila and Knudstorp [AK25].

[4]For more on the might-operator, see Veltman [Vel96], Hella and Stumpf [HS15], Yan [Yan23], and Anttila and Knudstorp [AK25].

Definition 3.1 (Well-formed Formulas)**.** Fix a (countably infinite) set $PROP$ of propositional variables. We recursively define the set of well-formed formulas of **PL**, $\mathbf{PL}(=(\cdot))$ and $\mathbf{PL}(=(\cdot), \blacklozenge)$:

$$\begin{array}{ll}
\alpha ::= p \mid \neg\alpha \mid \alpha \wedge \alpha \mid \alpha \vee \alpha & \mathbf{PL} \\
\varphi ::= p \mid \neg\alpha \mid =(p_1, \ldots, p_n, q) \mid \varphi \wedge \varphi \mid \varphi \vee \varphi & \mathbf{PL}(=(\cdot)) \\
\varphi ::= p \mid \neg\alpha \mid =(p_1, \ldots, p_n, q) \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \blacklozenge\varphi & \mathbf{PL}(=(\cdot), \blacklozenge)
\end{array}$$

where $p, p_1, \ldots, p_n, q \in PROP$.

In the following α always denotes a formula of **PL**. As previously discussed, the dependence atom $=(p_1, \ldots, p_n, q)$ expresses "q depends on $p_1, \ldots, p_n$". Further, the constancy atom $=(q)$ denotes "q is constant". The connective $\vee$ is called the tensor disjunction and generalizes the classical disjunction in the team-based setting. Lastly, $\blacklozenge$ is the might-operator where $\blacklozenge\varphi$ is interpreted as "might φ".

We define an implication $\alpha \to \psi := \neg\alpha \vee \varphi$ where $\alpha \in \mathbf{PL}$ and $\varphi \in \mathbf{PL}(=(\cdot), \blacklozenge)$. This means that the antecedent of the implication has to be a classical formula, whereas the consequent can be any formula of $\mathbf{PL}(=(\cdot), \blacklozenge)$. This restriction is necessary because the negation is only defined for formulas of **PL**.

To define the semantics of the dependence atom $=(\cdot)$, the semantics of the connectives $\neg, \wedge, \vee$ and the semantics of the might-operator $\blacklozenge$ within team semantics, we first need to define valuations and teams.

Definition 3.2 (Valuation and Team)**.** Given a set of propositional variables $P \subseteq PROP$, a *valuation v on P* is a function $v : P \to \{\mathbf{1}, \mathbf{0}\}$, which means that a valuation assigns a Boolean value to the propositions of P. A *team T on P* is a set of valuations on P and hence $T \subseteq \{\mathbf{1}, \mathbf{0}\}^P$.

The empty team $\emptyset$ does not contain any valuation, whereas the full team on P contains every possible valuation on P.

Team semantics stipulates that the satisfiability of well-formed formulas is defined with respect to teams rather than individual valuations.

Definition 3.3 (Team Semantics)**.** Given a set of propositions $P \subseteq PROP$, a team T on P, and a well-formed formula φ of **PL**, $\mathbf{PL}(=(\cdot))$, or $\mathbf{PL}(=(\cdot), \blacklozenge)$ such that every propositional variable occurring in φ is contained in P, the satisfiability of φ in T is recursively defined:

(1) $T \models p$ if and only if for all $v \in T$, $v(p) = \mathbf{1}$

(2) $T \models \neg\alpha$ if and only if for all $v \in T$, $\{v\} \not\models \alpha$

(3) $T \models =(p_1, \ldots, p_n, q)$ if and only if for all $v_i, v_j \in T$ with $v_i(p_k) = v_j(p_k)$ for $1 \leq k \leq n$, it is the case that $v_i(q) = v_j(q)$

(4) $T \models \varphi \wedge \psi$ if and only if $T \models \varphi$ and $T \models \psi$

(5) $T \models \varphi \vee \psi$ if and only if there exist subteams $T_1, T_2 \subseteq T$ with $T_1 \cup T_2 = T$ such that $T_1 \models \varphi$ and $T_2 \models \psi$

(6) $T \models \blacklozenge\varphi$ if and only if there exists a subteam $T' \subseteq T$ with $T' \neq \emptyset$ and $T' \models \varphi$

The constancy atom $=(q)$ reduces the satisfiability conditions of the dependence atom to the specification that $T \models =(q)$ if and only if for every $v_j, v_i \in T$ it is the case that $v_i(q) = v_j(q)$.

If a team T satisfies a formula φ, then φ is said to be *true* in T. Alternatively, φ is said to *hold* in T. In contrast, if φ is not satisfied by a team T, then φ is said to be *false* in T and to *not hold* in T. Further, if a formula φ is satisfied by all teams, then φ is said to be *valid*.

To provide some intuition about the semantics of the tensor disjunction, dependence atom and the might-operator, consider the following example.

Example 3.4. Let us consider the following team T on $P = \{p_1, p_2, p_3, p_4\}$.

	p_1	p_2	p_3	p_4
v_1	**1**	**0**	**1**	**1**
v_2	**1**	**0**	**1**	**1**
v_3	**1**	**1**	**1**	**1**
v_4	**0**	**0**	**1**	**0**

Observe that the proposition p_3 is constant because for any valuations $v_i, v_j \in T$ it is the case that $v_i(p_3) = v_j(p_3) = \mathbf{1}$. However, neither the proposition p_1, p_2 nor p_3 is constant because for any $p \in \{p_1, p_2, p_4\}$, there exist $v_i, v_j \in T$ such that $v_i(p) \neq v_j(p)$. For instance, v_1 assigns the value **1** to proposition p_1, whereas v_3 assigns the value **0** to proposition p_1 so that $T \not\models =\!(p_1)$.

Note that $T \models =\!(p_1, p_3)$ and $T \models =\!(p_1, p_4)$. This means that for all valuations of the team with the same Boolean value assigned to p_1, the Boolean value for p_3 and p_4 are identical. In contrast, $T \not\models =\!(p_1, p_2)$ because v_2 and v_3 assign the same value to p_1 but not the same value to p_2.

Consider the following subteams $T_1 = \{v_1, v_2, v_4\}$ and $T_2 = \{v_3\}$. Since for any $v \in T_1$ it is the case that $v(p_2) = \mathbf{0}$, it follows that $T_1 \models \neg p_2$. Similarly, since for any $v \in T_1$ it is the case that $v(p_2) = \mathbf{1}$ it follows that $T_2 \models p_2$. Due to the fact that $T_1 \cup T_2 = T$, it is the case that $T \models \neg p_2 \vee p_2$.

Lastly, $T \models \blacklozenge\, (p_1 \wedge p_3)$. This means that there exists a non-empty subteam $T' \subseteq T$ such that $T \models p_1 \wedge p_3$. This holds because $v_2(p_1) = \mathbf{1}$ and $v_2(p_3) = \mathbf{1}$, so that $T' = \{v_2\}$ satisfies these constraints. $\triangle$

Well-formed formulas of $\mathbf{PL}(=\!(\cdot), \blacklozenge)$ satisfy certain properties with respect to teams, indicating whether a formula remains satisfied after set-theoretic manipulation of the team.

Theorem 3.5. Let α be a formula of **PL** and φ be a formula of $\mathbf{PL}(=\!(\cdot))$. Let T and T' be teams.

(1) $\emptyset \models \varphi$ — Empty Team Property

(2) If $T' \subseteq T$ and $T \models \varphi$, then $T' \models \varphi$ — Downward Closure Property

(3) $T \models \alpha$ if and only if for all $v \in T$ it is the case that $\{v\} \models \alpha$ — Flatness Property

Proof. Proof by induction on the complexity of α and φ. $\square$

The empty team property states that every formula of $\mathbf{PL}(=\!(\cdot))$ is satisfied by the empty team. Hence, contradictions like $p \wedge \neg p$ are true in the empty team. The downward closure property states that if a team satisfies a formula of $\mathbf{PL}(=\!(\cdot))$, then any subset of that team satisfies the formula as well. The flatness property articulates that a formula of **PL** is true in a team if and only if for every valuation of the team it holds that the formula is true in the singleton team containing only this valuation.

Importantly, the flatness property does not hold for every formula of $\mathbf{PL}(=\!(\cdot))$ and $\mathbf{PL}(=\!(\cdot), \blacklozenge)$. This is the case because any stated dependence is true in a singleton team, but might not hold in a team containing more than one valuation. Further, the empty team property and the downward closure property do not hold for formulas of $\mathbf{PL}(=\!(\cdot), \blacklozenge)$. To see this, observe that any $\blacklozenge\varphi$ where φ is a well-formed formula of $\mathbf{PL}(=\!(\cdot), \blacklozenge)$ does not hold in the empty team due to the semantics of the might-operator.

Furthermore, the team-based semantics for **PL** over singleton teams coincide with the semantics of classical propositional logic. Let $v \models_{CPL} \alpha$ mean that v satisfies α in the sense of classical propositional logic. It is easy to prove that the following proposition holds.

Proposition 3.6. Given a formula α of **PL** and a team T, $T \models \alpha$ if and only if for all $v \in T$, $\{v\} \models \alpha$ if and only if for all $v \in T$, $v \models_{CPL} \alpha$.

Proof. Follows from Theorem 3.5 and induction on the complexity of α. □

Proposition 3.6 implies that the set of validities of **PL** is identical to the set of validities of classical propositional logic. Thus, validities of classical propositional logic, such as double negation elimination, can be utilized when working with **PL**. Furthermore, this means that $\mathbf{PL}(=(\cdot))$ and $\mathbf{PL}(=(\cdot), \blacklozenge)$ are conservative extensions of classical propositional logic.

3.2 $\mathbf{PL}(=(\cdot), \blacklozenge)$ for Legal Concept Applicability

Legal reasoning about the application of a concept is based on legal information about its application that is provided by meaning postulates. To represent this information, we will use tables where each entry of the table characterizes a decisive legal possibility and the information whether the concept is applicable to this decisive legal possibility. To logically reason about these tables, they will be specified using propositional dependence logic with the might-operator $\mathbf{PL}(=(\cdot), \blacklozenge)$ based on team semantics.

Definition 3.7 (Propositions)**.** Given the set of concepts $\mathcal{C}$ and the set of conditions $\mathcal{B}$, the set of propositions $PROP$ contains $\mathcal{C}$ and $\mathcal{B}$. That is, $B \cup \mathcal{C} \subseteq PROP$.

This means that we will treat conditions and concepts as propositions of the logic. These propositions are the foundations of the syntax of $\mathbf{PL}(=(\cdot), \blacklozenge)$.

In a table representation, the conditions and concepts are the columns of our table. In an entry of the table, these propositions are mapped to a specific value. This mapping is done by a valuation which specifies a row of the table. A valuation represents a meaning postulate and hence a data point that specifies the applicability of a concept. According to Definition 3.2, a valuation v on a set of propositional variables $P \subseteq \mathcal{B} \cup \mathcal{C}$, where $\mathcal{C}$ is the set of concepts and $\mathcal{B}$ is the set of conditions, is a function that maps each propositional variable to a Boolean value. That is, $v : P \to \{\mathbf{1}, \mathbf{0}\}$. Further, a team T is a set of valuations on P. Thus, a team T on P is a tabular representation of multiple data points that indicate the applicability or non-applicability of the concepts $c_1, \ldots, c_i$ relative to the conditions $b_1, \ldots, b_j$.

The Boolean values **1** and **0** are interpreted differently with respect to conditions and concepts. Given a condition b and a valuation v, $v(b) = \mathbf{1}$ denotes "The legal system knows that b is satisfied" and $v(b) = \mathbf{0}$ denotes "The legal system does not know that b is satisfied". This means that neither $v(b) = \mathbf{1}$ nor $v(b) = \mathbf{1}$ have any bearing on whether or not condition b is actually satisfied. As discussed in Section 2.1, this ensures that the proposed formalism focuses on what the legal system knows instead of what is actually the case. Further, $v(c) = \mathbf{1}$ denotes "Concept c is applicable according to the legal system" and $v(c) = \mathbf{0}$ denotes "Concept c is not applicable according the legal system".

To reason about the applicability of legal concepts within $\mathbf{PL}(=(\cdot), \blacklozenge)$, we will proceed as follows. First, we will provide a translation of concept applicability functions into teams. Second, we will define *legal teams* as those teams that are translations of concept applicability functions, and third, we will characterize the properties of legal teams.

Definition 3.8 (Translation of Concept Applicability Function)**.** Given a concept applicability function $f : \mathcal{D} \to \{\mathbf{1}, \mathbf{0}\}$, a team T translates f just in case there is some concept $c \in \mathcal{C}$ such that the domain of T is $\bigcup \mathcal{D} \cup \{c\}$ and the following conditions are satisfied:

(1) There exists a valuation $v \in T$ with truth set $Y_v = \{b \in \bigcup \mathcal{D} \mid v(b) = \mathbf{1}\}$ and falsity set $Z_v = \{b \in \bigcup \mathcal{D} \mid v(b) = \mathbf{0}\}$ if and only if there exists a legal possibility $X \in \mathcal{D}$ such that $Y_v = X$ and $Z_v = \bigcup \mathcal{D} \setminus X$.

(2) For each valuation $v \in T$ with truth set $Y_v = \{b \in \bigcup \mathcal{D} \mid v(b) = \mathbf{1}\}$, it is the case that $v(c) = \mathbf{1}$ if and only if $f(Y_v) = \mathbf{1}$.

Recall that, according to Definition 2.3, the domain $\mathcal{D}$ of any concept applicability function f is a non-empty set of legal possibilities such that $\bigcup \mathcal{D}$ is a non-empty set of conditions. Consequently, a team T that translates f is a team on $B \cup \{c\}$, where c is a concept and B is a non-empty set of conditions such that $\bigcup \mathcal{D} = B$.

Based on these translations of concept applicability functions, we can define *legal teams*.

Definition 3.9 (Legal Team)**.** Let $c \in \mathcal{C}$ be a concept and $B \subseteq \mathcal{B}$ be a non-empty set of conditions.

- A team T on $B \cup \{c\}$ is a *legal team* if and only if there exists a concept applicability function f such that T translates f. We denote a legal team T on $B \cup \{c\}$ by T^c.
- Given a legal system l, T^{c_j} is *the legal team of* c_i *in* l if and only if T^{c_j} translates $l(c_i)$ and $c_i = c_j$.

As defined in Definition 2.4, a legal system l is a partial function assigning concept applicability functions to legal concepts. Thus, there exists (at most) one concept applicability function f^{c_i} for c_i relative to this legal system, where $l(c_i) = f^{c_i}$. Given this concept applicability function f^{c_i}, it is easy to see that there can be only one legal team T^{c_j} such that T^{c_j} translates f^{c_i} and $c_i = c_j$. This is why we are justified in speaking of *the* legal team of c_i in l. Since we are only concerned with one legal system, we omit mentioning the legal system when referring to the legal team of a concept.

Given a legal team T^c and a valuation $v \in T^c$, we say that the legal possibility X, defined as $X = \{b \in B \mid v(b) = \mathbf{1}\}$, is the *legal possibility expressed by* v. If a legal possibility X is expressed by some $v \in T^c$, we say that *the legal possibility* X *is expressed by* T^c. If a legal team T^c is the legal team of c, then the legal possibilities expressed by T^c are *decisive legal possibilities* for c.

Further, as for any concept applicability function $f : \mathcal{D} \to \{\mathbf{1}, \mathbf{0}\}$, $\bigcup \mathcal{D}$ is its set of ground conditions, we will call any condition $b \in B$ a *ground condition of* T^c where T^c is a legal team on $B \cup \{c\}$.

According to Definitions 3.8 and 3.9, legal teams are a tabular representation of a concept applicability function. To illustrate this, let us consider the following example.

Example 3.10. Let us examine the concept c_1 expressing "possession of shirt$_1$". As before, the set of conditions includes the condition "wearing shirt$_1$", the condition "not wearing shirt$_1$", the condition "storing shirt$_1$ in the wardrobe", and the condition "not storing shirt$_1$ in the wardrobe". These conditions are denoted by b_1, b_2, b_3, and b_4, respectively. Let $\mathcal{B}_1^\perp = \{b_2, b_3\}, \mathcal{B}_2^\perp = \{b_1\}, \mathcal{B}_3^\perp = \{b_1, b_4\}$ and $\mathcal{B}_4^\perp = \{b_3\}$ define the sets of opposing conditions. Further, consider the following concept applicability function $f^{c_1} : \mathcal{D}^{c_1} \to \{\mathbf{1}, \mathbf{0}\}$ that specifies the applicability of c_1:

$$f^{c_1}(X) = \begin{cases} \mathbf{1} \text{ if } X = \{b_1\}, X = \{b_3\}, X = \{b_1, b_4\}, \text{ or } X = \{b_2, b_3\} \\ \mathbf{0} \text{ if } X = \{b_2, b_4\} \end{cases}$$

Next, we turn to the legal team T^{c_1} on $B \cup \{c_1\}$ where $B = \{b_1, b_2, b_3, b_4\}$, shown in Table 3.2.

valuation v_i	legal possibility X_i	b_1	b_2	b_3	b_4	concept c_1
v_1	$\{b_1\}$	**1**	**0**	**0**	**0**	**1**
v_2	$\{b_3\}$	**0**	**0**	**1**	**0**	**1**
v_3	$\{b_1, b_4\}$	**1**	**0**	**0**	**1**	**1**
v_4	$\{b_2, b_3\}$	**0**	**1**	**1**	**0**	**1**
v_5	$\{b_2, b_4\}$	**0**	**1**	**0**	**1**	**0**

Table 3.2: Legal team T^{c_1} of the concept "possession of shirt$_1$"

Let us sketch why T^{c_1} translates f^{c_1} and hence is the legal team of c_1. First, it is straightforward to verify that $\bigcup \mathcal{D}^{c_1} = B$.

Second, note that every truth set of a valuation $v_i \in T^{c_1}$ and hence every legal possibility X_i expressed by T^{c_1} is an element of the domain of f^{c_1}. That is, $X_i \in \mathcal{D}^{c_1}$. Further, for each $X_i \in \mathcal{D}^{c_1}$ there exists a valuation $v_i \in T^{c_1}$ such that X_i is the legal possibility expressed by v_i. Moreover, for each valuation of T^{c_1} the falsity set is the same as the set of ground conditions for c_1, excluding the legal possibility expressed by this valuation. Thus, the first condition of translation is satisfied.

Third, the concept c_1 is applicable to a legal possibility according to f^{c_1} if and only if the same legal possibility is expressed by some valuation of T^{c_1} which assigns the concept c_1 the value **1**. Therefore, the second condition of translation is satisfied. $\triangle$

We now present a key result on legal teams: Legal teams can be precisely characterized by the satisfaction of conditions corresponding to properties of concept applicability functions.

Theorem 3.11 (Characterization of Legal Teams)**.** Let $B \subseteq \mathcal{B}$ be a non-empty set of conditions, $c \in \mathcal{C}$ be a concept and T be a team on $B \cup \{c\}$ where $B = \{b_1, \dots, b_n\}$. The team T is a legal team if and only if it satisfies the following constraints:

(1) $T \models =\!(b_1, \dots, b_n, c)$ (2) $T \neq \emptyset$

(3) $T \models \bigvee_{b_i \in B} b_i$ (4) for each $b_i \in B$, $T \models \blacklozenge b_i$

(5) for each $b_i \in B$, $T \models \bigwedge_{b_j \in B_i^{\perp}} \neg(b_i \wedge b_j)$

where for each condition b_i, $B_i^{\perp}$ is the set of opposing conditions relative to B.

Proof. To show the left-to-right direction, suppose that T is a legal team. Therefore, there exists a concept applicability function $f : \mathcal{D} \to \{\mathbf{1}, \mathbf{0}\}$ such that T translates f. It needs to be shown that T satisfies the conditions (1) to (5).

To show that (1) is satisfied, let $v_i, v_j \in T$ be two valuations such that $v_i(b) = v_j(b)$ for all $b \in B$. It needs to be shown that $v_i(c) = v_j(c)$. By Definition 3.8, it follows that there exist $X_i, X_j \in \mathcal{D}$ such that $X_i = \{b \in B \mid v_i(b) = \mathbf{1}\}$ and $X_j = \{b \in B \mid v_j(b) = \mathbf{1}\}$. Since $v_i(b) = v_j(b)$ for all $b \in B$, it follows that $X_i = X_j$ and since f is a function it is the case that $f(X_i) = f(X_j)$. By Definition 3.8, it follows that $v_i(c) = v_j(c)$.

To see that T satisfies condition (2), see that due to Definition 2.3 it is the case $\mathcal{D}$ is not empty. This means that there exists $X \in \mathcal{D}$. By Definition 3.8, it follows that there exists a valuation $v \in T$. It follows that $T \neq \emptyset$.

To see that T satisfies condition (3), observe that any $X \in \mathcal{D}$ is a legal possibility due to Definition 2.3. By Definition 2.2, for any $X \in \mathcal{D}$, X is informative and hence $X \neq \emptyset$. By Definition 3.8, for any valuation $v \in T$, $\{b \in B \mid v(b) = \mathbf{1}\} \neq \emptyset$. It follows that for any $v \in T$, $\{v\} \models \bigvee_{b \in B} b$. Due to the fact that $\bigvee_{b \in B} b \in \mathbf{PL}$ and by flatness of formulas of $\mathbf{PL}$, $T \models \bigvee_{b \in B} b$.

To prove that T satisfies condition (4), let $b_i \in B$ be arbitrary. By Definition 3.8, $b_i \in \bigcup \mathcal{D}$. This means that there exists $X \in \mathcal{D}$ such that $b_i \in X$. By Definition 3.8, there exists $v \in T$ such that $\{b \in B \mid v(b) = \mathbf{1}\} = X$. Therefore, $\{v\} \models b_i$. Since $\{v\} \subseteq T$ it holds that $T \models \blacklozenge b_i$. Since b_i was arbitrary it follows that for any $b \in B$, $T \models \blacklozenge b$.

To see that T satisfies condition (5), let $b_i, b_j \in B$ be two opposing conditions. By Definition 2.3 it is the case for any $X \in \mathcal{D}$, X is a legal possibility. By Definition 2.2, it follows that X is conflict-free and hence $\{b_i, b_j\} \not\subseteq X$. By Definition 3.8, it is the case that for any valuation $v \in T$, $\{b_i, b_j\} \not\subseteq \{b \in B \mid v(b) = \mathbf{1}\}$. It follows that, for any $v \in T$ it is the case that $\{v\} \not\models (b_i \wedge b_j)$. It follows that $T \models \neg(b_i \wedge b_j)$. Since b_i, b_j were arbitrary, it follows that for each $b_i \in B$, $T \models \bigwedge_{b_j \in B_i^{\perp}} \neg(b_i \wedge b_j)$.

To prove the right-to-left direction, let T be an arbitrary team on $B \cup \{c\}$ where $B \subseteq \mathcal{B}$ is a non-empty set of conditions. Further, suppose that T satisfies the conditions (1) to (5). It is sufficient to show that there exists a concept applicability function $f : \mathcal{D} \to \{\mathbf{1}, \mathbf{0}\}$ such that T translates f.

Let $\mathcal{D}$ be defined as $\mathcal{D} := \{\{b \in B \mid v(b) = \mathbf{1}\} \mid v \in T\}$. Further, let $f : \mathcal{D} \to \{\mathbf{1}, \mathbf{0}\}$ be defined as: $f(X) = \mathbf{1}$ if and only if $v(c) = \mathbf{1}$ where $X = \{b \in B \mid v(b) = \mathbf{1}\}$. It needs to be shown that f is a concept applicability function.

Since $T \models =\!(p_1, \ldots, p_n, c)$, it follows that f is a function. Due to the fact that $T \models \bigvee_{b \in B} b$ and by Definition 2.2, any $X \in \mathcal{D}$ is informative. Further, due to the fact that $T \models \bigwedge_{b_j \in B_i^{\perp}} \neg(b_i \wedge b_j)$ for any $b_i \in B$ and by Definition 2.2, any $X \in \mathcal{D}$ is conflict-free. By Definition 2.2, it follows that every X in $\mathcal{D}$ is a legal possibility. Since $T \neq \emptyset$ it follows that $\mathcal{D} \neq \emptyset$. It follows by Definition 2.3 that f is a concept applicability function.

Lastly, it needs to be shown that T translates f. To do this, let us first show that $\bigcup \mathcal{D} = B$. To show that $\bigcup \mathcal{D} \subseteq B$, let $b_i \in \bigcup \mathcal{D}$ be arbitrary. As $b_i \in \bigcup \mathcal{D}$, there exist $X \in \mathcal{D}$ such that $b_i \in X$. By Definition 3.8, there exists $v \in T$ such that $\{b \in B \mid v(b) = \mathbf{1}\} = X$. Hence, $b_i \in \{b \in B \mid v(b) = \mathbf{1}\}$. Since $\{b \in B \mid v(b) = \mathbf{1}\} \subseteq B$, $b_i \in B$. To show that $B \subseteq \bigcup \mathcal{D}$, let $b_i \in B$ be arbitrary. Since for each $b \in B$ it is the case that $T \models \blacklozenge b$, it follows that $T \models \blacklozenge b_i$. Thus, $b_i \in \{b \in B \mid v(b) = \mathbf{1}\}$ for some $v \in T^c$. By Definition 3.8, $b_i \in X$ for some $X \in \mathcal{D}$. Thus, $b_i \in \bigcup \mathcal{D}$. It follows that $\bigcup \mathcal{D} = B$.

Now, let us show that the conditions (1) and (2) of Definition 3.8 are satisfied.

By the definition of $\mathcal{D}$, there exists a valuation $v \in T$ with truth set $Y_v = \{b \in B \mid v(b) = \mathbf{1}\}$ if and only if there exists a set of conditions $X \in \mathcal{D}$ such that $Y_v = X$. Further, since $\bigcup \mathcal{D} = B$, it holds that the falsity set $Z_v = \{b \in B \mid v(b) = \mathbf{0}\} = B \setminus Y_v = \bigcup \mathcal{D} \setminus X$. Therefore, the first condition of Definition 3.8 is satisfied.

Since f is a function and by its definition, it follows that for each valuation $v \in T$ with truth set

$Y_v = \{b \in B \mid v(b) = \mathbf{1}\}$, it is the case that $v(c) = \mathbf{1}$ if and only if $f(Y_v) = \mathbf{1}$. Therefore, the second condition of Definition 3.8 is satisfied.

By Definition 3.8, it follows that T translates f and hence T is a legal team. □

Let us shortly sketch the conceptual implication of these conditions. By Definition 2.3, a concept applicability function is a function from a non-empty set of legal possibilities to Boolean values. By Definition 2.2, every legal possibility is informative and conflict-free. Constraints (3) and (5) define analogous restrictions for legal teams. Constraint (5) states that no valuation of a legal team specifies that opposing conditions are known to be satisfied. Thus, no legal possibility characterized by a valuation contains two conditions which oppose each other. This highlights that legal reasoning is not based on conflicting information.

Constraint (3) expresses that for each valuation, it is the case that the legal system knows that at least one condition is satisfied. That is, every legal possibility expressed by some valuation v of the legal team is not empty. This illustrates that the applicability of a concept cannot be determined without knowing at least something.

Example 3.12. Let shirt_1 be a specific shirt, and let us consider the concept "possession of shirt_1", expressed by c_1, and the conditions "wearing shirt_1", "not wearing shirt_1", "storing shirt_1 in the wardrobe", and "not storing shirt_1" that are denoted by b_1, b_2, b_3, and b_4, respectively. Table 3.3 defines a team T_1 on $B_1 \cup \{c\}$, where $B_1 = \{b_2, b_3, b_4\}$. Hence, T_1 does not contain any information on b_1, so it does not include the information represented in the light gray column. Additionally, Table 3.3 specifies a team T_2 on the set $B_2 \cup \{c\}$, where $B_2 = \{b_1, b_2, b_3, b_4\}$, that includes information about b_1 and therefore contains the information of the light gray column.

valuation v_i	legal possibility X_i	b_1	b_2	b_3	b_4	concept c_1
v_1	$\{b_1\}$	**1**	**0**	**0**	**0**	**1**
v_2	$\{b_2, b_3\}$	**0**	**1**	**1**	**0**	**1**
v_3	$\{b_2, b_4\}$	**0**	**1**	**0**	**1**	**0**

Table 3.3: Team T_1 and T_2 specifying the applicability of the concept "possession of shirt_1"

T_1 is not a legal team because $T_1 \not\models \bigvee_{b \in B_1} b$. This is due to the fact that v_1 specifies a data point where the applicability of the concept c_1 is determined without knowing that any condition is satisfied. In contrast, this is not the case for T_2 because it considers condition b_1. Therefore, $T_2 \models \bigvee_{b \in B_2} b$. △

Constraint (2) states that any legal team is not empty. In other words, any legal team contains at least one valuation. This ensures that there is at least some information about the applicability of the considered concept. This reflects the condition of a concept applicability function that the domain of the concept applicability function is non-empty. Thus, constraints (2) and (3) encode that the applicability of a legal concept requires some known information about satisfied conditions.

The constraint (1) articulates a functional dependence between the legal possibilities expressed by some valuation and the judgment on the applicability. It states that the value of a concept is given by the values for the conditions so that two identical legal possibilities expressed by some valuations are mapped to the same value for the applicability of the concept by the valuations. This encodes the Aristotelian idea of formal equality, which states that equal cases should be treated equally [Ari20].[5]

[5] The significance of formal equality for legal reasoning has been intensively debated. Westen [Wes90] and Kelsen

If this condition is not satisfied by a team representing some information, then this information characterizes a concept as both applicable and inapplicable to an identical set of conditions that are known to be satisfied. Then, the requirement of formal equality indicates that there must be some conditions that, if known to be satisfied, cause the different judgment on the applicability of the concept. However, this condition is not considered by the team leading to the fact that the constraint (1) is not satisfied. This is exemplified by the following.

Example 3.13. Suppose that we are investigating a team that appears to specify the applicability of the concept "ticketed bike", denoted by c_2, of the Dutch train operator. This concept defines bikes that you are allowed to travel with in trains, but require a specific ticket. The following conditions are considered: "city bike", "folding bike" and "folded bigger than 45 x 86 x 80 cm". These conditions are denoted by b_1, b_2 and b_3, respectively. Further, consider the following table Table 3.4 sketching the team T_1 on $B_1 \cup \{c_2\}$ where $B_1 = \{b_1, b_2\}$. Thus, T_1 lacks any information on b_3 and hence does not consider any information provided in the light gray column. Further, Table 3.4 characterizes the team T_2 on $B_2 \cup \{c_2\}$ where $B_1 = \{b_1, b_2, b_3\}$. Consequently, T_2 contains information on b_3 and hence considers the information of the light gray column.

valuation v_i	legal possibility X_i	b_1	b_2	b_3	concept c_2
v_1	$\{b_1\}$	**1**	**0**	**0**	**1**
v_2	$\{b_1, b_2\}$	**1**	**1**	**0**	**0**
v_3	$\{b_1, b_2, b_3\}$	**1**	**1**	**1**	**1**

Table 3.4: Team T_1 and T_2 specifying the applicability of the concept "ticketed bike"

If a conductor wants to reason about the applicability of c_2 based on T_1, thus using the information characterized by the table without the column for b_3, they realize that this table does not contain sufficient information about the applicability of c_2. Based on the acquired information, a bike is not a "ticketed bike" if it is known to be a city bike and a folding bike due to v_2. However, v_3 conveys the information that a bike is a "ticketed bike" if the same conditions are known to be satisfied. This illustrates that T_1 does not provide information about why the concept is applicable to the latter but not the former. This is reflected by the fact that $T_1 \not\models =\!(b_1, b_2, c_2)$ so that T_1 is not a legal team. When T_2 is considered, these legal possibilities are distinguishable. According to T_2, the difference in applicability is due to the fact that condition b_3 is known to be satisfied in the latter, which is not the case in the former. This exemplifies why $T_2 \models =\!(b_1, b_2, b_3, c_2)$. △

Lastly, constraint (4) ensures that any condition considered by a legal team is contained within a legal possibility expressed by some valuation of the legal team. This captures the fact that any concept applicability function only refers to conditions that are contained within a legal possibility mapped to a Boolean value. Furthermore, if the legal team translates the concept applicability function of a concept, then this condition ensures that the set of ground conditions is identical to the set of conditions considered by the legal team.

Recall that Definition 2.6 defines properties of legal possibilities relative to concept applicability functions, such as sufficiency or positive necessity, in a set-theoretic fashion. Importantly, these properties are expressible using formulas of $\mathbf{PL}(=\!(\cdot), \blacklozenge)$. Proposition 3.14 provides these formulas and

[Kel17; KW24] articulate that formal equality is an insubstantial legal norm without normative effect. In contrast, Alexy [Ale20] and Somek [Som06] argue that formal equality as a norm prescribes equal treatment as the norm and unequal treatment as the derogation. Accordingly, there is a burden of justification for unequal treatment.

shows that they correspond to the relationships of legal possibilities expressed by legal teams.

Proposition 3.14. Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions, $f : \mathcal{D} \to \{\mathbf{1}, \mathbf{0}\}$ be a concept applicability function and T^c be a legal team on $B \cup \{c\}$ that translates f.

(1) A legal possibility $X_i \in \mathcal{D}$ is sufficient relative to f if and only if $T^c \models \varphi_{\mathsf{suf}}^{X_i}$ where

$$\varphi_{\mathsf{suf}}^{X_i} := \bigwedge_{b_i \in X_i} b_i \to c$$

(2) A legal possibility $X_i \in \mathcal{D}$ is an exception relative to f if and only if $T^c \models \varphi_{\mathsf{exc}}^{X_i}$ where

$$\varphi_{\mathsf{exc}}^{X_i} := \bigwedge_{b_i \in X_i} b_i \to \neg c$$

(3) A legal possibility $X_i \in \mathcal{D}$ is positively necessary relative to f if and only if $T^c \models \varphi_{\mathsf{nec+}}^{X_i}$ where

$$\varphi_{\mathsf{nec+}}^{X_i} := c \to \bigwedge_{b_i \in X_i} b_i$$

(4) A legal possibility $X_i \in \mathcal{D}$ is negatively necessary relative to f if and only if $T^c \models \varphi_{\mathsf{nec-}}^{X_i}$ where

$$\varphi_{\mathsf{nec-}}^{X_i} := \neg c \to \bigwedge_{b_i \in X_i} b_i$$

(5) A legal possibility relative to f $X_i \in \mathcal{D}$ is contingent if and only if $T^c \models \varphi_{\mathsf{con}}^{X_i}$ where

$$\varphi_{\mathsf{con}}^{X_i} := \blacklozenge \left(\bigwedge_{b_i \in X_i} b_i \wedge c \right) \wedge \blacklozenge \left(\bigwedge_{b_i \in X_i} b_i \wedge \neg c \right)$$

Proof. Proof is given in Appendix A.1. $\square$

Note that $\varphi_{\mathsf{suf}}^{X_i}, \varphi_{\mathsf{exc}}^{X_i}, \varphi_{\mathsf{nec+}}^{X_i}, \varphi_{\mathsf{nec-}}^{X_i} \in \mathbf{PL}$, while $\varphi_{\mathsf{con}}^{X_i} \in \mathbf{PL}(=(\cdot), \blacklozenge)$. Given a legal team T^c and a legal possibility X_i expressed by T^c, we say that X_i is sufficient relative to T^c if and only if $T^c \models \varphi_{\mathsf{suf}}^{X_i}$. Similarly, we say that X_i is an exception, positively necessary, negatively necessary, or contingent relative to T^c if and only if $T^c \models \varphi_{\mathsf{exc}}^{X_i}$, $T^c \models \varphi_{\mathsf{nec+}}^{X_i}$, $T^c \models \varphi_{\mathsf{nec-}}^{X_i}$, or $T^c \models \varphi_{\mathsf{con}}^{X_i}$, respectively.

If a legal team translates the concept applicability function f^c specifying the applicability of a legal concept c, then these properties characterize properties of decisive legal possibilities. In this setting, a decisive legal possibility is sufficient for c, a decisive legal possibility is an exception for c, and so on.

Since the discussed properties of legal possibilities hold for a legal possibility if and only if the corresponding formula is satisfied by the team, we derive the following results.

Fact 3.15. Let $c \in \mathcal{C}$ be a concept and $B \subseteq \mathcal{B}$ be a non-empty set of conditions. Further, let concept applicability function $f : \mathcal{D} \to \{\mathbf{1}, \mathbf{0}\}$ be concept applicability function and T^c be a legal team on $B \cup \{c\}$ that translates f. Let $X_i \in \mathcal{D}$ be a legal possibility expressed by some valuation $v_i \in T^c$.

(1) If $T^c \models \varphi_{\mathsf{suf}}^{X_i}$, then for any legal possibility $X_j \in \mathcal{D}^c$ expressed by some $v_j \in T^c$ with $v_j(b) = \mathbf{1}$ for all $b \in X_i$ it holds that $T^c \models \varphi_{\mathsf{suf}}^{X_j}$

(2) If $T^c \models \varphi^{X_i}_{\mathsf{exc}}$, then for any legal possibility $X_j \in \mathcal{D}^c$ expressed by some $v_j \in T^c$ with $v_j(b) = \mathbf{1}$ for all $b \in X_i$ it holds that $T^c \models \varphi^{X_j}_{\mathsf{exc}}$

(3) If $T^c \models \varphi^{X_i}_{\mathsf{nec}+}$, then for any legal possibility $X_j \in \mathcal{D}^c$ expressed by some $v_j \in T^c$ with $v_i(b) = \mathbf{1}$ for all $b \in X_j$ it holds that $T^c \models \varphi^{X_j}_{\mathsf{nec}+}$

(4) If $T^c \models \varphi^{X_i}_{\mathsf{nec}-}$, then for any legal possibility $X_j \in \mathcal{D}^c$ expressed by some $v_j \in T^c$ with $v_i(b) = \mathbf{1}$ for all $b \in X_j$ it holds that $T^c \models \varphi^{X_j}_{\mathsf{nec}-}$

(5) $T^c \models \varphi^{X_i}_{\mathsf{con}}$ if and only if $T^c \not\models \varphi^{X_i}_{\mathsf{suf}}$ and $T^c \not\models \varphi^{X_i}_{\mathsf{exc}}$

Proof. (1) to (5) follow immediately from Proposition 3.14 and Fact 2.9. □

To summarize, in this section we showed how $\mathbf{PL}(=\!(\cdot), \blacklozenge)$ allows us to model legal information on the applicability of concepts. This information is modeled using legal teams, which represent concept applicability functions in $\mathbf{PL}(=\!(\cdot), \blacklozenge)$. Like concept applicability functions, legal teams specify whether a concept is applicable to legal possibilities. In Theorem 3.11, we characterized when a team qualifies as a legal team. This logical framework lays the foundation for reasoning about the derivation of the applicability of a concept to a case, which is discussed in the next chapter.

Chapter 4

Derivation of Legal Concept Applicability

Assessing whether a legal concept is applicable to a case based on some known information leads to two scenarios. First, the legal team provides explicit information on the applicability of the concept to the legal case so that the applicability can be directly derived. This is defined in Section 4.2. Second, the legal team does not contain explicit information about the applicability of the concept to the legal case, but it might contain information to predict the concept's applicability. We define a mechanism capable of determining the applicability based on such information in Section 4.3 and investigate its requirements in Section 4.4. Together, these mechanisms accomplish the objective of Task 2. Before formally exploring these mechanisms, we provide an intuition for them in Section 4.1.

4.1 An Intuition about the Derivation of Legal Concept Applicability

Before providing a formal characterization of the proposed mechanisms of direct derivation and heuristic derivation, let us begin with an informal description using an example.

Example 4.1. Let us examine the concept of "tax relief eligibility", denoted by c_3. Let "employed by a Dutch company", "resident in the Netherlands", "student of a Dutch university", and "Dutch civil servant" be conditions, denoted by b_1, b_2, b_3, and b_4, respectively. Further, let $B = \{b_1, b_2, b_3, b_4\}$ and let $B_i^{\perp} = \emptyset$ for all $b_i \in B$ so that none of these conditions oppose one another. Let T^{c_3} on $B \cup \{c_3\}$, as shown in Table 4.1, be the legal team of c_3 and hence characterize the applicability of c_3.

valuation v_i	legal possibility X_i	b_1	b_2	b_3	b_4	concept c_3
v_1	$\{b_1\}$	**1**	**0**	**0**	**0**	**1**
v_2	$\{b_3\}$	**0**	**0**	**1**	**0**	**0**
v_3	$\{b_4\}$	**0**	**0**	**0**	**1**	**1**
v_4	$\{b_1, b_2\}$	**1**	**1**	**0**	**0**	**0**
v_5	$\{b_2, b_3\}$	**0**	**1**	**1**	**0**	**1**

Table 4.1: Legal team T^{c_3} for the concept "tax relief eligibility"

Table 4.1 states that if someone is known to be employed by a Dutch company, they are eligible for tax relief. However, if it is also known that they reside in the Netherlands, they are not eligible. In contrast, if someone is known to be a student at a Dutch university and resides in the Netherlands, they are eligible, but not if it is only known that they are a student at a Dutch university. Lastly, if it is known that someone is a Dutch civil servant, they are also eligible.

Table 4.1 allows us to deduce certain properties of decisive legal possibilities. It is easy to verify that the following are the properties possessed by expressed decisive legal possibilities, excluding contingency:

(1) The decisive legal possibility $\{b_4\}$ is sufficient for c_3. This means that knowing someone is a Dutch civil servant is sufficient to conclude that they are eligible for tax relief.

(2) The decisive legal possibility $\{b_1, b_2\}$ is an exception for c_3. That is, knowing someone is employed by a Dutch company and resides in the Netherlands is sufficient to conclude that they are not eligible for tax relief.

(3) The decisive legal possibility $\{b_2, b_3\}$ is sufficient for c_3. Hence, knowing that someone is a resident in the Netherlands and a student at a Dutch university is sufficient to conclude that the person is eligible for tax relief.

Based on these properties, we want to investigate the applicability of c_3 to several individuals: a_1, a_2, a_3, and a_4, where $K_{a_1} = \{b_2, b_3\}$, $K_{a_2} = \{b_2, b_4\}$, $K_{a_3} = \{b_1, b_2, b_3\}$, and $K_{a_4} = \{b_2\}$.

This means that in the case of a_1, it is known that they are a resident of the Netherlands and a student at a Dutch university, while for a_2, it is known that they are a Dutch civil servant and a resident of the Netherlands. Furthermore, in the case of a_3, it is known that they are employed by a Dutch company, reside in the Netherlands, and are a student at a Dutch university. Lastly, for a_4, it is only known that they are a resident of the Netherlands.

The applicability of c_3 to a_1 can be derived directly. The characterization states that if it is known about someone that they are a resident of the Netherlands and a student at a Dutch university, then they are eligible for tax relief. Since this is exactly what is known about a_1, one can conclude that c_3 is applicable to this case. That is, this is a scenario in which the legal team contains explicit information about the applicability of the concept to the considered case.

In contrast, the legal team does not contain explicit information about the cases a_2, a_3, and a_4—namely, it does not characterize the applicability of the concept to legal possibilities about which the same is known as about these cases. Thus, these cases constitute a scenario where the legal team does not contain explicit information about the applicability of the concept. Therefore, the properties of the legal possibilities must be used to derive applicability. This gives rise to a more fine-grained distinction among these cases.

According to (1), knowing that someone is a Dutch civil servant is sufficient to conclude that they are eligible for tax relief. Since it is known about a_2 that they are a Dutch civil servant, it is reasonable to derive, based on (1), that c_3 is applicable to a_2. Thus, in this situation, we can heuristically derive the applicability of the concept based on the properties of expressed legal possibilities. However, this is not the case for the remaining two cases.

Since it is known about a_3 that they are employed by a Dutch company, reside in the Netherlands, and are a student at a Dutch university, we can conclude based on (2) that a_3 is not eligible for tax relief. In contrast, according to (3), we can conclude that they are eligible for tax relief. This is the case because it is known that a_3 satisfies the conditions of both a legal possibility that is sufficient and an exception for c_3. Therefore, we have conflicting information regarding the applicability of c_3 to a_3, and we are not in a position to heuristically derive its applicability.

Lastly, for a_4, it is only known that they are a resident of the Netherlands. Importantly, this case does not fall under any of the legal possibilities discussed in (1), (2), or (3), as it does not satisfy all the conditions. Consequently, these properties cannot be used to derive the applicability of the concept, and there is no indication as to how to infer the applicability of c_3. In this sense, we do not have enough information to determine applicability. $\triangle$

Importantly, this example indicates that an investigation into the applicability of a concept to a case, based on legal information, can lead to four possible outcomes. First, the applicability may be directly derived from the legal information. Second, while it cannot be directly derived, it may still be heuristically derived. Alternatively, the applicability may be neither directly nor heuristically derivable. This gives rise to two further possibilities: either the legal team contains too little information, or it contains conflicting information for heuristic derivations.

Given a legal team T^c that specifies the applicability of a concept c, we will refer to cases that fall under the first group of possible outcomes as being *determined by* T^c. If a case falls into the second group, we will say that *c is predictable for this case with respect to* T^c. If a case qualifies as belonging to one of the last types of outcomes, it will be said that *c is inconsistently unpredictable* with respect to T^c if the legal team contains conflicting information, or *indeterminably unpredictable* with respect to T^c if the legal team lacks sufficient information about the applicability of the concept to the case.

We will formally define these notions in the following sections, and they will serve as the foundation for the direct and heuristic derivation of legal concept applicability.

4.2 Direct Derivation of Legal Concept Applicability

Direct derivation of the applicability of a concept to a legal case is based on the fact that the legal information explicitly states whether the concept is applicable to that case. That is, the legal information includes whether the legal concept is applicable to a legal possibility that contains the same conditions as the legal knowledge associated with the case. This means that when the applicability of a legal concept to a legal case is decided by direct derivation, the legal team that provides the legal information functions as a database against which the legal case is compared.

Definition 4.2 (Determination and Applicability)**.** Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and let T^c be a legal team on $B \cup \{c\}$. Further, let $a \in \mathcal{A}$ be a legal case and K_a the associated legal knowledge.

- Case a is *determined* by T^c if and only if there exists $v_i \in T^c$ such that $X_i = B \cap K_a$ where X_i is the legal possibility expressed by v_i.
- Concept c is *applicable* to a relative to T^c if and only if there exists $v_i \in T^c$ such that $X_i = B \cap K_a$ and $v_i(c) = \mathbf{1}$ where X_i is the legal possibility expressed by v_i.
- Concept c is *not applicable* to a relative to T^c if and only if there exists $v_i \in T^c$ such that $X_i = B \cap K_a$ and $v_i(c) = \mathbf{0}$ where X_i is the legal possibility expressed by v_i.

Let us highlight two important characteristics about Definition 4.2. First, observe that case a is determined by a legal team T^c if and only if the concept c is applicable to a or not applicable to a relative to T^c. Second, note that "c is not applicable to a relative to T^c" is not the same as "it is not the case that c is applicable to a relative to T^c".

According to Definition 4.2, the conclusion that a legal case is determined by a legal team is based on the fact that the legal team expresses a legal possibility that is identical to the legal knowledge about the case relative to the conditions considered by the legal team. If that holds, then the information stored in the legal team about the applicability of the concept to this legal possibility can be used to directly derive whether the concept is applicable or not applicable to this legal case.

Example 4.3. Let c_1 express "possession of shirt$_1$", b_1 denote wearing shirt$_1$, b_2 denote "not wearing shirt$_1$", b_3 denote "storing shirt$_1$ in the wardrobe", and b_4 denote "not storing shirt$_1$ in the wardrobe". Let $B = \{b_1, b_2, b_3, b_4\}$. Further, let $B_1^\perp = \{b_2, b_3\}, B_2^\perp = \{b_1\}, B_3^\perp = \{b_1, b_4\}$ and $B_4^\perp = \{b_3\}$ define the sets of opposing conditions. The following legal team T^{c_1} on $B \cup \{c_1\}$, given by Table 4.2, specify the applicability of c_1.

valuation v_i	legal possibility X_i	b_1	b_2	b_3	b_4	concept c_1
v_1	$\{b_1\}$	**1**	**0**	**0**	**0**	**1**
v_2	$\{b_3\}$	**0**	**0**	**1**	**0**	**1**
v_3	$\{b_1, b_4\}$	**1**	**0**	**0**	**1**	**1**
v_4	$\{b_2, b_3\}$	**0**	**1**	**1**	**0**	**1**
v_5	$\{b_2, b_4\}$	**0**	**1**	**0**	**1**	**0**

Table 4.2: Legal team T^{c_1} for the concept "possession of shirt$_1$"

In addition, let b_5 denote the condition "wearing hat$_1$" where hat$_1$ is a specific hat. Suppose that b_5 and b_1 do not oppose each other.

Let us consider the legal cases a_1, a_2 and a_3 where $K_{a_1} = \{b_1, b_5\}$, $K_{a_2} = \{b_2, b_4\}$ and $K_{a_3} = \{b_2\}$. This means that it is known that a_1 wears shirt$_1$ and hat$_1$, a_2 neither wears shirt$_1$ nor stores it, and a_3 does not wear shirt$_1$. Case a_1 and a_2 are determined by $T_1^{c_1}$ because $X_1 = \{b_1\} = B \cap K_{a_1}$ and $X_5 = \{b_2, b_4\} = B \cap K_{a_2}$ where X_1 and X_5 are the legal possibilities expressed by v_1 and v_5, respectively. Since $v_1(c_1) = \mathbf{1}$ and $v_5(c_1) = \mathbf{0}$, c_1 is applicable to a_1 and c_1 is not applicable to a_2. In contrast, a_3 is not determined by T^{c_1} because there does not exist a legal possibility X_i expressed by some valuation $v_i \in T^{c_1}$ such that $X_i = B \cap K_{a_3}$. △

The case a_3 illustrates that no conclusions can be drawn about the applicability of a concept to a legal case if the case is not determined by the legal team. This demonstrates that the mechanism for assessing the applicability of a concept is robust in the sense that it only derives conclusions based on information explicitly specified by the legal team. However, it is also limited, as it cannot derive the applicability of a concept beyond what is explicitly provided. Consequently, there is a need for a mechanism that extends beyond the explicitly available information, while still being grounded in it.

4.3 Heuristic Derivation of Legal Concept Applicability

The aim of a heuristic derivation of the applicability of a concept to a case is to go beyond the information explicitly provided about that applicability. Since such a mechanism derives a conclusion about the applicability of a concept that is not contained in the given information, it effectively predicts applicability based on that information.

As indicated in Example 4.1, the heuristic derivation of legal concept applicability is based on properties of legal possibilities, such as sufficiency and negative necessity. Recall that these properties are relative to the provided legal information on the applicability of the concept, which is represented by legal teams. Importantly, these properties might provide essential information about whether applying or not applying this concept to a legal case is sensible.

A legal possibility might be sufficient for the applicability of the concept according to the provided legal information, even though the information might be incomplete—namely, it does not specify whether the concept is applicable to any legal possibility. Predicting the applicability of a concept to a case that is not explicitly considered by the legal information should align with the sufficiency of this legal possibility. That is, it is a reasonable inference that this legal possibility should remain sufficient, and that the information simply did not include all cases.

In this sense, a prediction regarding the applicability of a legal concept to a case must not deviate from the provided information about the properties of legal possibilities.

To do this, we first integrate the legal case into the legal team. This allows us to investigate whether the legal team continues to satisfy the properties of being a legal team, and whether it continues to satisfy the same properties of expressed legal possibilities. This step lays the foundation for predicting whether a legal concept is applicable. A failure to satisfy these properties allows us to draw conclusions about whether the applicability or non-applicability is compatible with the given legal information.

To integrate a legal case into a legal team, a translation of a legal case into a valuation with respect to a concept and a non-empty set of conditions needs to be defined.

Definition 4.4 (Positive Valuation and Negative Valuation)**.** Let $c \in \mathcal{C}$ be a concept and $B \subseteq \mathcal{B}$ be a non-empty set of conditions. Further, let $a \in \mathcal{A}$ be a legal case and K_a the associated legal knowledge. The positive valuation $v^{c+}_{K_a} : B \cup \{c\} \to \{\mathbf{1}, \mathbf{0}\}$ and negative valuation $v^{c-}_{K_a} : B \cup \{c\} \to \{\mathbf{1}, \mathbf{0}\}$ of a with respect to c and B are two functions such that

(i) for all $b_i \in B \cap K_a$, $v^{c+}_{K_a}(b_i) = \mathbf{1}$ and $v^{c-}_{K_a}(b_i) = \mathbf{1}$

(ii) for all $b_j \in B \setminus K_a$, $v^{c+}_{K_a}(b_j) = \mathbf{0}$ and $v^{c-}_{K_a}(b_j) = \mathbf{0}$

(iii) $v^{c+}_{K_a}(c) = \mathbf{1}$ and $v^{c-}_{K_a}(c) = \mathbf{0}$

The positive valuation $v^{c+}_{K_a}$ and the negative valuation $v^{c-}_{K_a}$ of a with respect to a concept c and a non-empty set of conditions translate the information about the conditions that a satisfies and that are ground conditions. Further, the positive valuation assigns the Boolean value **1** to the concept, while the negative valuation assigns the Boolean value **0** to the concept. Thus, the positive valuation captures the case where the concept is applicable to the legal possibility expressed by a and the negative valuation illustrates the case where the concept is not applicable to it. Given these valuations, an extension of a legal team by a legal case can be defined.

Definition 4.5 (Extension of Legal Team)**.** Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and T^c be a legal team on $B \cup \{c\}$. Further, let $a \in \mathcal{A}$ be a legal case and K_a the associated legal knowledge.

(1) A *positive extension* of T^c by a is the team $T = T^c \cup \{v^{c+}_{K_a}\}$ where $v^{c+}_{K_a}$ is the positive valuation of a with respect to concept c and non-empty set of conditions B.

(2) A *negative extension* of T^c by a is the team $T = T^c \cup \{v^{c-}_{K_a}\}$ where $v^{c-}_{K_a}$ is the negative valuation of a with respect to concept c and non-empty set of conditions B.

We say that T is an *extension* of T^c by a if it is a positive extension of T^c by a, or a negative extension of T^c by a. This definition of an extension of a legal team by a case does not guarantee that the extension is itself a legal team. This is because the positive or negative valuation of a case might contain information that conflicts with the properties required for being a legal team. Consider, for instance, the following example.

Example 4.6. Let c_1 express "possession of shirt$_1$", b_1 denote "wearing shirt$_1$", b_2 denote "not wearing shirt$_1$", b_3 denote "storing shirt$_1$ in the wardrobe", b_4 denote "not storing shirt$_1$ in the wardrobe" and $B = \{b_1, b_2, b_3, b_4\}$. Let $B_1^\perp = \{b_2, b_3\}, B_2^\perp = \{b_1\}, B_3^\perp = \{b_1, b_4\}$ and $B_4^\perp = \{b_3\}$ define the sets of opposing conditions. Table 4.3, excluding the last two rows, defines the legal team T^{c_1} on $B \cup \{c_1\}$ that specifies the applicability of c_1.

valuation v_i	legal possibility X_i	b_1	b_2	b_3	b_4	concept c_1
v_1	$\{b_1\}$	**1**	**0**	**0**	**0**	**1**
v_2	$\{b_3\}$	**0**	**0**	**1**	**0**	**1**
v_3	$\{b_1, b_4\}$	**1**	**0**	**0**	**1**	**1**
v_4	$\{b_2, b_3\}$	**0**	**1**	**1**	**0**	**1**
v_5	$\{b_2, b_4\}$	**0**	**1**	**0**	**1**	**0**
$v_{K_{a_4}}^{c_1-}$	$\{b_1\}$	**1**	**0**	**0**	**0**	**0**
$v_{K_{a_5}}^{c_1+}$	$\{b_2\}$	**0**	**1**	**0**	**0**	**1**

Table 4.3: Legal team T^{c_1} for the concept "possession of shirt$_1$" with extensions by a_4 and a_5

Furthermore, consider the legal case a_4 with associated legal knowledge $K_{a_4} = \{b_1\}$. According to Definition 4.4, the negative valuation $v_{K_{a_4}}^{c_1-}$ of a_4 with respect to c_1 and B is defined as: $v_{K_{a_4}}^{c_1-}(b_1) = \mathbf{1}$ and $v_{K_{a_4}}^{c_1-}(c_1) = v_{K_{a_4}}^{c_1-}(b) = \mathbf{0}$ for $b \in \{b_2, b_3, b_4\}$. The negative extension T of T^{c_1} by a_4 is defined by the previous table, including the valuation $v_{K_{a_4}}^{c_1-}$ but excluding the last row.

Observe that $v_1(b) = v_{K_{a_4}}^{c_1-}(b)$ for all $b \in B$, and $v_1(c_1) \neq v_{K_{a_4}}^{c_1-}(c_1)$. Thus, the two valuations assign the same values to the conditions and different values to the concept. This means that $T \not\models =\!(b_1, \ldots, b_4, c_1)$. By Theorem 3.11, it follows that T is not a legal team, and hence the negative extension of T^{c_1} by a_4 is not a legal team.

Now, let us consider the legal case a_5 with associated legal knowledge $K_{a_5} = \{b_2\}$. The positive valuation $v_{K_{a_5}}^{c_1+}$ of a_5 with respect to c_1 is defined as: $v_{K_{a_5}}^{c_1+}(c_1) = v_{K_{a_5}}^{c_1+}(b_2) = \mathbf{1}$ and $v_{K_{a_5}}^{c_1+}(b) = \mathbf{0}$ for $b \in \{b_1, b_3, b_4\}$. Therefore, the positive extension T of T^{c_1} by a_5 is defined by the previous table, including the valuation $v_{K_{a_5}}^{c_1+}$ but excluding $v_{K_{a_4}}^{c_1-}$. It is easy to verify that this positive extension T is a legal team. That is, T satisfies all the conditions of a legal team provided in Theorem 3.11. △

Fact 4.7 indicates some conditions under which an extension of a legal team is a legal team as well.

Fact 4.7. Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and T^c be a legal team on $B \cup \{c\}$. Further, let $a \in \mathcal{A}$ be a legal case and K_a the associated legal knowledge. Lastly, let T_{pos} be the positive extension of T^c by case a and T_{neg} be the negative extension of T^c by case a.

(1) $T_{\mathsf{pos}} \neq \emptyset$, $T_{\mathsf{pos}} \models \blacklozenge b_i$ and $T_{\mathsf{pos}} \models \bigwedge_{b_j \in B_i^\perp} \neg(b_i \wedge b_j)$ for each $b_i \in B$.

(2) $T_{\mathsf{neg}} \neq \emptyset$, $T_{\mathsf{neg}} \models \blacklozenge b_i$ and $T_{\mathsf{neg}} \models \bigwedge_{b_j \in B_i^\perp} \neg(b_i \wedge b_j)$ for each $b_i \in B$.

(3) If $B \cap K_a = \emptyset$, then neither T_{pos} nor T_{neg} is a legal team.

(4) Concept c is applicable to a relative to T^c if and only if T_{pos} is a legal team and T_{neg} is not a legal team.

(5) Concept c is not applicable to a relative to T^c if and only if T_{pos} is not a legal team and T_{neg} is a legal team.

(6) Case a is not determined by T^c and $B \cap K_a \neq \emptyset$ if and only if T_{pos} and T_{neg} are legal teams.

Proof. Proof is given in Appendix A.2. □

Statement (6) sketches the circumstances under which the applicability of a concept to a case is predictable by going beyond what is explicitly stated in the legal team. Such a prediction can occur when it is known that a legal case satisfies some conditions considered by the legal team, and no explicit legal information about the applicability of the concept to this case is given.

Recall that a prediction about the applicability of a concept to a case should not change the properties of the legal possibilities expressed by the legal team. To formalize this idea, a notion of constancy needs to be defined—one which captures the idea that a prediction preserves these specific features of the provided legal information.

Definition 4.8 (Constant Extension of a Legal Team)**.** Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and T^c be a legal team on $B \cup \{c\}$. Further, let $a \in \mathcal{A}$ be a legal case. An extension T of T^c by a is a *constant extension* of T^c if and only if the following conditions are satisfied:

(1) T is a legal team

(2) for any legal possibility X_i expressed by some $v_i \in T^c$, if $T^c \models \psi$, then $T \models \psi$ for $\psi \in \{\varphi^{X_i}_{\mathsf{suf}}, \varphi^{X_i}_{\mathsf{exc}}, \varphi^{X_i}_{\mathsf{nec}+}, \varphi^{X_i}_{\mathsf{nec}-}, \varphi^{X_i}_{\mathsf{con}}\}$

This definition states that any constant extension of a legal team preserves any property of a legal possibility expressed by a valuation of the legal team. This means that the information encoded in a legal team are preserved under this extension.

Fact 4.9. Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and T^c be a legal team on $B \cup \{c\}$. Further, let $a \in \mathcal{A}$ be a legal case, K_a the associated legal knowledge and T be a constant extension of T^c by a. For any legal possibility X_i expressed by some $v_i \in T^c$, if $T \models \psi$, then $T^c \models \psi$ for $\psi \in \{\varphi^{X_i}_{\mathsf{suf}}, \varphi^{X_i}_{\mathsf{exc}}, \varphi^{X_i}_{\mathsf{nec}+}, \varphi^{X_i}_{\mathsf{nec}-}, \varphi^{X_i}_{\mathsf{con}}\}$.

Proof. Proof is given in Appendix A.2. □

Definition 4.8 requires that the properties of legal possibilities are preserved if a legal team is extended to a constant extension. Fact 4.9 investigates the opposite direction. It states that no legal possibility expressed by a valuation of a legal team gains properties in a constant extension. This means that a constant extension preserves information in two senses: no properties of legal possibilities expressed by the legal team are lost or gained through a constant extension.

We can define how to predict the applicability of a concept to a legal case based on this information-preserving constant extension of a legal team.

Definition 4.10 (Prediction of Applicability)**.** Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and T^c be a legal team on $B \cup \{c\}$. Further, let $a \in \mathcal{A}$ be a legal case, T_{pos} be the positive extension of T^c by case a and T_{neg} be the negative extension of T^c by case a.

- Concept c is *predictable* for a with respect to T^c if and only if exactly one of T_{pos} or T_{neg} is a constant extension of T^c, but not both.

- Concept c is *predicted to be applicable* to a with respect to T^c if and only if T_{pos} is a constant extension of T^c, but T_{neg} is not a constant extension of T^c.
- Concept c is *predicted to be not applicable* to a with respect to T^c if and only if T_{neg} is a constant extension of T^c, but T_{pos} is not a constant extension of T^c.
- Concept c is *indeterminably unpredictable* for a with respect to T^c if and only if T_{pos} and T_{neg} are constant extension of T^c.
- Concept c is *inconsistently unpredictable* for a with respect to T^c if and only if neither T_{pos} nor T_{neg} is a constant extension of T^c.

If the relevant legal team T^c is clear from the context, the reference to this legal team is omitted when stating the predictability of the (non-)applicability of c to a legal case. Similarly, it is omitted when stating that c is indeterminably (inconsistently) unpredictable for a case. As an example, "c is predicted to be (not) applicable to a" means that c is predicted to be (not) applicable to a with respect to T^c.

Note that when a concept c is unpredictable for a with respect to a legal team, there are two cases to consider. Either the positive and negative extensions are both constant extensions, or neither the negative nor the positive extension is a constant extension. In the former case, the concept c is indeterminably unpredictable for a with. This means that the legal team does not contain enough information to determine whether the concept should be applied to a because the legal team does not provide enough information to rule out one alternative. On the other hand, In the latter case, the legal team contains conflicting information, so that neither applicability nor non-applicability preserves the properties given in the legal team. In such a case, the concept c is inconsistently unpredictable for a.

Fact 4.11. Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions, T^c be a legal team on $B \cup \{c\}$ and $a \in \mathcal{A}$ be a legal case. If a is determined by T^c, then

(1) c is applicable to a if and only if c is predicted to be applicable to a.

(2) c is not applicable to a if and only if c is predicted to be not applicable to a.

Proof. Proof is given in Appendix A.2. □

Fact 4.11 specifies the relationship between the prediction of the applicability of a concept to a legal case and its applicability to the legal case. Importantly, the prediction of the applicability extends the applicability, which is illustrated by the following insights. If it is given whether a concept is applicable to a case, then the prediction follows this judgment. This means that if the concept is applicable to the case, then it is also predicted that the concept is applicable to the case. Similarly, if the concept is not applicable to the case, then it is also predicted that the concept is not applicable to the case. Further, given that a concept is determined by a legal team, a concept is applicable to a case if it is predicted that the concept is applicable to the case. Similarly, once it is predicted that a concept is not applicable to a case determined by the legal team, the concept is not applicable.

Example 4.12. Let c_1 express "possession of shirt$_1$", b_1 denote "wearing shirt$_1$", b_2 denote "not wearing shirt$_1$", b_3 denote "storing shirt$_1$ in the wardrobe", b_4 denote "not storing shirt$_1$ in the wardrobe" and $B = \{b_1, b_2, b_3, b_4\}$. Let $B_1^\perp = \{b_2, b_3\}, B_2^\perp = \{b_1\}, B_3^\perp = \{b_1, b_4\}$ and $B_4^\perp = \{b_3\}$ define the sets of opposing conditions. The legal team T^{c_1} on $B \cup \{c_1\}$, given by Table 4.4 excluding the last two rows, specifies the applicability of c_1.

valuation v_i	legal possibility X_i	b_1	b_2	b_3	b_4	concept c_1
v_1	$\{b_1\}$	**1**	**0**	**0**	**0**	**1**
v_2	$\{b_3\}$	**0**	**0**	**1**	**0**	**1**
v_4	$\{b_2, b_3\}$	**0**	**1**	**1**	**0**	**1**
v_5	$\{b_2, b_4\}$	**0**	**1**	**0**	**1**	**0**
$v^{c_1+}_{K_{a_6}}$	$\{b_1, b_4\}$	**1**	**0**	**0**	**1**	**1**
$v^{c_1-}_{K_{a_6}}$	$\{b_1, b_4\}$	**1**	**0**	**0**	**1**	**0**

Table 4.4: Legal team T^{c_1} for the concept "possession of shirt$_1$" with extensions by a_6

Note that this specification of the concept c_1 differs from the specification discussed in Examples 4.3 and 4.6 due to the fact that this specification omits v_3. It is easy to verify that $T^{c_1} \models \varphi^{X_1}_{\mathsf{suf}}$, $T^{c_1} \models \varphi^{X_2}_{\mathsf{suf}}$, $T^{c_1} \models \varphi^{X_4}_{\mathsf{suf}}$, $T^{c_1} \models \varphi^{X_5}_{\mathsf{exc}}$ and $T^{c_1} \models \varphi^{X_5}_{\mathsf{nec}-}$. In addition, T^{c_1} does not satisfy any other formula defining a property for any legal possibility expressed by any valuation.

Let a_6 be a legal case that is associated with the legal knowledge $K_{a_6} = \{b_1, b_4\}$. Note that there does exist a valuation $v \in T^{c_1}$ such that $\{b \in B \mid v(b) = \mathbf{1}\} = K_{a_6} \cap B$ where $B = \{b_1, b_2, b_3, b_4\}$. According to Definition 4.2, this means that a_6 is not determined by T^{c_1}.

A positive extension T_{pos} of T^{c_1} by a_6 is defined as $T_{\mathsf{pos}} = T^{c_1} \cup \{v^{c_1+}_{K_{a_6}}\}$, and the negative extension T_{neg} of T^{c_1} by a_6 is defined as $T_{\mathsf{neg}} = T^{c_1} \cup \{v^{c_1-}_{K_{a_6}}\}$. Observe that $B \cap K_{a_6} \neq \emptyset$. Due to this, and the fact that a_6 is not determined by T^{c_1}, it follows by Fact 4.7 that both T_{pos} and T_{neg} are legal teams. Alternatively, one can directly verify that T_{pos} and T_{neg} satisfy the properties of legal teams as specified in Theorem 3.11.

It is easy to verify that $T_{\mathsf{pos}} \models \varphi^{X_1}_{\mathsf{suf}}$, $T_{\mathsf{pos}} \models \varphi^{X_3}_{\mathsf{suf}}$, $T_{\mathsf{pos}} \models \varphi^{X_4}_{\mathsf{suf}}$, $T_{\mathsf{pos}} \models \varphi^{X_5}_{\mathsf{exc}}$ and $T_{\mathsf{pos}} \models \varphi^{X_5}_{\mathsf{nec}-}$. Thus, T_{pos} is a constant extension of T^{c_1}.

Further, note that $T_{\mathsf{neg}} \not\models \neg b_1 \vee c$. Since $\{v^{c_1-}_{K_{a_6}}\} \models b_1$ and $\{v^{c_1-}_{K_{a_6}}\} \models \neg c_1$, there does not exist $T_1, T_2 \subseteq T_{\mathsf{neg}}$ such that $T_1 \cup T_2 = T_{\mathsf{neg}}$, $T_1 \models \neg b_1$ and $T_2 \models c_1$. This means that $T_{\mathsf{neg}} \not\models \bigwedge_{b \in X_1} b \to c$. It follows that $T_{\mathsf{neg}} \not\models \varphi^{X_1}_{\mathsf{suf}}$ which implies that T_{neg} is not a constant extension.

Since T_{pos} is a constant extension and T_{pos} is not a constant extension, it follows due to Definition 4.10 that c_1 is predicted to be applicable to a_6. This means that it is predicted that the individual, who is known to be wearing shirt$_1$ but not storing shirt$_1$ (case a_6), possesses the shirt.

Let us briefly explain why predicting the applicability of c_1 to a_6 is reasonable. According to Table 4.4, the legal possibility $\{b_1\}$ is sufficient for the applicability of c_1. This means that, if it is known that someone wears shirt$_1$, then they possess shirt$_1$. Moreover, it is known that a_6 wears shirt$_1$ and does not store shirt$_1$ in the wardrobe. The legal information does not explicitly address this combination of information. Nevertheless, the information about a_6 contains the legal possibility $\{b_1\}$. Since this legal possibility is sufficient for the applicability of c_1, only the prediction that c_1 is applicable to the case is consistent with the legal information. Thus, this is the only prediction grounded in the obtained legal information, making it a sensible conclusion. △

4.4 Criteria of Heuristic Derivation of Legal Concept Applicability

In this section, we investigate when the applicability of a legal concept can be predicted with respect to a legal team. Therefore, we examine what kind of information a legal team needs to articulate so that such a prediction can be established.

Intuitively, the legal team must contain some information in order to be able to predict whether a concept is applicable. This information is necessary so that the concept is not indeterminably unpredictable for a case. Therefore, we need to establish a lower bound of required information.

However, the legal team cannot contain just any amount of information to surpass this lower bound. That is, it must be ensured that the legal team does not contain information that renders a prediction of the applicability impossible in the sense that the concept becomes inconsistently unpredictable.

To characterize the information a legal team needs to possess in order to establish that a concept is predicted to be applicable (not applicable) to a case, we proceed as follows. Firstly, we will establish a general result about the satisfaction of properties expressed by a legal team if this legal team is extended. This will be done in Proposition 4.13. Based on these results, we will characterize the circumstances under which a concept is unpredictable for a case with respect to a legal team. That is, we will secondly examine under which conditions a concept is inconsistently unpredictable for a legal case. This is addressed in Lemma 4.14. Thirdly, we will investigate under which conditions a concept is indeterminably unpredictable for a legal case. This is done in Lemma 4.15. Based on these results, we can establish the requirements of predictability and hence the requirements of heuristic derivations in Theorem 4.16. We will illustrate these results in Example 4.17.

Proposition 4.13. Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions, T^c be a legal team on $B \cup \{c\}$ and X a legal possibility expressed by some valuation of T^c. Further, let $a \in \mathcal{A}$ be a legal case and K_a the associated legal knowledge. For any $\psi \in \{\varphi^X_{\mathsf{suf}}, \varphi^X_{\mathsf{exc}}, \varphi^X_{\mathsf{nec}+}, \varphi^X_{\mathsf{nec}-}, \varphi^X_{\mathsf{con}}\}$,

(1a) $\underset{\text{pos}}{\checkmark}$ denotes that if $T^c \models \psi$, then $T_{\mathsf{pos}} \models \psi$ (1b) $\underset{\text{neg}}{\checkmark}$ denotes that if $T^c \models \psi$, then $T_{\mathsf{neg}} \models \psi$

(2a) $\underset{\text{pos}}{\text{✗}}$ denotes that if $T^c \models \psi$, then $T_{\mathsf{pos}} \not\models \psi$ (2b) $\underset{\text{neg}}{\text{✗}}$ denotes that if $T^c \models \psi$, then $T_{\mathsf{neg}} \not\models \psi$

where $T_{\mathsf{pos}} = T^c \cup \{v^{c+}_{K_a}\}$ is the positive extension of T^c by a, $T_{\mathsf{neg}} = T^c \cup \{v^{c-}_{K_a}\}$ is the negative extension of T^c by a and X is a legal possibility expressed by some $v \in T^c$.

Table 4.5 specifies whether $\psi \in \{\varphi^X_{\mathsf{suf}}, \varphi^X_{\mathsf{exc}}, \varphi^X_{\mathsf{nec}+}, \varphi^X_{\mathsf{nec}-}, \varphi^X_{\mathsf{con}}\}$ remains satisfied by the positive or negative extension of T^c by a, depending on whether $X \subseteq B \cap K_a$ or $X \not\subseteq B \cap K_a$, assuming that $T^c \models \psi$.

	$\psi := \varphi^X_{\mathsf{suf}}$		$\psi := \varphi^X_{\mathsf{exc}}$		$\psi := \varphi^X_{\mathsf{nec}+}$		$\psi := \varphi^X_{\mathsf{nec}-}$		$\psi := \varphi^X_{\mathsf{con}}$	
$X \subseteq B \cap K_a$	✓ pos	✗ neg	✗ pos	✓ neg	✓ pos	✓ neg	✓ pos	✓ neg	✓ pos	✓ neg
$X \not\subseteq B \cap K_a$	✓ pos	✓ neg	✓ pos	✓ neg	✗ pos	✓ neg	✓ pos	✗ neg	✓ pos	✓ neg

Table 4.5: Preservation of satisfaction for ψ after extending by a

Proof. Proof is given in Appendix A.2. □

Proposition 4.13 analyzes whether a certain property of a legal possibility—expressed by some valuation—remains satisfied when a legal team is extended by a legal case. The table distinguishes between whether the legal possibility is contained within the legal knowledge associated with the legal case or not. Assuming the legal team satisfies a given property about the legal possibility, and depending on whether the legal possibility is contained in the case's legal knowledge, each row indicates whether the positive or negative extension by the legal case also satisfies this property. Importantly, any case where either the positive or negative extension fails to satisfy the property is marked by a cross (✗).

This result allows us to characterize the conditions under which a concept is inconsistently unpredictable.

Lemma 4.14 (Inconsistently Unpredictable)**.** Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and T^c be a legal team on $B \cup \{c\}$. Further, let $a \in \mathcal{A}$ be a legal case and K_a the associated legal knowledge such that a is not determined by T^c. The concept c is inconsistently unpredictable for a with respect to T^c if and only if one of the following conditions is satisfied:

(1) $T^c \models \varphi^{X_i}_{\mathsf{suf}}$ and $T^c \models \varphi^{X_j}_{\mathsf{exc}}$ where X_i and X_j are legal possibilities expressed by some valuations $v_i, v_j \in T^c$ such that $X_i \subseteq B \cap K_a$ and $X_j \subseteq B \cap K_a$

(2) $T^c \models \varphi^{Y_i}_{\mathsf{nec}+}$ and $T^c \models \varphi^{Y_j}_{\mathsf{nec}-}$ where Y_i and Y_j are legal possibilities expressed by some valuations $v_i, v_j \in T^c$ such that $Y_i \nsubseteq B \cap K_a$ and $Y_j \nsubseteq B \cap K_a$

(3) $B \cap K_a = \emptyset$

Proof. Let $T_{\mathsf{pos}} = T^c \cup \{v^{c+}_{K_a}\}$ be the the positive extension of T^c by legal case a and $T_{\mathsf{neg}} = T^c \cup \{v^{c-}_{K_a}\}$ be the the negative extension of T^c by legal case a.

For the left-to-right direction, suppose that c is inconsistently unpredictable for a with respect to T^c. This means that neither T_{pos} nor T_{neg} is a constant extension. Since a is not determined by T^c and due to Definition 4.8 and Fact 4.7, either (a) $B \cap K_a = \emptyset$ whence by Fact 4.7 (3) the extensions are not legal teams or (b) $B \cap K_a \neq \emptyset$ whence by Fact 4.7 (6), the extensions are legal teams, but then because the extensions are not constant, for each extension there exists one formula $\psi \in \{\varphi^X_{\mathsf{suf}}, \varphi^X_{\mathsf{exc}}, \varphi^X_{\mathsf{nec}+}, \varphi^X_{\mathsf{nec}-}, \varphi^X_{\mathsf{con}}\}$ such that $T^c \models \psi$ and the the extension does not satisfy ψ where X is a legal possibility expressed by some $v \in T^c$.

If (a) is the case, then condition (3) is satisfied.

If (b) is the case, then there exists ψ_1 such that $T^c \models \psi_1$ and $T_{\mathsf{pos}} \not\models \psi_1$ where $\psi_1 \in \{\varphi^{X_i}_{\mathsf{suf}}, \varphi^{X_i}_{\mathsf{exc}}, \varphi^{X_i}_{\mathsf{nec}+}, \varphi^{X_i}_{\mathsf{nec}-}, \varphi^{X_i}_{\mathsf{con}}\}$ and X_i is the legal possibility expressed by some $v_i \in T^c$. Further, there exists ψ_2 such that $T^c \models \psi_2$ and $T_{\mathsf{neg}} \not\models \psi_2$ where $\psi_2 \in \{\varphi^{X_j}_{\mathsf{suf}}, \varphi^{X_j}_{\mathsf{exc}}, \varphi^{X_j}_{\mathsf{nec}+}, \varphi^{X_j}_{\mathsf{nec}-}, \varphi^{X_j}_{\mathsf{con}}\}$ and X_j is the legal possibility expressed by some $v_j \in T^c$. It needs to be shown that (1) $\psi_1 := \varphi^{X_i}_{\mathsf{exc}}$ and $\psi_2 := \varphi^{X_j}_{\mathsf{suf}}$ where $X_i, X_j \subseteq B \cap K_a$ or (2) $\psi_1 := \varphi^{Y_i}_{\mathsf{nec}+}$ and $\psi_2 := \varphi^{Y_j}_{\mathsf{nec}-}$ where $Y_i, Y_j \nsubseteq B \cap K_a$.

By Proposition 4.13, it follows that $\psi_1 := \varphi^{X_i}_{\mathsf{exc}}$ where $X_i \subseteq B \cap K_a$ or $\psi_1 := \varphi^{Y_i}_{\mathsf{nec}+}$ where $Y_i \nsubseteq B \cap K_a$. Further, $\psi_2 := \varphi^{X_j}_{\mathsf{suf}}$ where $X_j \subseteq B \cap K_a$ or $\psi_2 := \varphi^{Y_j}_{\mathsf{nec}-}$ where $Y_j \nsubseteq B \cap K_a$.

Note that it cannot hold that $\psi_1 := \varphi^{X_i}_{\mathsf{exc}}$ and $\psi_2 := \varphi^{Y_j}_{\mathsf{nec}-}$ where $X_i \subseteq B \cap K_a$ and $Y_j \nsubseteq B \cap K_a$. Similarly, it is not possible that $\psi_1 := \varphi^{Y_i}_{\mathsf{nec}+}$ and $\psi_2 := \varphi^{X_j}_{\mathsf{suf}}$ where $Y_i \nsubseteq B \cap K_a$ and $X_j \subseteq B \cap K_a$. This is established in Fact A.1. In contrast, it is possible that $\psi_1 := \varphi^{X_i}_{\mathsf{exc}}$ and $\psi_2 := \varphi^{X_j}_{\mathsf{suf}}$ where $X_i, X_j \subseteq B \cap K_a$. Likewise, one can have $\psi_1 := \varphi^{Y_i}_{\mathsf{nec}+}$ and $\psi_2 := \varphi^{Y_j}_{\mathsf{nec}-}$ where $Y_i, Y_j \nsubseteq B \cap K_a$. This is shown in Fact A.2.

Therefore, it follows that $\psi_1 := \varphi^{X_i}_{\mathsf{exc}}$ and $\psi_2 := \varphi^{X_j}_{\mathsf{suf}}$ where $X_i, X_j \subseteq B \cap K_a$, or $\psi_1 := \varphi^{Y_i}_{\mathsf{nec}+}$ and $\psi_2 := \varphi^{Y_j}_{\mathsf{nec}-}$ where $Y_i, Y_j \nsubseteq B \cap K_a$. Hence, if (b) holds, then condition (1) or (2) is satisfied.

For the right-to-left direction, suppose that one of the conditions (1), (2) or (3) holds.

First, suppose that $T^c \models \varphi^{X_i}_{\mathsf{suf}}$ and $T^c \models \varphi^{X_j}_{\mathsf{exc}}$ where X_i and X_j are legal possibilities expressed by some valuations $v_i, v_j \in T^c$ such that $X_i, X_j \subseteq B \cap K_a$. It needs to be shown that neither T_{pos}

nor T_{neg} is a constant extension. Since $X_i \subseteq B \cap K_a$ and due to Proposition 4.13, it follows that $T_{\mathsf{neg}} \not\models \varphi^{X_i}_{\mathsf{suf}}$. Analogously, since $X_j \subseteq B \cap K_a$ and due to Proposition 4.13, it follows that $T_{\mathsf{pos}} \not\models \varphi^{X_j}_{\mathsf{exc}}$. By Definition 4.8, neither T_{pos} nor T_{neg} is a constant extension.

Second, suppose that $T^c \models \varphi^{Y_i}_{\mathsf{nec}+}$ and $T^c \models \varphi^{Y_j}_{\mathsf{nec}-}$ where Y_i and Y_j are legal possibilities expressed by some valuations $v_i, v_j \in T^c$ such that $Y_i, Y_j \not\subseteq B \cap K_a$. It needs to be shown that neither T_{pos} nor T_{neg} is a constant extension. Since $Y_i \not\subseteq B \cap K_a$ and due to Proposition 4.13, it follows that $T_{\mathsf{pos}} \not\models \varphi^{Y_i}_{\mathsf{nec}+}$. Analogously, since $Y_j \not\subseteq B \cap K_a$ and due to Proposition 4.13, it follows that $T_{\mathsf{neg}} \not\models \varphi^{Y_j}_{\mathsf{nec}-}$. By Definition 4.8, neither T_{pos} nor T_{neg} is a constant extension.

Last, suppose that $B \cap K_a = \emptyset$. Therefore, $\{v^{c+}_{K_a}\} \not\models \bigvee_{b \in B} b$ and $\{v^{c-}_{K_a}\} \not\models \bigvee_{b \in B} b$. By Theorem 3.11, it follows that T_{neg} and T_{pos} are not legal teams. By Definition 4.8, neither T_{pos} nor T_{neg} is a constant extension. $\square$

Lemma 4.14 states that there are three scenarios in which a concept is inconsistently unpredictable for a legal case. The first scenario is characterized by a legal team that contains two legal possibilities such that one legal possibility is sufficient for the concept and the other legal possibility is an exception for the concept, and the legal case satisfies the conditions given by both legal possibilities. As a consequence, a positive or negative extension by the legal case results in a violation of the fact that one legal possibility is an exception or of the fact that the other is sufficient. An example of this scenario is discussed in Example 4.17.

The second scenario specifies that the legal team contains two legal possibilities such that one is positively necessary, the other is negatively necessary and the legal case does not satisfy the conditions of either of these legal possibilities. This implies that any extension of the legal team by the case invalidates the necessity of one of the legal possibilities—either positively or negatively—for the concept.

The third scenario specifies that it is not known that a case satisfies any condition that is considered in the legal team. If this is the case, then any extension by this case is not a legal team and hence not a constant extension.

Next, let us examine the situations where a concept is indeterminably unpredictable.

Lemma 4.15 (Indeterminably Unpredictable)**.** Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and T^c be a legal team on $B \cup \{c\}$. Further, let $a \in \mathcal{A}$ be a legal case and K_a the associated legal knowledge such that a is not determined by T^c. The concept c is indeterminably unpredictable for a with respect to T^c if and only if the following conditions are satisfied:

(1) $T^c \not\models \varphi^{X_i}_{\mathsf{suf}}$ and $T^c \not\models \varphi^{X_i}_{\mathsf{exc}}$ for any legal possibility expressed by some $v_i \in T^c$ with $X_i \subseteq B \cap K_a$

(2) $T^c \not\models \varphi^{Y_j}_{\mathsf{nec}+}$ and $T^c \not\models \varphi^{Y_j}_{\mathsf{nec}-}$ for any legal possibility expressed by some $v_j \in T^c$ with $Y_j \not\subseteq B \cap K_a$

Proof. Let $T_{\mathsf{pos}} = T^c \cup \{v^{c+}_{K_a}\}$ be the the positive extension of T^c by legal case a and $T_{\mathsf{neg}} = T^c \cup \{v^{c-}_{K_a}\}$ be the the negative extension of T^c by legal case a.

For the left-to-right direction, suppose that c is indeterminably unpredictable for a. This means that T_{pos} and T_{neg} are constant extensions. By Definition 4.8, it follows that for all $\psi \in \{\varphi^X_{\mathsf{suf}}, \varphi^X_{\mathsf{exc}}, \varphi^X_{\mathsf{nec}+}, \varphi^X_{\mathsf{nec}-}, \varphi^X_{\mathsf{con}}\}$ with $T^c \models \psi$, it is the case that $T_{\mathsf{pos}} \models \psi$ where X is the legal possibility expressed by some $v \in T^c$. By Proposition 4.13, it follows that for any legal possibility X_i expressed by some $v_i \in T^c$ with $X_i \subseteq B \cap K_a$, $T^c \not\models \varphi^{X_i}_{\mathsf{exc}}$. Further, for any legal possibility Y_j expressed by some $v_j \in T^c$ with $Y_j \not\subseteq B \cap K_a$, $T^c \not\models \varphi^{Y_j}_{\mathsf{nec}+}$.

Since T_{neg} is constant and by Definition 4.8, it follows that for all $\psi \in \{\varphi^X_{\mathsf{suf}}, \varphi^X_{\mathsf{exc}}, \varphi^X_{\mathsf{nec+}}, \varphi^X_{\mathsf{nec-}}, \varphi^X_{\mathsf{con}}\}$ with $T^c \models \psi$, it is the case that $T_{\mathsf{neg}} \models \psi$ where X is the legal possibility expressed by some $v \in T^c$. By Proposition 4.13, it follows that for any legal possibility X_i expressed by some $v_i \in T^c$ with $X_i \subseteq B \cap K_a$, $T^c \not\models \varphi^{X_i}_{\mathsf{suf}}$. Further, for any legal possibility Y_j expressed by some $v_j \in T^c$ with $Y_j \not\subseteq B \cap K_a$, $T^c \not\models \varphi^{Y_j}_{\mathsf{nec-}}$. Hence, condition (1) and (2) are satisfied.

For the right-to-left direction, suppose that condition (1) and (2) are satisfied. By Definition 4.8 and Proposition 4.13, it follows that T_{pos} and T_{neg} are constant extensions. This means that c is indeterminably unpredictable for a. □

Lemma 4.15 characterizes when a legal team contains too little information to predict the applicability of a concept to a case. Example 4.17 provides an example of such a situation. To be able to predict the applicability of a concept, it needs to be known that the case satisfies a sufficient legal possibility or a legal possibility that is an exception. Alternatively, it needs to be known that the case does not satisfy a positively or negatively necessary legal possibility.

Based on Lemmas 4.14 and 4.15, we can determine when a concept is predictable for a legal case with respect to the properties of the legal team and the legal case.

Theorem 4.16 (Predictability of a Concept)**.** Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and T^c be a legal team on $B \cup \{c\}$. Further, let $a \in \mathcal{A}$ be a legal case and K_a the associated legal knowledge such that a is not determined by T^c. The concept c is predictable for a with respect to T^c if and only if the following three conditions are satisfied:

(1) The legal case contains enough information with respect to c which means that the following is satisfied: $B \cap K_a \neq \emptyset$

(2) The legal team contains enough information with respect to a which means that one of the following conditions is satisfied:

(2a) For some legal possibility X_i expressed by some $v_i \in T^c$ such that $X_i \subseteq B \cap K_a$, $T^c \models \varphi^{X_i}_{\mathsf{suf}}$

(2b) For some legal possibility X_i expressed by some $v_i \in T^c$ such that $X_i \subseteq B \cap K_a$, $T^c \models \varphi^{X_i}_{\mathsf{exc}}$

(2c) For some legal possibility Y_j expressed by some $v_j \in T^c$ such that $Y_j \not\subseteq B \cap K_a$, $T^c \models \varphi^{Y_j}_{\mathsf{nec+}}$

(2d) For some legal possibility Y_j expressed by some $v_j \in T^c$ such that $Y_j \not\subseteq B \cap K_a$, $T^c \models \varphi^{Y_j}_{\mathsf{nec-}}$

(3) The legal team does not contain conflicting information with respect to a which means that none of the following conditions are satisfied:

(3a) $T^c \models \varphi^{X_i}_{\mathsf{suf}}$ and $T^c \models \varphi^{X_j}_{\mathsf{exc}}$ where X_i and X_j are legal possibilities expressed by some valuations $v_i, v_j \in T^c$ such that $X_i \subseteq B \cap K_a$ and $X_j \subseteq B \cap K_a$

(3b) $T^c \models \varphi^{Y_i}_{\mathsf{nec+}}$ and $T^c \models \varphi^{Y_j}_{\mathsf{nec-}}$ where Y_i and Y_j are legal possibilities expressed by some valuations $v_i, v_j \in T^c$ such that $Y_i \not\subseteq B \cap K_a$ and $Y_j \not\subseteq B \cap K_a$

Proof. Follows immediately from Lemmas 4.14 and 4.15 and Definition 4.10. □

This theorem specifies the previously established negative results about the predictability of concepts in a positive way by defining three conditions that a legal team has to fulfill in order for the concept to be predictable for a legal case. The first condition demands that it be known that a case satisfies some ground conditions of the legal team. If this does not hold, then the concept is inconsistently unpredictable for the case, according to Lemma 4.14. While Lemma 4.15 examines when too little information is given, the second condition establishes positively how much information is needed to be able to predict the applicability of a concept. The third condition specifies the circumstances in which a legal team has conflicting information about the predictability of the applicability of a concept, as established in Lemma 4.14.

Example 4.17. As before, let the concept "tax relief eligibility" be denoted by c_3. Let "employed by a Dutch company", "resident in the Netherlands", "student of a Dutch university", and "Dutch civil servant" be conditions, denoted by b_1, b_2, b_3, and b_4, respectively. Let $B = \{b_1, b_2, b_3, b_4\}$. Note that none of these conditions oppose one another, so that $B_i^{\perp} = \emptyset$ for all $b_i \in B$. Let the legal team T^{c_3} on $B \cup \{c_3\}$, as shown in Table 4.6, characterize the applicability of c_3.

valuation v_i	legal possibility X_i	b_1	b_2	b_3	b_4	concept c_3
v_1	$\{b_1\}$	**1**	**0**	**0**	**0**	**1**
v_2	$\{b_3\}$	**0**	**0**	**1**	**0**	**0**
v_3	$\{b_4\}$	**0**	**0**	**0**	**1**	**1**
v_4	$\{b_1, b_2\}$	**1**	**1**	**0**	**0**	**0**
v_5	$\{b_2, b_3\}$	**0**	**1**	**1**	**0**	**1**

Table 4.6: Legal team T^{c_3} for the concept "tax relief eligibility"

As illustrated in Example 4.1, the expressed legal possibilities satisfy the following properties, excluding contingency: $T^{c_3} \models \varphi^{X_3}_{\mathsf{suf}}$, $T^{c_3} \models \varphi^{X_4}_{\mathsf{exc}}$, and $T^{c_3} \models \varphi^{X_5}_{\mathsf{suf}}$. As before, let us investigate the applicability of c_3 to several individuals: a_1, a_2, a_3, and a_4, where $K_{a_1} = \{b_2, b_3\}$, $K_{a_2} = \{b_2, b_4\}$, $K_{a_3} = \{b_1, b_2, b_3\}$, and $K_{a_4} = \{b_2\}$.

Observe that a_1 is determined by T^{c_3} because $\{b \in B \mid v_5(b) = \mathbf{1}\} = B \cap K_{a_1}$. This case is not interesting to exemplify the previous results, which are concerned with legal cases that are not determined by a legal team. Nevertheless, it is easy to verify that a_2, a_3, and a_4 are not determined by T^{c_3}.

Concept c_3 is not indeterminably unpredictable for a_2. This is the case because the legal possibility X_3 expressed by v_3 is sufficient for c_3 and $X_3 \subseteq B \cap K_{a_2}$. By Lemma 4.15, this means that c_3 is not indeterminably unpredictable for a_2. Similarly, c_3 is not inconsistently unpredictable for a_2. This is the case because there does not exist another expressed legal possibility X_j such that X_j is an exception for c_3 and $X_j \subseteq B \cap K_{a_2}$. Further, it holds that there do not exist expressed legal possibilities Y_i and Y_j such that the former is positively necessary, the latter negatively necessary for c_3, and $Y_i, Y_j \not\subseteq B \cap K_{a_2}$. By Lemma 4.14, this implies that c_3 is not inconsistently unpredictable for a_2. Therefore, c_3 is predictable for a_2 due to Theorem 4.16. In fact, it is easy to validate that c_3 is predicted to be applicable to a_2.

In contrast, c_3 is inconsistently unpredictable for a_3. Observe that for the expressed legal possibilities X_4, which is an exception for c_3, and X_5, which is sufficient for c_3, it holds that $X_4, X_5 \subseteq B \cap K_{a_3}$. According to Lemma 4.14 and Theorem 4.16, this implies that c_3 is inconsistently unpredictable for a_3 and hence not predictable for a_3.

Further, c_3 is indeterminably unpredictable for a_4. Observe that there does not exist any legal possibility X_i such that it is sufficient or an exception for c_3 and $X_i \subseteq B \cap K_{a_4}$. Similarly, there does not exist any legal possibility Y_i such that it is positively or negatively necessary for c_3 and $Y_i \not\subseteq B \cap K_{a_4}$. According to Lemma 4.15 and Theorem 4.16, this implies that c_3 is indeterminably unpredictable for a_4 and hence not predictable for a_4. △

Chapter 5

Tasks in the Formal Analysis of Legal Concept Applicability

In this chapter, we investigate how the proposed formalism and the mechanisms for directly and heuristically deriving the applicability of legal concepts to cases can address the remaining tasks introduced in Section 1.3. We proceed as follows. First, we define and examine instability in applicability derivations in Section 5.1. Second, we explore how a sequence of cases can be predicted with respect to a legal team in Section 5.2. Lastly, we discuss the conditions relevant to the applicability of legal concepts, and formally characterize them in Section 5.3. Therefore, these sections address Tasks 3, 4 and 5, respectively.

5.1 Stability of Derived Legal Concept Applicability

Given a prediction that a concept is or is not applicable to a legal case based on a legal team specifying the applicability of that concept, one might wonder whether the prediction might be reversed if more information about the legal case is acquired. Accordingly, it is important to investigate whether the prediction remains stable as more information is acquired—this addresses Task 3.

We introduce two notions of stability: *stability* and *strong stability*. Stability refers to the circumstance in which no acquired information can lead to a reversed prediction of the concept's applicability. That is, given that a concept is predicted to be applicable to a legal case, no further information about the legal case can lead to the conclusion that the concept is predicted not to be applicable to that legal case. In contrast, strong stability means that no acquired information can change either the predictability or the prediction of the applicability of the concept. This means that if further information is acquired about the legal case, then neither can the prediction of applicability be reversed, nor can the concept become unpredictable for the legal case.

To formally define these notions, we will formalize notions of instability. Similarly to stability, we introduce two notions of instability: *strong* and *weak instability*. Strong instability refers to the reversal of the prediction of a concept's applicability if further information is acquired, while weak instability denotes the reversal of predictability if further information is acquired. Hence, weak instability means that further information can result in a concept becoming unpredictable for a legal case after more information about the legal case is acquired.

Definition 5.1 (Instability)**.** Let $c \in \mathcal{C}$ be a legal concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and T^c be a legal team on $B \cup \{c\}$. Further, let $a_i \in \mathcal{A}$ be a case such that K_{a_i} is legal knowledge and c is predictable for a_i with respect to T^c.

(1) The prediction of c for a_i with respect to T^c is *strongly unstable* if and only if there exists $a_j \in \mathcal{A}$ such that K_{a_j} is legal knowledge, $B \cap K_{a_i} \subseteq B \cap K_{a_j}$ and one of the following is satisfied:

  (1a) c is predicted to be applicable to a_i and c is predicted to be not applicable to a_j with respect to T^c

  (1b) c is predicted to be not applicable to a_i and c is predicted to be applicable to a_j with respect to T^c

(2) The prediction of c for a_i with respect to T^c is *weakly unstable* if and only if there exists $a_j \in \mathcal{A}$ such that K_{a_j} is legal knowledge and $B \cap K_{a_i} \subseteq B \cap K_{a_j}$ and c is not predictable for a_j with respect to T^c.

We will say that a prediction of concept c for a case a such that K_a is legal knowledge is *strongly stable* if and only if the prediction is neither strongly unstable nor weakly unstable. Further, we will say that a prediction of concept c for a legal case a is *stable* if and only if the prediction is not strongly unstable. Hence, a stable prediction might be weakly unstable. Further, it is easy to see that any strongly stable prediction is a stable prediction. To exemplify these notions, let us consider the following example.

Example 5.2. Let us consider the concept "ticketed bike", denoted by c_2. Further, let b_1 refer to the condition "city bike", b_2 refer to the condition "folding bike", and b_3 refer to the condition "folded bigger than 45 × 86 × 80 cm". Let $B = \{b_1, b_2, b_3\}$ and Let $B_i^{\perp} = \emptyset$ for all $b_i \in B$. Further, the following legal team T^{c_2} on $B \cup \{c_2\}$ given by Table 5.1.

valuation v_i	legal possibility X_i	b_1	b_2	b_3	concept c_2
v_1	$\{b_1\}$	**1**	**0**	**0**	**1**
v_2	$\{b_2\}$	**0**	**1**	**0**	**0**
v_3	$\{b_1, b_2\}$	**1**	**1**	**0**	**0**
v_4	$\{b_2, b_3\}$	**0**	**1**	**1**	**1**

Table 5.1: Legal team T^{c_2} for the concept "ticketed bike"

Let us consider the case a_1, where the legal information is $K_{a_1} = \{b_1\}$. Hence, a_1 represents a bike that is known to be a city bike. It is easy to verify that the concept c_2 is applicable to a_1 and is thus predicted to be applicable to a_1. Suppose that we learn that a_1 satisfies b_2 as well so that $K'_{a_1} = \{b_1, b_2\}$. It is straightforward to validate that in this situation the concept c_2 is not applicable to a_1.

Hence, if the legal system acquires the information that a_1 is a folding bike, then it is not a "ticketed bike". Thus, the prediction that c_2 is applicable to a_1 is strongly unstable and therefore not stable.

Further, let us consider the case a_2, where the legal information is $K_{a_2} = \{b_2, b_3\}$. Observe that c_2 is applicable to a_2, so that c_2 is predicted to be applicable to a_2. This prediction is weakly unstable, but not strongly unstable. To see this, realize that the only additional information the legal system could acquire that would affect the prediction in this context is whether a_2 also satisfies b_1. Thus, suppose that the legal system acquires the information that a_2 is a city bike, so that $K'_{a_2} = \{b_1, b_2, b_3\}$. It is easy to verify that, in this case, c_2 is inconsistently unpredictable for a_2 due to the fact that $T^{c_2} \models \varphi^{X_3}_{\mathsf{exc}}$, $T^{c_2} \models \varphi^{X_4}_{\mathsf{suf}}$ and $X_3, X_4 \subseteq B \cap K'_{a_2}$.

Therefore, if the legal system acquires the information that a_2 is a folding bike, then c_2 is no longer predictable for a_2. This means that the prediction that c_2 is applicable to a_2 is weakly unstable, but not strongly unstable. Thus, the prediction is stable, though not strongly stable. △

Further, let us investigate some properties of strong and weak instability.

Fact 5.3. Let $c \in \mathcal{C}$ be a legal concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and T^c be a legal team on $B \cup \{c\}$. Further, let $a \in \mathcal{A}$ be a case such that K_a is legal knowledge and c is predictable for a with respect to T^c.

(1) If the prediction that c is applicable to a is strongly unstable, then there does not exist a legal possibility X expressed by some $v \in T^c$ such that $T^c \models \varphi^X_{\mathsf{suf}}$ and $X \subseteq B \cap K_a$.

(2) If the prediction that c is not applicable to a is strongly unstable, then there does not exist a legal possibility X expressed by some $v \in T^c$ such that $T^c \models \varphi^X_{\mathsf{exc}}$ and $X \subseteq B \cap K_a$.

Proof. Proof is given in Appendix A.3. □

Fact 5.3 states that if a prediction about the applicability of a concept to a case is strongly unstable, then the legal information about the case is not a superset of any legal possibility—expressed by some valuation of the legal team—that is either sufficient or an exception. In other words, if it is known that the case satisfies a legal possibility that qualifies as sufficient or an exception, then any prediction about the concept's applicability to the case cannot be reversed.

Building on this result, the Proposition 5.4 provides further insights about the situations in which a prediction about the applicability of a concept to a case can be reversed.

Proposition 5.4. Let $c \in \mathcal{C}$ be a legal concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and T^c be a legal team on $B \cup \{c\}$. Further, let $a \in \mathcal{A}$ be a case such that K_a is legal knowledge and c is predictable for a with respect to T^c.

(1) If $T^c \models \blacklozenge \left(\bigwedge_{b \in B \cap K_a} b \wedge \neg c\right)$, then the prediction that c is applicable to a is strongly unstable.

(2) If the prediction that c is applicable to a is strongly unstable and $T^c \not\models \blacklozenge \left(\bigwedge_{b \in B \cap K_a} b \wedge \neg c\right)$, then legal possibilities Y_i and X_j exist, expressed by some valuations $v_i, v_j \in T^c$, such that $T^c \models \varphi^{Y_i}_{\mathsf{nec}-} \wedge \varphi^{X_j}_{\mathsf{exc}}$, $Y_i \nsubseteq (B \cap K_a)$ and $Y_i \cup X_j \cup (B \cap K_a)$ is conflict-free.

(3) If $T^c \models \blacklozenge \left(\bigwedge_{b \in B \cap K_a} b \wedge c\right)$, then the prediction that c is not applicable to a is strongly unstable.

(4) If the prediction that c is not applicable to a is strongly unstable and $T^c \not\models \blacklozenge \left(\bigwedge_{b \in B \cap K_a} b \wedge c\right)$, then legal possibilities Y_i and X_j exist, expressed by some valuations $v_i, v_j \in T^c$, such that $T^c \models \varphi^{Y_i}_{\mathsf{nec}+} \wedge \varphi^{X_j}_{\mathsf{suf}}$, $Y_i \nsubseteq (B \cap K_a)$, and $Y_i \cup X_j \cup (B \cap K_a)$ is conflict-free.

Proof. Proof is given in Appendix A.3. □

To illustrate these results, we consider the prediction that a concept is applicable to a case according to a legal team. Statement (1) states that this prediction can be reversed if the legal team includes a legal possibility in which the concept is not applicable, and this possibility contains at least as much information as the legal system has about the case. Example 5.2 illustrates such a situation.

Importantly, this does not mean that every strongly unstable prediction of applicability requires the presence of such a legal possibility. In fact, some strongly unstable predictions arise even when the legal team does not include such a legal possibility.

Statement (2) characterizes these situations. It states that if the prediction can be reversed, but the legal team does not provide explicit information about a legal possibility that contains the information about the considered case and to which the concept is not applicable, then the following holds. First, the case does not meet the conditions of a negatively necessary legal possibility. Second, the information known to the legal system about the case conflicts neither with the conditions of this negatively necessary legal possibility nor with those of another legal possibility that constitutes an exception. For an illustration of such a situation, consider Example 5.5.

The statements (3) and (4) articulate analogous results for predictions that a concept is not applicable to a case according to a legal team.

Example 5.5. Let us again consider the concept "ticketed bike" and the conditions "city bike", "folding bike", "child bike", and "electric bike". These are denoted by c_2, b_1, b_2, b_3, and b_4, where $B = \{b_1, b_2, b_3, b_4\}$. None of these conditions oppose each other, so that $B_i^{\perp} = \emptyset$ for all $b_i \in B$. Further, let legal team T^{c_2}, given by Table 5.2, characterize the applicability of the concept.

valuation v_i	legal possibility X_i	b_1	b_2	b_3	b_4	concept c_2
v_1	$\{b_2\}$	**0**	**1**	**0**	**0**	**0**
v_2	$\{b_4\}$	**0**	**0**	**0**	**1**	**1**
v_3	$\{b_1, b_3\}$	**1**	**0**	**1**	**0**	**1**

Table 5.2: Legal team T^{c_2} of concept "ticketed bike"

According to this characterization of ticketed bikes, bikes known to be folding bikes are not "ticketed bikes", while any bike known to be a city bike for children is a "ticketed bike". Furthermore, a bike that is known to be an electric bike is a "ticketed bike" as well.

Furthermore, it is easy to verify that $T^{c_2} \models \varphi^{X_1}_{\mathsf{nec}-} \wedge \varphi^{X_1}_{\mathsf{exc}} \wedge \varphi^{X_2}_{\mathsf{suf}} \wedge \varphi^{X_3}_{\mathsf{suf}}$ and that no legal possibility expressed by T^{c_2} satisfies any other property.

Consider the following case a_1 with the associated knowledge $K_{a_1} = \{b_1\}$. Clearly, a_1 is not determined by T^{c_2}. According to Fact 4.7, this means that the positive and negative extensions of T^{c_2} by a_1 are legal teams. The negative extension of T^{c_2} by a_1 is not constant, due to the fact that $T^{c_2} \cup \{v^{c_2-}_{K_{a_1}}\} \not\models \varphi^{X_1}_{\mathsf{nec}-}$. This follows from the fact that $X_1 \not\subseteq B \cap K_{a_1}$ and by Proposition 4.13. In contrast, the positive extension of T^{c_2} by a_1 is constant because of Proposition 4.13 and the fact that $X_1, X_2, X_3 \not\subseteq B \cap K_{a_1}$. It follows that c_2 is predicted to be applicable to a_1 with respect to T^{c_2}.

In addition, $T^{c_2} \not\models \blacklozenge(b_1 \wedge \neg c_2)$. This means that T^{c_2} does not provide explicit information about a bike that is known to be a city-bike and does not qualify as a "ticketed bike". Nevertheless, the prediction that c_2 is applicable to a_1 is strongly unstable. To see this, consider the legal case a_2 with associated legal knowledge $K_{a_2} = \{b_1, b_2\}$.

As before, a_2 is not determined by T^{c_2}, and Fact 4.7 implies that the positive and negative extensions of T^{c_2} by a_2 are legal teams. However, the positive extension of T^{c_2} by a_2 is not constant because $T^{c_2} \cup \{v^{c_2+}_{K_{a_2}}\} \not\models \varphi^{X_1}_{\mathsf{exc}}$. This is due to the fact that $X_1 \subseteq B \cap K_{a_2}$ and by Proposition 4.13. In contrast, the negative extension of T^{c_2} by a_2 is constant because of Proposition 4.13 and the fact that $X_1 \subseteq B \cap K_{a_2}$ and $X_2, X_3 \not\subseteq B \cap K_{a_2}$. It follows that c_2 is predicted to be not applicable to a_2 with respect to T^{c_2}.

This exemplifies that there is some information one might acquire about a_1—namely, that it satisfies condition b_2—which leads to updated legal knowledge K_{a_2}. This updated knowledge results in a case where the concept is predicted to be not applicable, even though the legal team does not provide explicit information how to handle legal possibilities like K_{a_2}. $\triangle$

Further, let us examine weak instability. Since weak instability refers to unpredictability, the characterization of unpredictability, given by Theorem 4.16, can be used to specify this notion.

Proposition 5.6. Let $c \in \mathcal{C}$ be a legal concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and T^c be a legal team on $B \cup \{c\}$. Further, let $a \in \mathcal{A}$ be a case such that K_a is legal knowledge and c is predictable for a with respect to T^c. The prediction that a is applicable (not applicable) to a is weakly unstable if and only if there exists $a_i \in \mathcal{A}$ such that K_{a_i} is legal knowledge and $B \cap K_a \subseteq B \cap K_{a_i}$ and further that one of the following conditions is satisfied:

(1) c is indeterminably unpredictable for a_i

(2) $T^c \models \varphi^{X_i}_{\mathsf{suf}} \wedge \varphi^{X_j}_{\mathsf{exc}}$ and $X_i, X_j \subseteq B \cap K_{a_i}$ where X_i and X_j are legal possibilities expressed by some $v_i, v_j \in T^c$

(3) $T^c \models \varphi^{Y_i}_{\mathsf{nec+}} \wedge \varphi^{Y_j}_{\mathsf{nec-}}$ and $Y_i, Y_j \nsubseteq B \cap K_{a_i}$ where Y_i and Y_j are legal possibilities expressed by some $v_i, v_j \in T^c$

Proof. Proof is given in Appendix A.3. $\square$

If this characterization is compared to Theorem 4.16 and Lemma 4.14, one might observe that it refers only to two conditions under which c is inconsistently unpredictable for a. These conditions are given by statements (2) and (3). However, it does not address inconsistent unpredictability arising from the fact that $B \cap K_{a_i} = \emptyset$. This situation occurs when the legal system does not know that any of the ground conditions are satisfied by a. This possibility is ruled out because a is predictable.[6]

5.2 Sequential Derivations of Legal Concept Applicability

Previously, only the decision whether a concept is predicted to be applicable (not applicable) to one legal case was investigated. To overcome this restriction and to tackle Task 4, it needs to be investigated how the prediction of the applicability of a concept to a sequence of cases can be determined, where a sequence of legal cases is defined as follows.

Definition 5.7 (Sequence of Legal Cases)**.** Let $A \subseteq \mathcal{A}$ be a non-empty set of legal cases. A sequence of A is a bijective function $seq : \{n \in \mathbb{N}_{>0} \mid n \leq |A|\} \rightarrow A$.

For a sequence seq of A, $seq(i)$ denotes the ith legal case of this sequence. As a sequence is a bijective function, each case occurs exactly once in the sequence.

To determine the applicability of a concept to a sequence, a mechanism that sequentially determines whether a concept is predicted to be applicable or not applicable for each case is needed. Conceptually, there are two distinct approaches: either the underlying legal team, specifying the conditions under which a concept is applicable, is updated after the predictability is decided for a case or the legal team is not updated. The former means that a previously determined prediction of the applicability of a concept to a case should influence the prediction of the applicability of another case whereas the latter rejects this. Due to this characteristic, the latter is referred to as static sequential predictions and the former as dynamic sequential predictions. Formally, these are defined as the following.

[6] This is discussed in the proof of Proposition 5.6.

Definition 5.8 (Static Sequential Predictions)**.** Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and T^c be a legal B team on $B \cup \{c\}$. Further, let $A \subseteq \mathcal{A}$ be a non-empty set of legal cases and $seq : \{n \in \mathbb{N}_{>0} \mid n \leq |A|\} \to A$ be a sequence of A. For each case $seq(i) = a_i$ of the sequence seq of A, *c is statically sequentially predicted to be applicable (not applicable) to a_i* for the sequence seq of legal cases A with respect to T^c if and only if c is predicted to be applicable (not applicable) to a_i with respect to T^c.

This formalizes the intuition that any prediction of the applicability of a concept to a case in the finite sequence is independent of the decisions made previously and is exemplified by the following example.

Example 5.9. Let us investigate a situation where it needs to be determined whether several individuals possess shirt$_1$ and the concept "possession of shirt$_1$" is denoted by c_1. Further, b_1 denotes the condition "wearing shirt$_1$", b_2 denotes the condition "not wearing shirt$_1$", b_3 denotes the condition "storing shirt$_1$ in the wardrobe", and b_4 denotes the condition "not storing shirt$_1$ in the wardrobe" so that $B = \{b_1, b_2, b_3, b_4\}$. Let $B_1^{\perp} = \{b_2, b_3\}, B_2^{\perp} = \{b_1\}, B_3^{\perp} = \{b_1, b_4\}$ and $B_4^{\perp} = \{b_3\}$ define the sets of opposing conditions. Let the legal team T^{c_1}, given by Table 5.3, specify the applicability of c_1.

valuation v_i	legal possibility X_i	b_1	b_2	b_3	b_4	concept c_1
v_1	$\{b_1\}$	**1**	**0**	**0**	**0**	**1**
v_2	$\{b_3\}$	**0**	**0**	**1**	**0**	**1**
v_3	$\{b_1, b_4\}$	**1**	**0**	**0**	**1**	**1**
v_4	$\{b_2, b_3\}$	**0**	**1**	**1**	**0**	**1**
v_5	$\{b_2, b_4\}$	**0**	**1**	**0**	**1**	**0**

Table 5.3: Legal team T^{c_1} of concept "possession of shirt$_1$"

Given this information, the applicability of the concept needs to be determined for a sequence seq on the following set of legal cases $A = \{a_1, a_2\}$ where $K_{a_1} = \{b_1, b_4\}$ and $K_{a_2} = \{b_2, b_3\}$. Let $seq : \{1, 2\} \to \{a_1, a_2\}$ be defined by $seq(1) = a_1$ and $seq(2) = a_2$. Observe that a_1 and a_2 are determined by T^{c_1} and that the concept is applicable to both cases due to v_3 and v_4. According to Fact 4.11, c_1 is predicted to be applicable to a_1 and a_2 with respect to T^{c_1}. It follows that c_1 is statically sequentially predicted to be applicable to a_1 and a_2 with respect to T^{c_1}. △

This example illustrates that static sequential predictions for a sequence of legal cases operate as independent checks that determine whether or not the concept is predicted to be applicable to these cases with respect to the underlying legal team.

Since dynamic sequential predictions are based on the idea of updating the legal team to consider previously made predictions about the applicability of the concept, this updating needs to be defined.

Definition 5.10 (Sequential Update)**.** Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and T^c be a legal team on $B \cup \{c\}$. Further, let $A \subseteq \mathcal{A}$ be a non-empty set of legal cases, K_{a_i} the associated legal knowledge for each $a_i \in A$ and $seq : \{n \in \mathbb{N}_{>0} \mid n \leq |A|\} \to A$ be a sequence of A. The *sequential updates* of T^c by the sequence seq is defined inductively:

(1) $T_0^c = T^c$

(2) $$T_i^c = \begin{cases} T_{i-1}^c \cup \{v_{K_{a_i}}^{c+}\}, & \text{if } c \text{ is predicted to be applicable to } seq(i) = a_i \text{ with respect to } T_{i-1}^c \\ T_{i-1}^c \cup \{v_{K_{a_i}}^{c-}\}, & \text{if } c \text{ is predicted to be not applicable to } seq(i) = a_i \text{ with respect to } T_{i-1}^c \\ T_{i-1}^c, & \text{otherwise} \end{cases}$$

Hence, a sequential update results in a positive (negative) extension of the legal team by the legal case, if the concept is predicted to be applicable (not applicable) to the case. Otherwise, the sequential update does not modify the legal team to be updated. These updates provide sufficient means to introduce dynamic sequential predictions.

Definition 5.11 (Dynamic Sequential Predictions)**.** Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and T^c be a legal team on $B \cup \{c\}$. Further, let $A \subseteq \mathcal{A}$ be a non-empty set of legal cases, $seq : \{n \in \mathbb{N}_{>0} \mid n \leq |A|\} \to A$ be a sequence of A and $T^c_0, \ldots, T^c_n$ be the sequential updates of T^c by seq. For each case $seq(i) = a_i$ of the sequence seq of A, c is *dynamically sequentially predicted to be applicable (not applicable) to* a_i for the sequence seq of A with respect to T^c if and only if c is predicted to be applicable (not applicable) to a_i according to T^c_{i-1}.

In contrast to static sequential predictions, the sequentially dynamic predictions of the applicability are not with respect to the same legal team. After each prediction, this legal team might change. Namely, the positive or negative valuation of the legal case, over which the prediction was made, might be added to the legal team. Example 5.15 illustrates how dynamic sequential predictions operate.

Since the underlying legal team might be updated after each prediction, one might wonder whether the predictions depend on the order of the legal cases in the sequence. Formally, we define order-independence as follows:

Definition 5.12 (Order-Independence)**.** Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions, T^c be a legal team on $B \cup \{c\}$ and $A \subseteq \mathcal{A}$ be a non-empty set of legal cases. Static (dynamic) sequential predictions are *order-independent* if and only if for any sequences seq_i and seq_j of A and for all $a \in A$, c is statically (dynamically) sequentially predicted to be applicable (not applicable) to a for the sequence seq_i of A with respect to T^c if and only if c is statically (dynamically) sequentially predicted to be applicable (not applicable) to a for the sequence seq_j of A with respect to T^c.

Based on this definition, the following proposition illustrates that static sequential predictions are order-independent for any set of conditions, set of legal cases, and legal teams. However, dynamic sequential predictions are not order-independent. Importantly, there exist some sets of conditions, sets of legal cases, and legal teams where the order of cases for which the applicability of the concept needs to be determined will affect the predictions.

Proposition 5.13. Let $c \in \mathcal{C}$ be a concept.

(1) Static sequential predictions are order-independent for any non-empty set of conditions $B \subseteq \mathcal{B}$, any legal team on $B \cup \{c\}$, and any non-empty set of legal cases $A \subseteq \mathcal{A}$.

(2) Dynamic sequential predictions are not order-independent, for some non-empty set of conditions $B \subseteq \mathcal{B}$, a legal team on $B \cup \{c\}$, and some non-empty set of legal cases $A \subseteq \mathcal{A}$.

Proof. Proof is given in Appendix A.3. □

According to Proposition 5.13, the order of predicting the applicability of a concept to a legal case matters when the applicability is dynamically sequentially predicted. Importantly, the impact of the sequence does not only affect legal cases with distinct legal knowledge relative to the conditions considered by the case. Since multiple cases can be associated to the same legal knowledge, the applicability of a concept to cases associated to the same legal knowledge might need to be determined in a sequence of cases. This is characterized by the following fact.

Fact 5.14. Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and T^c be a legal team on $B \cup \{c\}$. Further, let $A \subseteq \mathcal{A}$ be a non-empty set of legal cases and $seq : \{n \in \mathbb{N}_{>0} \mid n \leq |A|\} \to A$ be a sequence of A. For any two legal cases $a_i, a_j \in A$ such that $B \cap K_{a_i} = B \cap K_{a_j}$ and $seq^{-1}(a_i) < seq^{-1}(a_j)$, the following holds:

(1) The concept c is statically sequentially predicted to be applicable (not applicable) to a_i for the sequence seq with respect to T^c if and only if c is statically sequentially predicted to be applicable (not applicable) to a_j for the sequence seq with respect to T^c.

(2) If concept c is dynamically sequentially predicted to be applicable (not applicable) to a_i for the sequence seq with respect to T^c, then concept c is dynamically sequentially predicted to be applicable (not applicable) to a_j for the sequence seq with respect to T^c.

Proof. Follows immediately from Definition 5.8 and Definition 5.11. □

Observe how the left-to-right and right-to-left direction for static sequential predictions hold, whereas the right-to-left direction does not hold for dynamic sequential predictions. This entails an important property of dynamic sequential prediction. Even if a legal case might not be sequentially dynamically predictable, another legal case associated to the same legal knowledge (relative to the ground conditions of the legal team) as the first case might be dynamically sequentially predictable. Hence, dynamic sequential predictions can lead to the fact that a case becomes predictable. This is characterized by the following example.

Example 5.15. Let us consider a more fine-grained analysis of the concept "ticketed bike", denoted by c_2. Let b_1 denote the condition "city bike", b_2 denote the condition "folding bike", b_3 denote the condition "folded bigger than 45 cm wide x 86 cm long x 80 cm high", b_4 denote the condition "not folded" and $B = \{b_1, b_2, b_3, b_4\}$. Let $B_i^{\perp} = \emptyset$ for all $b_i \in B$. Further, consider the following legal team T_2^c on $B \cup \{c_2\}$, as specified in Table 5.4, excluding the last two rows.

valuation v_i	legal possibility X_i	b_1	b_2	b_3	b_4	concept c_2
v_1	$\{b_1\}$	**1**	**0**	**0**	**0**	**1**
v_2	$\{b_1, b_2\}$	**1**	**1**	**0**	**0**	**0**
v_3	$\{b_3, b_4\}$	**0**	**0**	**1**	**1**	**1**
v_4	$\{b_1, b_2, b_4\}$	**1**	**1**	**0**	**1**	**1**
$v_5 = v^{c_2+}_{K_{a_2}}$	$\{b_1, b_3\}$	**1**	**0**	**1**	**0**	**1**
$v_6 = v^{c_2+}_{K_{a_3}}$	$\{b_1, b_2, b_3\}$	**1**	**1**	**1**	**0**	**1**

Table 5.4: Legal team T^{c_2} for the concept "ticketed bike" with extensions by a_2 and a_3

This means that if a bike is known to be a city bike, then "ticketed bike" is applicable to the bike. If a bike is known to be not folded and folded larger than 45 cm wide x 86 cm long x 80 cm high, then "ticketed bike" is applicable to this bike as well. Alternatively, if a bike is known to be a folding city bike that is folded bigger than 45 cm wide x 86 cm long x 80 cm high, then it is a "ticketed bike" too. However, if it is only known that a bike is a folding city bike, then "ticketed bike" is not applicable.

Furthermore, it is easy to verify that $T^{c_2} \models \varphi^{X_1}_{\mathsf{nec-}} \wedge \varphi^{X_2}_{\mathsf{nec-}} \wedge \varphi^{X_3}_{\mathsf{suf}} \wedge \varphi^{X_4}_{\mathsf{suf}}$ and that no legal possibility expressed by T^{c_2} satisfies any other property besides contingency.

Let a_1, a_2 and a_3, where $K_{a_1} = K_{a_3} = \{b_1, b_2, b_3\}$ and $K_{a_2} = \{b_1, b_3\}$, be legal cases. Further, consider the following sequence of these cases: $seq : \{1, 2, 3\} \to A$ defined by $seq(1) = a_1, seq(2) = a_2$ and

$seq(3) = a_3$. Thus, it must be dynamically predicted sequentially whether the concept applies to a bike known to be a city bike that is folded larger than 45 cm wide x 86 cm long x 80 cm high (a_2), and to bikes known to be folding city bikes that are folded larger than 45 cm wide x 86 cm long x 80 cm high (a_1 and a_3).

The dynamically sequential prediction of the applicability of c_2 to a_1 is determined with respect to T^{c_2}. It is clearly the case that a_1 is not determined by T^{c_2}. Based on Fact 4.7, it follows that the positive and negative extensions of T^{c_2} by a_1 are legal teams. Since $X_1, X_2 \subseteq B \cap K_{a_1}$ and $X_3, X_4 \nsubseteq B \cap K_{a_1}$, it follows due to Proposition 4.13 that the positive and the negative extension are constant extensions. Thus, c_2 is indeterminably unpredictable for a_1.

Consequently, the applicability of c_2 is then dynamically sequentially predicted to a_2 with respect to $T_1^{c_2} = T^{c_2}$. The positive and negative extensions of $T_1^{c_2}$ by a_2 are legal teams due to Fact 4.7. Since $X_2 \nsubseteq B \cap K_{a_2}$ and by Proposition 4.13, it holds that $T_1^{c_2} \cup \{v_{K_{a_2}}^{c_2-}\} \not\models \varphi_{\mathsf{nec-}}^{X_2}$. This means that the negative extension of $T_1^{c_2}$ by a_2 is not constant. In contrast, it is easy to confirm that the positive extension of $T_1^{c_2}$ by a_2 is constant. Thus, c_2 is predicted to be applicable to a_2 with respect to T_1^{c}. This means that c_2 is dynamically sequentially predicted to be applicable to a_2.

Further, this means that the applicability of c_2 to a_3 is dynamically sequentially predicted with respect to $T_2^{c_2} = T^{c_2} \cup \{v_{K_{a_2}}^{c_2+}\}$. This legal team is specified by Table 5.4, excluding only the last row. This legal team further specifies that a bike known to be a city bike and folded to dimensions larger than 45 cm wide $\times$ 86 cm long $\times$ 80 cm high is a "ticketed bike". In addition, $T_2^{c_2} \models \varphi_{\mathsf{nec-}}^{X_1} \wedge \varphi_{\mathsf{nec-}}^{X_2} \wedge \varphi_{\mathsf{suf}}^{X_3} \wedge \varphi_{\mathsf{suf}}^{X_4}$ and $T_2^{c_2} \models \varphi_{\mathsf{suf}}^{X_5}$.

Again, the positive and negative extensions of $T_2^{c_2}$ by a_3 are legal teams due to Fact 4.7. Further, the negative extension of $T_2^{c_2}$ by a_3 is not constant, due to the fact that $T_2^{c_2} \cup \{v_{K_{a_3}}^{c_2-}\} \not\models \varphi_{\mathsf{suf}}^{X_5}$. This is a conclusion of Proposition 4.13 and the fact that $X_5 \subseteq B \cap K_{a_3}$. In contrast, it can be easily confirmed that the positive extension of $T_2^{c_2}$ by a_3 is constant. It follows that c_2 is predicted to be applicable to a_3 with respect to $T_2^{c_2}$. This means that c_2 is sequentially dynamically predicted to be applicable to a_3, inducing the updated legal team $T_3^{c_2}$. Legal team $T_3^{c_2}$ is specified by Table 5.4 including the last two rows.

This legal team specifies that bikes known to be folding city bikes, which when folded are bigger than 45 cm wide x 86 cm long x 80 cm, are a "ticketed bike". This differs from the specification of the applicability of c_2 by $T^{c_2} \cup \{v_{K_{a_2}}^{c_2+}\}$ in that the folded measurements now refer to folding city bikes and not to city bikes in general. $\triangle$

5.3 Relevant Conditions of Legal Concept Applicability

Given a legal team and its ground conditions, one might wonder whether every ground condition is used to distinguish between the legal possibilities to which the concept is applicable and those to which it is not. Thus, we aim to investigate Task 5. To illustrate this task, let us consider an example.

Example 5.16. Let us assume that Nederlandse Spoorwegen only considers the following two meaning postulates for the concept "ticketed bike":

(1) A city bike is a "ticketed bike".

(2) A *folding* city bike is not a "ticketed bike".

These meaning postulates refer to the conditions "city bike" and "folding bike". However, only the second condition allows to distinguish between the legal possibility to which the concept is applicable and the legal possibility to which it is not applicable. That is, the concept is not applicable to legal possibilities about which it is known that they satisfy the condition of being a "folding bike", while it is applicable to legal possibilities about which this is not known. △

This leads to an important conceptual insight about legal reasoning. Even though meaning postulates might refer to several conditions, not every condition is capable of differentiating the legal possibilities with respect to concept applicability. Conditions that are capable of such differentiations of legal possibilities are *applicability-relevant* conditions.

To identify the applicability-relevant conditions, it is sufficient to identify a function that determines the applicability of the concept based on a minimal set of ground conditions. This is the case because the information about the satisfaction of applicability-relevant conditions differentiates between the legal possibilities to which the concept is applicable and those to which it is not applicable, so that these conditions constitute a minimal set of ground conditions that determines the applicability of the concept to these legal possibilities.

To simplify the search for this function, we do not aim to find the function itself, but rather to identify the existence of a functional dependency that underlies it. Hence, the purpose of determining the applicability-relevant conditions within a legal team is to find the smallest set of conditions that establish such a dependency.

For this purpose, we build on the work of Väänänen [Vää07a; Vää07b] and adapt dependency notions introduced by Väänänen.[7] These notions are based on dependency sets, defined as follows: a set of conditions $X = \{b_1, \ldots, b_n\}$ is a dependency set of a concept c with respect to a team T if and only if $T \models =\!(b_1, \ldots, b_n, c)$.

Definition 5.17 (Applicability-Relevance)**.** Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and T^c be a legal team on $B \cup \{c\}$.

(1) $X \subseteq B$ is *applicability-relevant* for c relative to T^c if and only if there is some minimal (with respect to set-inclusion) dependency set Y of c with respect to T^c such that $X \subseteq Y$.

(2) $X \subseteq B$ is *applicability-irrelevant* for c relative to T^c if and only if X is not applicability-relevant for c relative to T^c.

(3) $X \subseteq B$ is *totally applicability-relevant* for c relative to T^c if and only if X is applicability-relevant for c relative to T^c and for every dependency set Y of c with respect to T^c it holds that $X \subseteq Y$.

A condition b is said to be applicability-relevant for c if $\{b\}$ is applicability-relevant for c. This applies similarly to the notions of applicability-irrelevant and totally applicability-relevant.

According to Definition 5.17, determining the applicability-relevant sets of conditions focuses on identifying the minimal sets that functionally determine the applicability of the concept under consideration. Any subset of such a set is applicability-relevant to that concept and might be among the conditions that guide legal decisions regarding the concept's applicability. Knowing that such a condition is satisfied might influence whether the concept is applicable or not, and therefore distinguishes the legal possibilities to which the concept is applicable from those to which it is not. If a condition is totally applicability-relevant, then it is certain that knowledge about the satisfaction of this condition affects

[7] Our notion of applicability-relevant conditions corresponds to Väänänen's notions of dependence.

the applicability of the concept. Thus, it is one of the conditions that differentiates between the legal possibilities to which the concept is applicable and those to which it is not. In contrast, if a condition is applicability-irrelevant, then knowledge of this condition does not affect the applicability of the concept and cannot help to differentiate legal possibilities.

Example 5.18. Let us consider the concept "ticketed bike", denoted by c_2. Further, let b_1 refer to the condition "city bike", b_2 refer to the condition "folding bike", b_3 refer to the condition "child bike" and b_4 refer to the condition "folded bigger than 45 x 86 x 80 cm". Let $B = \{b_1, b_2, b_3, b_4\}$ and $B_i^{\perp} = \emptyset$ for all $b_i \in B$. Based on the legal team T_2^c given by Table 5.5, the aim is to investigate the applicability-relevant conditions.

valuation v_i	legal possibility X_i	b_1	b_2	b_3	b_4	concept c_2
v_1	$\{b_1, b_3\}$	**1**	**0**	**1**	**0**	**1**
v_2	$\{b_1, b_2, b_4\}$	**1**	**1**	**0**	**1**	**1**
v_3	$\{b_1, b_2, b_3\}$	**1**	**1**	**1**	**0**	**0**

Table 5.5: Legal team T^{c_2} for the concept "ticketed bike"

Accordingly, if it is known that a city bike is a folding bike and folded bigger than 45 x 86 x 80 cm, then it is a "ticketed bike". If it is known that a bike is a city-bike for children, then it is a "ticketed bike", but if it is known that a bike is a *folding* city-bike for children, it is not a "ticketed bike".

Figure 5.1 characterizes the dependency sets for c_2 relative to T^{c_2} and specifies the inclusion order.

$$\begin{tikzcd}
 & & \{b_1, b_2, b_3, b_4\} & & \\
\{b_1, b_2, b_3\} \arrow[urr] & & \{b_2, b_3, b_4\} \arrow[u] & & \{b_1, b_2, b_4\} \arrow[ull] \\
 & \{b_2, b_3\} \arrow[ul] \arrow[ur] & & \{b_2, b_4\} \arrow[ul] \arrow[ur] &
\end{tikzcd}$$

Figure 5.1: Set inclusion order on the dependency sets of the concept "ticketed bike"

No other set of conditions is a dependency set for c_2 relative to T^{c_2}. Condition b_2 is applicability-relevant and totally applicability-relevant for c_2, because it is contained in any dependency set, including the minimal dependency sets $\{b_2, b_3\}$ and $\{b_2, b_4\}$. The conditions b_3 and b_4 are applicability-relevant for c_2, but not totally applicability-relevant, as they are only contained in one minimal dependency set and not in both. Lastly, the condition b_1 is applicability-irrelevant for c_2 because it is not contained in any minimal dependency set.

This means that the question of whether the bike in question is a folding bike has, first, an impact on determining whether this bike is a "ticketed bike", and second, it is one of the conditions that allows us to differentiate between the legal possibilities regarding the applicability of the concept "ticketed bike". The conditions b_3 and b_4 might be taken into account when deciding the applicability of this concept. This means that they might be used to differentiate between legal possibilities. That is, either b_2 and b_3 collectively, or b_2 and b_4 collectively, provide enough information to distinguish between the legal possibilities to which the concept c_2 is applicable and those to which it is not. Hence, the question of whether the bike is a child bike or folded to a size larger than 45 $\times$ 86 $\times$ 80 cm might have an impact on whether this bike is a "ticketed bike". In contrast, condition b_1 does not affect the applicability of c_2. Thus, the information on whether a bike is a city bike does not help to distinguish between legal possibilities with respect to the applicability of the concept "ticketed bike". $\triangle$

Additionally, let us illustrate some properties about applicability-relevance and applicability-irrelevance. Importantly, if a concept is constant, which means it is either always defined to be applicable or inapplicable, then only the empty team is applicability-relevant, whereas any other non-empty set of conditions is applicability-irrelevant. This is formally shown in Fact 5.19 and illustrates an important feature of this definition of relevance conditions: If the legal information specified in a legal team does not contain a specification of the applicability and non-applicability of a concept, then any condition is applicability-irrelevant.

Fact 5.19. Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and T^c be a legal team on $B \cup \{c\}$.

(1) If X is totally applicability-relevant for c relative to T^c, then X is applicability-relevant for c and is not applicability-irrelevant for c relative to T^c.

(2) $\emptyset$ is applicability-relevant for c and totally applicability-relevant for c relative to T^c.

(3) c is not constant if and only if there exists some non-empty $X \subseteq B$ such that X is applicability-relevant for c relative to T^c.

Proof. Proof is given in Appendix A.3. $\square$

Furthermore, the properties of sufficiency, being an exception, and positive and negative necessity, can be used to determine the applicability-relevance of ground conditions of a legal team.

Fact 5.20. Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and T^c be a legal team on $B \cup \{c\}$. Suppose that c is not constant with respect to T^c, then the following facts hold:

(1) If a legal possibility $X \subseteq B$ expressed by some valuation of T^c is sufficient and positively necessary relative to T^c, then some non-empty set of conditions $Y \subseteq X$ is applicability-relevant for c relative to T^c.

(2) If a legal possibility $X \subseteq B$ expressed by some valuation of T^c is an exception and negatively necessary relative to T^c, then some non-empty set of conditions $Y \subseteq X$ is applicability-relevant for c relative to T^c.

Proof. Proof is given in Appendix A.3. $\square$

As an immediate consequence of Fact 5.20, any ground condition of a legal team that is sufficient and positively necessary relative to this legal team is applicability-relevant for the considered concept. Similarly, a condition is applicability-relevant for the considered concept, if the condition is an exception and negatively necessary.

Lastly, let us now focus on how this notion of applicability-relevant conditions formalizes the intuition illustrated at the beginning of this section. For this purpose, let us formally explore Example 5.16.

Example 5.21. The concept "ticketed bike" is denoted by c_2. Let the conditions be "city bike" and "folding bike". These conditions are denoted by b_1 and b_2 so that $B = \{b_1, b_2\}$. Further, suppose that

$B_1^\perp = \emptyset$ and $B_2^\perp = \emptyset$. Table 5.6 characterizes the legal team T^{c_2} on $B \cup \{c_2\}$ that specifies the concept "ticketed bike" according to the meaning postulates given in Example 5.16.

valuation v_i	legal possibility X_i	b_1	b_2	concept c_2
v_1	$\{b_1\}$	**1**	**0**	**1**
v_2	$\{b_1, b_2\}$	**1**	**1**	**0**

Table 5.6: Legal team T^{c_2} for the concept "ticketed bike"

According to T^{c_2}, only the sets $\{b_2\}$ and $\{b_1, b_2\}$ are dependency sets. This is the case because $T^{c_2} \models {=}(b_2, c_2)$ and $T^{c_2} \models {=}(b_1, b_2, c_2)$, but $T^{c_2} \not\models {=}(c_2)$ and $T^{c_2} \not\models {=}(b_1, c_2)$. It follows that the only minimal dependency set is $\{b_2\}$, so b_2 is applicability-relevant and totally applicability-relevant. Since $\{b_1\} \not\subseteq \{b_2\}$, it follows that b_1 is not contained in any minimal dependency set. Therefore, b_1 is applicability-irrelevant.

This means that information about the satisfaction of b_2 differentiates the legal possibilities to which the concept c_2 is applicable and those to which it is not. This captures the intuition that the concept is not applicable to legal possibilities about which it is known that they satisfy the condition of being a "folding bike", while it is applicable to legal possibilities about which this is not known. $\triangle$

In addition, this example shows that ground conditions of a legal team might not be applicability-relevant. This leads to a distinction between ground conditions and applicability-relevant conditions, which is due to their different functions in legal reasoning.

Ground conditions constitute legal possibilities, which means they are the entities that are reasoned about. In this example, the legal system only provides rules determining the applicability of the concept c_2 with respect to entities about which it is known that they are "city bikes". Therefore, any legal possibility contains the information that the condition "city bike" is satisfied. Nevertheless, the condition "city bike" does not influence the applicability of the concept, as the applicability is solely determined by information about the applicability-relevant condition "folding bike".

This exemplifies that the formalism constitutes a two-step process. Ground conditions are used to determine the legal possibilities, and given these possibilities, (totally) applicability-relevant conditions determine the applicability of the concept. Since ground conditions might not be applicability-relevant, a legal system might require more information to establish that something is a legal possibility than the information required to determine the applicability.

Chapter 6

Conclusion and Future Work

As AI increasingly influences legal processes, the focus on logical models of legal thinking becomes all the more critical. These models help to ensure interpretability, accountability, and trust. This thesis introduced a new framework to reason about the applicability of legal concepts, based on propositional dependence logic with the might-operator $\mathbf{PL}(=(\cdot), \blacklozenge)$. The thesis not only demonstrated the usefulness of using this logic for this purpose, but also defined and investigated a new mechanism to determine the applicability of concepts to a case.

In Chapter 1, an overview of the legal theory on the applicability of legal concepts was provided, and approaches to model the applicability of concepts were outlined. In Sections 1.1 and 1.2, we clarified the applicability of a legal concepts to a case is based on the acquired knowledge about the case and on the meaning postulates that specify the applicability of the legal concept. Furthermore, we motivated that our formalism is concerned with the applicability of legal concepts whose applicability does not depend on other legal concepts.

We provided the foundation for modeling this understanding of applicability of concepts in Chapter 2. We introduced core notions such as legal possibilities (non-empty and non-conflicting sets of conditions) and legal cases (cases about which it is known that they satisfy some non-conflicting conditions) in Section 2.1. Further, we defined concept applicability functions in Section 2.2 which are functions which assigns Boolean values to legal possibilities to illustrate whether a concept is applicable or not. These concept applicability functions allow us to represent the meaning postulates of legal concepts. Moreover, we examined properties of legal possibilities with respect to a concept applicability function. Based on these formalizations, the main contributions of this thesis are presented in the following chapters.

In Chapter 3, we showed $\mathbf{PL}(=(\cdot), \blacklozenge)$ can be used to model the applicability of legal terms. To do this, we introduced $\mathbf{PL}(=(\cdot), \blacklozenge)$ in Section 3.1. In Section 3.2, we introduced *legal teams* that translate concept applicability functions into teams so that one can reason about these functions using $\mathbf{PL}(=(\cdot), \blacklozenge)$. Importantly, Theorem 3.11 defined the conditions under which a team translates a concept applicability function and hence is a legal team. Furthermore, we provide a translation of the properties of legal possibility relative to a concept applicability function into formulas of $\mathbf{PL}(=(\cdot), \blacklozenge)$.

In Chapter 4, we investigated how the applicability of a legal concept to a legal case can be directly or heuristically derived relative to a legal team. A mechanism capable of directly deriving the applicability is presented in Section 4.2. In this mechanism, the legal team functions as a database against which the acquired knowledge about a case is compared in order to determine the applicability of the concept to the case. If the applicability cannot be directly derived, we outlined a mechanism that can heuristically derive the applicability and hence predict the applicability in Section 4.3. This mechanism is based on the notion of *constant extensions*, defined in Definition 4.8. Due to the requirement that any prediction must be a constant extension, the prediction that a concept is applicable to a case ensures that it does not violate any properties of the legal possibilities expressed by the legal team. Therefore, these properties guide the prediction about the concept's applicability. Further, we examined the requirements necessary to predict applicability in Section 4.4. A characterization of the requirements is given by Theorem 4.16.

In Chapter 5, we examined how the proposed formalism is capable of addressing several tasks associated with legal concept applicability. We investigated the stability of predictions in Section 5.1, whereby we specified the conditions a legal team needs to satisfy so that a prediction is not stable and more information about the considered case leads to a different prediction about the applicability. In Section 5.2, we examined sequences of predictions. This allowed us to characterize how the order of predictions about the applicability of a concept to cases influences these predictions. This means that, depending on the position of the case in the sequence, the case might be predicted differently. Applicability-relevant conditions are investigated in Section 5.3. These conditions allow us to differentiate between legal possibilities to which a concept is applicable and legal possibilities to which it is not applicable. In this sense, they are the conditions relevant for the applicability. In this section, we defined applicability-relevant conditions and discussed why not every ground condition of a legal team is an applicability-relevant condition. This led to the conceptual insight that certain conditions are only used to establish an entity about which the legal system can reason.

The proposed framework gives rise to different directions for further research. In the following, we outline three promising directions.

Research Question 1: Hierarchies of legal concepts

Currently, our formalism is not capable of representing hierarchies of legal concepts. The applicability of any considered concept does not depend on any other concepts. To overcome this limitation, the formalism needs to be extended so that concepts can depend on further concepts.

A straightforward extension is a set of *states* $S \subseteq \mathcal{B} \cup \mathcal{C}$, which is a non-empty set of conditions and concepts. This allows us to express the dependence of the applicability of a concept on another concept. For example, consider the following table sketching the applicability of the concepts c_1 and c_2. Let T be a team on $S \cup \{c_2\}$, where $S = \{c_1, b_3, b_4\}$, and let T' be a team on $S' \cup \{c_1\}$, where $S' = \{b_1, b_2\}$.

valuation v_i	legal possibility X_i	b_1	b_2	concept c_1	b_3	b_4	concept c_2
v_1	$\{b_1, b_3\}$	**1**	**0**	**1**	**1**	**0**	**1**
v_2	$\{b_1, b_2, b_4\}$	**1**	**1**	**1**	**0**	**1**	**0**
v_3	$\{b_2, b_4\}$	**0**	**1**	**0**	**0**	**1**	**0**

This encodes that c_1 is determined by the conditions b_1 and b_2. Based on the applicability of c_1 and the conditions b_3 and b_4, the applicability of c_2 is then determined. Given a new case, one needs to first predict the applicability of c_1 to be in a position to predict c_2.

This intuitively ad hoc extension leads to several interesting questions. First, *sets of states* need to be formally defined and constrained. Based on this definition, concept applicability functions and legal teams need to be extended accordingly. Second, similar to the work by Canavotto and Horty [CH23a; CH23b], it needs to be investigated whether hierarchies of concepts can be flattened to the simple legal teams discussed in this work, and whether hierarchies and flattened hierarchies result in identical judgments on the applicability of concepts.

Research Question 2: Extending the set of preserved properties of constant extensions

Determining the prediction of the applicability of a concept is based on constant extensions. It is important that a constant extension not only is a legal team, but also preserves certain properties of the expressed legal possibilities. The investigated properties are sufficiency, being an exception, positive necessity, negative necessity, and contingency. These properties by a legal possibility X_i relative to a legal team are expressed by $\varphi_{\mathsf{suf}}^{X_i}$, $\varphi_{\mathsf{exc}}^{X_i}$, $\varphi_{\mathsf{nec+}}^{X_i}$, $\varphi_{\mathsf{nec-}}^{X_i}$, and $\varphi_{\mathsf{con}}^{X_i}$, respectively.

There are good reasons to extend this set of properties that are preserved under constant extensions. A natural extension would be an operator that expresses that a legal possibility X_i satisfies a property if another legal possibility X_j is not taken into account. Let us call this operator $\mathsf{except}_{X_j}(\psi)$, where $\psi \in \{\varphi_{\mathsf{suf}}^{X_i}, \varphi_{\mathsf{exc}}^{X_i}, \varphi_{\mathsf{nec+}}^{X_i}, \varphi_{\mathsf{nec-}}^{X_i}, \varphi_{\mathsf{con}}^{X_i}\}$. Given that $\psi := \varphi_{\mathsf{suf}}^{X_i}$, this expresses that the legal possibility X_i is sufficient if X_j is not considered.

This intuitive characterization gives rise to both a strict and a weak formalization of this operator. The former expresses that only the legal possibility X_j is excluded. The latter states that any legal possibility that includes X_j is excluded. Given a legal team on a set of conditions B and a concept, the strict operator can be formalized as:

$$\mathsf{except}_{X_j}^{\mathsf{strict}}(\psi) := \left(\bigwedge_{b_i \in X_j} b_i \wedge \bigwedge_{b_j \in B \setminus X_j} \neg b_j \right) \vee \psi \quad \text{where } \psi \in \{\varphi_{\mathsf{suf}}^{X_i}, \varphi_{\mathsf{exc}}^{X_i}, \varphi_{\mathsf{nec+}}^{X_i}, \varphi_{\mathsf{nec-}}^{X_i}, \varphi_{\mathsf{con}}^{X_i}\}$$

Intuitively, $\mathsf{except}_{X_j}^{\mathsf{strict}}(\varphi_{\mathsf{suf}}^{X_i})$ expresses that the concept c is applicable to any legal possibility that is distinct from X_j and includes the legal possibility X_i. Therefore, it expresses that X_i is sufficient, provided only X_j is excluded.

In contrast, the weak operator allows for the exclusion of more than one legal possibility. This is because it excludes any legal possibility that includes X_j. The weak operator can be defined as:

$$\mathsf{except}_{X_j}^{\mathsf{weak}}(\psi) := \bigwedge_{b_i \in X_j} b_i \vee \psi \quad \text{where } \psi \in \{\varphi_{\mathsf{suf}}^{X_i}, \varphi_{\mathsf{exc}}^{X_i}, \varphi_{\mathsf{nec+}}^{X_i}, \varphi_{\mathsf{nec-}}^{X_i}, \varphi_{\mathsf{con}}^{X_i}\}$$

Such operators raise several conceptual questions. It must be clarified what properties the weak and strict operators capture, and based on this understanding, which operator is better suited for the legal domain. The current formalization only allows the exclusion of one legal possibility. However, it must be discussed whether this is a sensible assumption. Thus, it should be investigated whether it is appropriate to define weak and strict exceptions for one or multiple legal possibilities. Based on these decisions, it is worth investigating how the implementation of such an operator affects the predictability of a concept's applicability to a case.

Furthermore, it is worth noting that an except-operator might not be the only viable extension. The framework provides sufficient flexibility to define new properties and to investigate how these properties affect the predictability of a concept in relation to a case. Therefore, defining new properties of legal possibilities represents a fruitful direction for further research.

Research Question 3: Non-monotonicity and concept applicability

Recently, non-monotonicity has attracted attention from researchers studying propositional dependence logic.[8] Our work does not investigate non-monotonicity explicitly, but it does play an important role in the modeled legal reasoning. The instability of predictions, discussed in Section 5.1, exemplifies the non-monotonic character of reasoning about the applicability of a concept. That is, further information about a case might result in a different judgment on the applicability of a concept.

Since we do not investigate this non-monotonic nature of legal reasoning in depth, an interesting question is to explore the proposed mechanism in light of the research on non-monotonicity and propositional dependence logic.

Such an investigation could deepen our understanding of meaning postulates. Currently, the formalism lacks a mechanism for expressing that a concept is typically, though not always, applicable to a given legal possibility. As a result, the formalism cannot differentiate between a detailed characterization of an exception and the general characterization of applicability. Such a distinction would allow for a more fine-grained analysis of meaning postulates by enabling differentiation between edge cases and those that reflect regular scenarios. This would lead to a better representation of meaning postulates.

[8] See, for instance, Sauerwald and Kontinen [SK24a; SK24b].

Bibliography

[AAB16] Latifa Al-Abdulkarim, Katie Atkinson, and Trevor Bench-Capon. "A methodology for designing systems to reason with legal cases using Abstract Dialectical Frameworks". In: *Artificial Intelligence and Law* 24.1 (2016), pages 1–49. `https://doi.org/10.1007/s10506-016-9178-1` (cited on pages 2, 13, 14).

[AB05a] Katie Atkinson and Trevor Bench-Capon. "Legal Case-based Reasoning as Practical Reasoning". In: *Artificial Intelligence and Law* 13.1 (2005), pages 93–131. `https://doi.org/10.1007/s10506-006-9003-3` (cited on pages 5, 6).

[AB05b] Katie Atkinson and Trevor Bench-Capon. "Levels of Reasoning With Legal Cases". In: *ICAIL 2005 Workshop on Argumentation in AI and Law*. IAAIL Workshop Series (2005), pages 1–11 (cited on pages 5, 6).

[AK25] Aleksi Anttila and Søren Brinck Knudstorp. *Convex Team Logics*. 2025. `https://doi.org/10.48550/arXiv.2503.21850` (cited on page 22).

[Al-+18] Latifa Al-Abdulkarim, Katie Atkinson, Trevor Bench-Capon, Stuart Whittle, Rob Williams, and Catriona Wolfenden. "Noise induced hearing loss: Building an application using the ANGELIC methodology". In: *Argument & Computation* 10.1 (2018), pages 5–22. `https://doi.org/10.3233/AAC-181005` (cited on pages 13, 14).

[Ale20] Robert Alexy. *Theorie der Grundrechte*. 9. Auflage. Suhrkamp-Taschenbuch Wissenschaft 582. Frankfurt am Main: Suhrkamp, 2020. ISBN: 978-3-518-28182-6 (cited on page 30).

[Ant25] Aleksi Anttila. "Not Nothing. Nonemptiness in Team Semantics". PhD thesis. Amsterdam, Netherlands: Universiteit van Amsterdam, 2025. `https://eprints.illc.uva.nl/id/eprint/2350/1/DS-2025-04.text.pdf` (cited on page 21).

[Ari20] Aristoteles. *Nikomachische Ethik (Griechisch/Deutsch): Reclams Universal-Bibliothek*. Edited and translated by Gernot Krapinger. 1. Auflage. Reclams Universal-Bibliothek. Ditzingen: Reclam Verlag, 2020. ISBN: 978-3-15-019670-0 (cited on page 29).

[Ash17] Kevin D. Ashley. *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age*. 1st edition. Cambridge University Press, July 10, 2017. ISBN: 978-1-316-76138-0. `https://doi.org/10.1017/9781316761380` (cited on page 10).

[Ben22] Trevor Bench-Capon. "Thirty years of Artificial Intelligence and Law: Editor's Introduction". In: *Artificial Intelligence and Law* 30.4 (2022), pages 475–479. `https://doi.org/10.1007/s10506-022-09325-8` (cited on page 2).

[Ben23] Trevor Bench-Capon. "A Note on Hierarchical Constraints". In: *Frontiers in Artificial Intelligence and Applications*. Edited by Giovanni Sileno, Jerry Spanakis, and Gijs Van Dijck. IOS Press, 2023. ISBN: 978-1-64368-472-7. `https://doi.org/10.3233/FAIA230940` (cited on page 11).

[Ben24] Trevor Bench-Capon. "Intermediate factors and precedential constraint". In: *Artificial Intelligence and Law* (2024). https://doi.org/10.1007/s10506-024-09405-x (cited on page 11).

[Bre+13] Gerhard Brewka, Stefan Ellmauthaler, Hannes Strass, Johannes Wallner, and Stefan Woltran. "Abstract Dialectical Frameworks Revisited". In: *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence* (2013), pages 803–809 (cited on page 12).

[BT12] Jan Broersen and Leendert van der Torre. "Ten Problems of Deontic Logic and Normative Reasoning in Computer Science". In: *Lectures on Logic and Computation: ESSLLI 2010 Copenhagen, Denmark, August 2010, ESSLLI 2011, Ljubljana, Slovenia, August 2011, Selected Lecture Notes.* 2012, pages 55–88 (cited on page 4).

[BW10] Gerhard Brewka and Stefan Woltran. "Abstract Dialectical Frameworks". In: *Association for the Advancement of Artificial Intelligence.* 2010 (cited on page 12).

[CAB20] Joe Collenette, Katie Atkinson, and Trevor Bench-Capon. "An Explainable Approach to Deducing Outcomes in European Court of Human Rights Cases Using ADFs". In: *Frontiers in Artificial Intelligence and Applications.* IOS Press, 2020. https://doi.org/10.3233/FAIA200488 (cited on page 13).

[CAB23] Joe Collenette, Katie Atkinson, and Trevor Bench-Capon. "Explainable AI tools for legal reasoning about cases: A study on the European Court of Human Rights". In: *Artificial Intelligence* 317 (2023), page 103861. https://doi.org/10.1016/j.artint.2023.103861 (cited on pages 13, 14).

[Can25] Ilaria Canavotto. "Reasoning with inconsistent precedents". In: *Artificial Intelligence and Law* 33.1 (2025), pages 137–166. https://doi.org/10.1007/s10506-023-09382-7 (cited on page 11).

[CGR18] Ivano Ciardelli, J. A. G. Groenendijk, and Floris Roelofsen. *Inquisitive semantics.* First edition. Oxford: Oxford University Press, 2018. ISBN: 978-0-19-185247-3. https://doi.org/10.1093/oso/9780198814788.001.0001 (cited on page 2).

[CH23a] Ilaria Canavotto and John Horty. "Reasoning with hierarchies of open-textured predicates". In: *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law.* ICAIL 2023: Nineteenth International Conference on Artificial Intelligence and Law. Braga Portugal: ACM, June 19, 2023, pages 52–61. ISBN: 9798400701979. https://doi.org/10.1145/3594536.3595148. https://dl.acm.org/doi/10.1145/3594536.3595148 (cited on pages 11, 14, 62).

[CH23b] Ilaria Canavotto and John Horty. "The Importance of Intermediate Factors". In: *Frontiers in Artificial Intelligence and Applications.* Edited by Giovanni Sileno, Jerry Spanakis, and Gijs Van Dijck. IOS Press, 2023. ISBN: 978-1-64368-472-7. https://doi.org/10.3233/FAIA230941 (cited on pages 11, 14, 62).

[Deg24] Marco Degano. "Indefinites and their values". PhD thesis. Amsterdam, Netherlands: Universiteit van Amsterdam, 2024. https://eprints.illc.uva.nl/id/eprint/2319/1/DS-2024-07.text.pdf (cited on page 2).

[Di +23] Cecilia Di Florio, Xinghan Liu, Emiliano Lorini, Antonino Rotolo, and Giovanni Sartor. "Inferring New Classifications in Legal Case-Based Reasoning". In: *Frontiers in Artificial Intelligence and Applications.* Edited by Giovanni Sileno, Jerry Spanakis, and Gijs Van

Dijck. IOS Press, 2023. ISBN: 978-1-64368-472-7. https://doi.org/10.3233/FAIA230942 (cited on page 13).

[Dua19] Luís Duarte d'Almeida. "On the Legal Syllogism". In: *Dimensions of Normativity*. Edited by David Plunkett, Scott J. Shapiro, and Kevin Toh. 1st edition. New York: Oxford University Press, 2019, pages 335–364. ISBN: 978-0-19-064040-8. https://doi.org/10.1093/oso/9780190640408.003.0015 (cited on page 6).

[Dua21] Luís Duarte d'Almeida. "What Is It to Apply the Law?" In: *Law and Philosophy* 40.4 (2021), pages 361–386. https://doi.org/10.1007/s10982-021-09405-x (cited on page 6).

[Dun95] Phan Minh Dung. "On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games". In: *Artificial Intelligence* 77.2 (1995), pages 321–357. https://doi.org/10.1016/0004-3702(94)00041-X (cited on page 12).

[Fis91] Stanley Fish. "The Law Wishes to Have a Formal Existence". In: *The Fate of law*. Edited by Austin Sarat and Thomas R. Kearns. Ann Arbor: University of Michigan Press, 1991, pages 159–208. ISBN: 978-0-472-10244-0 (cited on page 6).

[Frä87] Åke Frändberg. "An Essay on the Systematics of Legal Concepts - A Study of Legal Concept Formation". In: *Scandinavian Studies in Law* 31 (1987), pages 81–116 (cited on page 3).

[GJ13] Davide Grossi and Andrew J. I. Jones. "Constitutive Norms and Counts-as Conditionals". In: *Handbook of deontic logic and normative systems*. Edited by Dov M. Gabbay and John Horty. Volume 1. London: College Publications, 2013, pages 407–441. ISBN: 978-1-84890-132-2 (cited on page 11).

[GK12] Tomasz Gizbert-Studnicki and Mateusz Klinowski. "Are Legal Concepts Embedded in Legal Norms?" In: *International Journal for the Semiotics of Law - Revue internationale de Sémiotique juridique* 25.4 (2012), pages 553–562. https://doi.org/10.1007/s11196-011-9238-5 (cited on page 4).

[GMD05] Davide Grossi, John-Jules Ch. Meyer, and Frank Dignum. "Modal logic investigations in the semantics of counts-as". In: *Proceedings of the 10th international conference on Artificial intelligence and law*. ICAIL05: 10th International Joint Conference on Artificial Intellignece and Law. Bologna Italy: ACM, 2005, pages 1–9. https://doi.org/10.1145/1165485.1165487 (cited on pages 11, 14).

[GMD08] Davide Grossi, John-Jules Ch. Meyer, and Frank Dignum. "The many faces of counts-as: A formal analysis of constitutive rules". In: *Journal of Applied Logic* 6.2 (2008), pages 192–217. https://doi.org/10.1016/j.jal.2007.06.008 (cited on pages 11, 14).

[HK16] Miika Hannula and Juha Kontinen. "A finite axiomatization of conditional independence and inclusion dependencies". In: *Information and Computation* 249 (2016), pages 121–137. https://doi.org/10.1016/j.ic.2016.04.001. https://linkinghub.elsevier.com/retrieve/pii/S0890540116300025 (cited on page 2).

[HM18] H. L. A. Hart and Christoph Möllers. *Der Begriff des Rechts: mit dem Postskriptum von 1994 und einem Nachwort von Christoph Möllers*. Translated by Alexander von Baeyer. 2. Auflage. Suhrkamp-Taschenbuch Wissenschaft 2009. Berlin: Suhrkamp, 2018. ISBN: 978-3-518-29609-7 (cited on page 6).

[Hor04] John F. Horty. “THE RESULT MODEL OF PRECEDENT”. In: *Legal Theory* 10.1 (2004), pages 19–31. https://doi.org/10.1017/S1352325204000151 (cited on pages 10, 14).

[Hor11] John F. Horty. “Reasons and precedent”. In: *Proceedings of the 13th International Conference on Artificial Intelligence and Law*. ICAIL ’11: International Conference on Artificial Intelligence and Law. Pittsburgh Pennsylvania: ACM, 2011, pages 41–50. https://doi.org/10.1145/2018358.2018363 (cited on pages 2, 10, 14).

[Hor19] John Horty. “Reasoning with dimensions and magnitudes”. In: *Artificial Intelligence and Law* 27.3 (2019), pages 309–345. https://doi.org/10.1007/s10506-019-09245-0 (cited on page 11).

[HPT07] Jörg Hansen, Gabriella Pigozzi, and Leendert van der Torre. “Ten Philosophical Problems in Deontic Logic”. In: *Dagstuhl Seminar Proceedings 07122*. Normative Multiagent Systems. Edited by Leendert van der Torre, Guido Boella, and Harko Verhagen. 2007 (cited on page 4).

[HS15] Lauri Hella and Johanna Stumpf. “The expressive power of modal logic with inclusion atoms”. In: *Electronic Proceedings in Theoretical Computer Science* 193 (2015), pages 129–143. https://doi.org/10.4204/EPTCS.193.10 (cited on page 22).

[JH21] Jan C. Joerden and Eric Hilgendorf. “Logik und Recht”. In: *Handbuch Rechtsphilosophie*. Edited by Jan C. Joerden. Stuttgart: J.B. Metzler, 2021, pages 240–243. ISBN: 978-3-476-05638-2. https://doi.org/10.1007/978-3-476-05639-9_34 (cited on page 6).

[JS96] Andrew J. I. Jones and Marek Sergot. “A Formal Characterization of Institutionalised Power”. In: *Journal of the IGPL* 3 (1996), pages 427–443 (cited on pages 11, 14).

[Kel17] Hans Kelsen. *Reine Rechtslehre: Mit einem Anhang: Das Problem der Gerechtigkeit (Studienausgabe der 2. Auflage 1960)*. Edited by Matthias Jestaedt. 1st ed. Tübingen: Mohr Siebeck, 2017. ISBN: 978-3-16-152973-3 (cited on pages 6, 30).

[KW24] Hans Kelsen and Robert Walter. *Was ist Gerechtigkeit?* 6. Auflage. Ditzingen: Reclam, 2024. ISBN: 978-3-15-019366-2 (cited on page 30).

[Lie87] Anne von der Lieth Gardner. *An artificial intelligence approach to legal reasoning*. A Bradford book. Cambridge, Mass.: MIT Press, 1987. ISBN: 978-0-262-07104-8 (cited on pages 4, 5).

[Lin04] Lars Lindahl. “Deduction and Justification in the Law. The Role of Legal Terms and Concepts”. In: *Ratio Juris* 17.2 (2004), pages 182–202. https://doi.org/10.1111/j.1467-9337.2004.00263.x (cited on page 5).

[Liu+22] Xinghan Liu, Emiliano Lorini, Antonino Rotolo, and Giovanni Sartor. “Modelling and Explaining Legal Case-Based Reasoners Through Classifiers”. In: *Frontiers in Artificial Intelligence and Applications*. Edited by Enrico Francesconi, Georg Borges, and Christoph Sorge. IOS Press, 2022. ISBN: 978-1-64368-364-5. https://doi.org/10.3233/FAIA220451 (cited on pages 13, 14).

[LL21] Xinghan Liu and Emiliano Lorini. “A Logic for Binary Classifiers and Their Explanation”. In: *Logic and Argumentation*. Edited by Pietro Baroni, Christoph Benzmüller, and Yì N Wáng. Volume 13040. Series Title: Lecture Notes in Computer Science. Cham: Springer International Publishing, 2021, pages 302–321. ISBN: 978-3-030-89390-3. https://doi.org/10.1007/978-3-030-89391-0_17 (cited on page 13).

[LO06] Lars Lindahl and Jan Odelstad. “Intermediate Concepts in Normative Systems”. In: *Deontic Logic and Artificial Normative Systems*. Edited by Lou Goble and John-Jules Ch. Meyer. Redacted by David Hutchison, Takeo Kanade, Josef Kittler, Jon M. Kleinberg, Friedemann Mattern, John C. Mitchell, Moni Naor, Oscar Nierstrasz, C. Pandu Rangan, Bernhard Steffen, Madhu Sudan, Demetri Terzopoulos, Dough Tygar, Moshe Y. Vardi, and Gerhard Weikum. Volume 4048. Series Title: Lecture Notes in Computer Science. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pages 187–200. ISBN: 978-3-540-35842-8. `https://doi.org/10.1007/11786849_16` (cited on pages 12, 14).

[LO08] Lars Lindahl and Jan Odelstad. “Intermediaries and intervenients in normative systems”. In: *Journal of Applied Logic* 6.2 (2008), pages 229–250. `https://doi.org/10.1016/j.jal.2007.06.010` (cited on pages 12, 14).

[LO13] Lars Lindahl and Jan Odelstad. “The Theory of Joining-Systems”. In: *Handbook of deontic logic and normative systems*. Edited by Dov M. Gabbay and John Horty. London: College Publications, 2013, pages 545–634. ISBN: 978-1-84890-132-2 (cited on pages 12, 14).

[Pra21] Henry Prakken. “A formal analysis of some factor- and precedent-based accounts of precedential constraint”. In: *Artificial Intelligence and Law* 29.4 (2021), pages 559–585. `https://doi.org/10.1007/s10506-021-09284-6` (cited on page 11).

[PS15] Henry Prakken and Giovanni Sartor. “Law and logic: A review from an argumentation perspective”. In: *Artificial Intelligence* 227 (2015), pages 214–245. `https://doi.org/10.1016/j.artint.2015.06.005` (cited on pages 2, 4).

[PT17] Gabriella Pigozzi and Leendert van der Torre. “Multiagent deontic logic and its challenges from a normative systems perspective”. In: *IfCoLog Journal of Logics and Their Applications* (2017), pages 2929–2993 (cited on page 4).

[PY16] Eric Pacuit and Fan Yang. “Dependence and Independence in Social Choice: Arrow’s Theorem”. In: *Dependence Logic*. Edited by Samson Abramsky, Juha Kontinen, Jouko Väänänen, and Heribert Vollmer. Cham: Springer International Publishing, 2016, pages 235–260. ISBN: 978-3-319-31801-1. `https://doi.org/10.1007/978-3-319-31803-5_11` (cited on page 2).

[RAB05] Edwina L. Rissland, Kevin D. Ashley, and L. Karl Branting. “Case-based reasoning and law”. In: *The Knowledge Engineering Review* 20.3 (2005), pages 293–298. `https://doi.org/10.1017/S0269888906000701` (cited on page 10).

[Ros57] Alf Ross. “Tû-tû”. In: *Harvard Law Review* 70.5 (1957), pages 812–825 (cited on pages 4, 5).

[Sar09] Giovanni Sartor. “Legal concepts as inferential nodes and ontological categories”. In: *Artificial Intelligence and Law* 17.3 (2009), pages 217–251. `https://doi.org/10.1007/s10506-009-9079-7` (cited on page 4).

[Sea11] John R. Searle. *Making the social world: the structure of human civilization*. First publ. in paperback. Oxford: Oxford University Press, 2011. ISBN: 978-0-19-957691-3 (cited on page 3).

[SK24a] Kai Sauerwald and Juha Kontinen. “A First Peek into Preferential Logics with Team Semantics”. In: *Proceedings of the 22nd International Workshop on Nonmonotonic Reasoning (NMR 2024) co-located with 21st International Conference on Principles of Knowledge Representation and Reasoning (KR 2024), Hanoi, Vietnam, November 2-4, 2024*. 22nd International Workshop on Nonmonotonic Reasoning, November 2-4, 2024, Hanoi, Vietnam. Edited by Nina Gierasimczuk and Jesse Heyninck. Volume 3835. Aachen: CEUR Workshop Proceedings, 2024, pages 94–100. https://ceur-ws.org/Vol-3835/paper10.pdf (cited on page 63).

[SK24b] Kai Sauerwald and Juha Kontinen. *A Primer for Preferential Non-Monotonic Propositional Team Logics*. 2024. https://doi.org/10.48550/arXiv.2405.06973 (cited on page 63).

[Som06] Alexander Somek. *Rechtliches Wissen*. 1. Aufl. Suhrkamp Taschenbuch Wissenschaft 1802. Frankfurt am Main: Suhrkamp, 2006. ISBN: 978-3-518-29402-4 (cited on page 30).

[Vää07a] Jouko Väänänen. *Dependence logic: a new approach to independence friendly logic*. London Mathematical Society student texts 70. Cambridge: Cambridge University Press, 2007. ISBN: 978-0-521-87659-9 (cited on pages 21, 22, 56).

[Vää07b] Jouko Väänänen. “Team Logic”. In: *Interactive Logic. Selected Papers from the 7th Augustus de Morgan Workshop, London*. Edited by Johan van Benthem, Dov Gabbay, and Benedikt Löwe. 2007, pages 281–302 (cited on pages 22, 56).

[Van+23a] Wijnand Van Woerkom, Davide Grossi, Henry Prakken, and Bart Verheij. “Hierarchical a Fortiori Reasoning with Dimensions”. In: *Frontiers in Artificial Intelligence and Applications*. Edited by Giovanni Sileno, Jerry Spanakis, and Gijs Van Dijck. IOS Press, 2023. ISBN: 978-1-64368-472-7. https://doi.org/10.3233/FAIA230944 (cited on pages 11, 14).

[Van+23b] Wijnand Van Woerkom, Davide Grossi, Henry Prakken, and Bart Verheij. “Hierarchical Precedential Constraint”. In: *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*. ICAIL 2023: Nineteenth International Conference on Artificial Intelligence and Law. Braga Portugal: ACM, 2023, pages 333–342. ISBN: 9798400701979. https://doi.org/10.1145/3594536.3595154 (cited on pages 11, 14).

[Vel96] Frank J. M. M. Veltman. “Defaults in update semantics”. In: *Journal of Philosophical Logic* 25 (1996), pages 221–261 (cited on page 22).

[Vil+22] Serena Villata, Michal Araszkiewicz, Kevin Ashley, Trevor Bench-Capon, L. Karl Branting, Jack G. Conrad, and Adam Wyner. “Thirty years of artificial intelligence and law: the third decade”. In: *Artificial Intelligence and Law* 30.4 (2022), pages 561–591. https://doi.org/10.1007/s10506-022-09327-6 (cited on page 2).

[Wes90] Peter Westen. *Speaking of equality: an analysis of the rhetorical force of ”equality” in moral and legal discourse*. Studies in moral, political, and legal philosophy. Princeton (N.J.): Princeton University Press, 1990. ISBN: 978-0-691-07858-8 (cited on page 29).

[Wró71] Jerzy Wróblewski. “Legal decision and its justification”. In: *Logique et Analyse* 14.53 (1971), pages 409–419 (cited on page 6).

[Wró74] Jerzy Wróblewski. “Legal Syllogism and Rationality of Judicial Decision”. In: *Rechtstheorie* 5.1 (1974), pages 33–46 (cited on page 6).

[Wyn08] Adam Wyner. “An ontology in OWL for legal case-based reasoning”. In: *Artificial Intelligence and Law* 16.4 (2008), pages 361–387. https://doi.org/10.1007/s10506-008-9070-8 (cited on pages 5, 6).

[Yan23] Jialiang Yan. “Monotonicity in Intensional Contexts: Weakening and Pragmatic Effects under Modals and Attitudes”. PhD thesis. Amsterdam, Netherlands: Tsinghua University and Universiteit van Amsterdam, 2023. https://eprints.illc.uva.nl/id/eprint/2254/1/DS-2023-06.text.pdf (cited on page 22).

[YV16] Fan Yang and Jouko Väänänen. "Propositional logics of dependence". In: *Annals of Pure and Applied Logic* 167.7 (2016), pages 557–589. https://doi.org/10.1016/j.apal.2016.03.003 (cited on page 22).

Appendix A

Appendix

A.1 Proofs of Chapter 3

Proposition 3.14. Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions, $f : \mathcal{D} \to \{\mathbf{1}, \mathbf{0}\}$ be a concept applicability function and T^c be a legal team on $B \cup \{c\}$ that translates f.

(1) A legal possibility $X_i \in \mathcal{D}$ is sufficient relative to f if and only if $T^c \models \varphi_{\mathsf{suf}}^{X_i}$ where

$$\varphi_{\mathsf{suf}}^{X_i} := \bigwedge_{b_i \in X_i} b_i \to c$$

(2) A legal possibility $X_i \in \mathcal{D}$ is an exception relative to f if and only if $T^c \models \varphi_{\mathsf{exc}}^{X_i}$ where

$$\varphi_{\mathsf{exc}}^{X_i} := \bigwedge_{b_i \in X_i} b_i \to \neg c$$

(3) A legal possibility $X_i \in \mathcal{D}$ is positively necessary relative to f if and only if $T^c \models \varphi_{\mathsf{nec+}}^{X_i}$ where

$$\varphi_{\mathsf{nec+}}^{X_i} := c \to \bigwedge_{b_i \in X_i} b_i$$

(4) A legal possibility $X_i \in \mathcal{D}$ is negatively necessary relative to f if and only if $T^c \models \varphi_{\mathsf{nec-}}^{X_i}$ where

$$\varphi_{\mathsf{nec-}}^{X_i} := \neg c \to \bigwedge_{b_i \in X_i} b_i$$

(5) A legal possibility $X_i \in \mathcal{D}$ is contingent relative to f if and only if $T^c \models \varphi_{\mathsf{con}}^{X_i}$ where

$$\varphi_{\mathsf{con}}^{X_i} := \blacklozenge \left(\bigwedge_{b_i \in X_i} b_i \wedge c \right) \wedge \blacklozenge \left(\bigwedge_{b_i \in X_i} b_i \wedge \neg c \right)$$

Proof. To prove the left-to-right direction of (1), suppose that $X_i \in \mathcal{D}$ is sufficient relative to f. By Definition 2.6, for every $X_j \in \mathcal{D}$ with $X_i \subseteq X_j$ it is the case that $f(X_j) = \mathbf{1}$. According to Definition 3.8, this means that each legal possibility expressed by some valuation $v \in T^c$ with $X_i \subseteq \{b \in B \mid v(b) = \mathbf{1}\}$ it is the case that $v(c) = \mathbf{1}$. Now, let $T_1 = \{v \in T^c \mid v(b_i) = \mathbf{1}$ for all $b_i \in X_i\}$ and $T_2 = \{v \in T^c \mid$

$v(b_i) = \mathbf{0}$ for some $b_i \in X_i\}$. Since for all $v_j \in T_1$ it is the case that $X_i \subseteq \{b \in B \mid v_j(b_i) = \mathbf{1}\}$, it follows that $v_j(c) = \mathbf{1}$. Since the formula $c \in \mathbf{PL}$ and by Theorem 3.5, it follows that $T_1 \models c$. Further, by the definition of T_2 it follows that for all $v_k \in T_2$, $\{v_k\} \not\models (\bigwedge_{b_i \in X_i} b_i)$. This implies that $T_2 \models \neg(\bigwedge_{b_i \in X_i} b_i)$. Since it is clearly the case that $T_1 \cup T_2 = T^c$, it follows that $T^c \models \neg(\bigwedge_{b_i \in X_i} b_i) \vee c$. This means that $T^c \models (\bigwedge_{b_i \in X_i} b_i) \rightarrow c$ and $T^c \models \varphi^{X_i}_{\mathsf{suf}}$.

To prove the right-to-left direction of (1), suppose that $T^c \models \varphi^{X_i}_{\mathsf{suf}}$ which means that $T^c \models (\bigwedge_{b_i \in X_i} b_i) \rightarrow c$ and hence $T^c \models \neg(\bigwedge_{b_i \in X_i} b_i) \vee c$. This implies that there exists T_1 and T_2 such that $T_1 \cup T_2 = T^c$, $T_1 \models \neg(\bigwedge_{b_i \in X_i} b_i)$ and $T_2 \models c$. Note that for any $v \in T^c$ such that $v(b_i) = \mathbf{1}$ for all $b_i \in X_i$ it is the case that $\{v\} \models \bigwedge_{b_i \in X_i} b_i$. This means that it is not the case that $\{v\} \not\models \bigwedge_{b_i \in X_i} b_i$ and hence $v \notin T_1$. Since $T^c = T_1 \cup T_2$, it follows that $v \in T_2$ and hence $v(c) = \mathbf{1}$. This means that for any legal possibilities $X_j \in \mathcal{D}$ expressed by some valuation $v_j \in T^c$ such that $X_i \subseteq X_j$ it is the case that $v_j(c) = \mathbf{1}$. By Definition 3.8, for any $X_j \in \mathcal{D}$ such that $X_i \subseteq X_j$ it is the case that $f(X_j) = \mathbf{1}$. According to Definition 2.6, this means that X_i is sufficient relative to f.

Statement (2) follows by analogous reasoning to (1).

To prove the left-to-right direction of (3), suppose that $X_i \in \mathcal{D}$ is positively necessary relative to f. By Definition 2.6 it follows that for every $X_j \in \mathcal{D}$ with $f(X_j) = \mathbf{1}$ it is the case that $X_i \subseteq X_j$. Let $T_1 = \{v \in T^c \mid v(c) = \mathbf{1}\}$ and $T_2 = \{v \in T^c \mid v(c) = \mathbf{0}\}$. It is clearly the case that $T_1 \cup T_2 = T^c$. Due to the construction of T_2, it holds that $T_2 \models \neg c$. Further, for any $v \in T_1$ it is the case that $v(c) = \mathbf{1}$. Since X_i is positively necessary it follows by Definition 3.8 that for all $v \in T^c$ with $v(c) = \mathbf{1}$ it is the case that $\{v\} \models \bigwedge_{b_i \in X_i} b_i$. Since the formula $\bigwedge_{b_i \in X_i} b_i \in \mathbf{PL}$ and by Theorem 3.5, it follows that $T_1 \models \bigwedge_{b_i \in X_i} b_i$. Therefore, it is the case that $T^c \models \neg c \vee (\bigwedge_{b_i \in X_i} b_i)$. This means that $T^c \models c \rightarrow (\bigwedge_{b_i \in X_i} b_i)$ and hence $T^c \models \varphi^{X_i}_{\mathsf{nec+}}$.

To prove the right-to-left direction of (3), suppose that $T^c \models \varphi^{X_i}_{\mathsf{nec+}}$. This means that $T^c \models c \rightarrow (\bigwedge_{b_i \in X_i} b_i)$ and hence $T^c \models \neg c \vee (\bigwedge_{b_i \in X_i} b_i)$. Therefore, there exist T_1 and T_2 such that $T_1 \cup T_2 = T^c$, $T_1 \models \neg c$ and $T_2 \models \bigwedge_{b_i \in X_i} b_i$. For any $v \in T^c$ with $v(c) = \mathbf{1}$ it is the case that $v \notin T_1$. It follows that $v \in T_2$. This means that for any $v \in T^c$ with $v(c) = \mathbf{1}$ it is the case that $\{v\} \models \bigwedge_{b_i \in X_i} b_i$. That is, $v(b_i) = \mathbf{1}$ for all $b_i \in X_i$. By Definition 3.8, it follows that for any $X_j \in \mathcal{D}$ with $f(X_j) = \mathbf{1}$ it is the case that $X_i \subseteq X_j$. According to Definition 2.6, this means that X_i is positively necessary relative to f.

Statement (4) follows by analogous reasoning to (3).

To prove the left-to-right direction of (5), suppose that X_i is contingent relative to f. According to Definition 2.6, there exist X_j and X_k such that $f(X_j) = \mathbf{1}$, $f(X_k) = \mathbf{0}$ and $X_i \subseteq X_j, X_k$. By Definition 3.8, it follows that there exist $v_j, v_k \in T^c$ such that $v_j(b_j) = \mathbf{1}$ for all conditions $b_j \in X_j$, $v_k(b_k) = \mathbf{1}$ for all conditions $b_k \in X_k$, $v_j(c) = \mathbf{1}$ and $v_k(c) = \mathbf{0}$. Since $X_i \subseteq X_j$ and $X_i \subseteq X_k$, it holds that $v_j(b_i) = v_k(b_i) = \mathbf{1}$ for all $b_i \in X_i$. Now. let $T_1 = \{v_j\}$ and $T_2 = \{v_k\}$ so that $T_1 \models \bigwedge_{b_i \in X_i} b_i \wedge c$ and $T_2 \models \bigwedge_{b_i \in X_i} b_i \wedge \neg c$. It is clearly the case that $T_1 \neq \emptyset$, $T_2 \neq \emptyset$ and $T_1, T_2 \subseteq T^c$. Therefore, it follows that $T^c \models \blacklozenge(\bigwedge_{b_i \in X_i} b_i \wedge c) \wedge \blacklozenge(\bigwedge_{b_i \in X_i} b_i \wedge \neg c)$. This means that $T^c \models \varphi^{X_i}_{\mathsf{con}}$.

To prove the right-to-left direction of (5), suppose that $T^c \models \varphi^{X_i}_{\mathsf{con}}$. This means that $T^c \models \blacklozenge(\bigwedge_{b_i \in X_i} b_i \wedge c) \wedge \blacklozenge(\bigwedge_{b_i \in X_i} b_i \wedge \neg c)$. It follows that there exist non-empty $T_1, T_2 \subseteq T^c$ such that $T_1 \models \bigwedge_{b_i \in X_i} b_i \wedge c$ and $T_2 \models \bigwedge_{b_i \in X_i} b_i \wedge \neg c$. Therefore, there exist two valuations $v_j, v_k \in T^c$ such that $v_j(b_i) = v_i(b_i) = \mathbf{1}$ for all $b_i \in X_i$, $v_j(c) = \mathbf{1}$ and $v_k(c) = \mathbf{0}$. By Definition 3.8, it follows that there exist $X_j, X_k \in \mathcal{D}$ such that $f(X_j) = \mathbf{1}$, $f(X_k) = \mathbf{0}$ and $X_i \subseteq X_j, X_k$. Due to Definition 2.6, it holds that X_i is contingent relative to f. $\square$

A.2 Proofs of Chapter 4

Fact 4.7. Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and T^c be a legal team on $B \cup \{c\}$. Further, let $a \in \mathcal{A}$ be a legal case and K_a the associated legal knowledge. Lastly, let T_{pos} be the positive extension of T^c by case a and T_{neg} be the negative extension of T^c by case a.

(1) $T_{\mathsf{pos}} \neq \emptyset$, $T_{\mathsf{pos}} \models \blacklozenge b_i$ and $T_{\mathsf{pos}} \models \bigwedge_{b_j \in B_i^{\perp}} \neg(b_i \wedge b_j)$ for each $b_i \in B$.

(2) $T_{\mathsf{neg}} \neq \emptyset$, $T_{\mathsf{neg}} \models \blacklozenge b_i$ and $T_{\mathsf{neg}} \models \bigwedge_{b_j \in B_i^{\perp}} \neg(b_i \wedge b_j)$ for each $b_i \in B$.

(3) If $B \cap K_a = \emptyset$, then neither T_{pos} nor T_{neg} is a legal team.

(4) Concept c is applicable to a relative to T^c if and only if T_{pos} is a legal team and T_{neg} is not a legal team.

(5) Concept c is not applicable to a relative to T^c if and only if T_{pos} is not a legal team and T_{neg} is a legal team.

(6) Case a is not determined by T^c and $B \cap K_a \neq \emptyset$ if and only if T_{pos} and T_{neg} are legal teams.

Proof. The statements are proven separately.

To prove (1), it needs to be shown that (a) $T_{\mathsf{pos}} \neq \emptyset$, (b) $T_{\mathsf{pos}} \models \blacklozenge b_i$ and (c) $T_{\mathsf{pos}} \models \bigwedge_{b_j \in B_i^{\perp}} \neg(b_i \wedge b_j)$ for each $b_i \in B$. Since T^c is a legal team and by Theorem 3.11, it follows that $T^c \neq \emptyset$. Therefore, $T_{\mathsf{pos}} \neq \emptyset$ because $T^c \subseteq T_{\mathsf{pos}}$. Furthermore, since T^c is a legal team and by Theorem 3.11, it holds that for each $b_i \in B$, there exists a non-empty subteam $T \subseteq T^c$ with $T \models b_i$. Since $T^c \subseteq T_{\mathsf{pos}}$, it follows that $T \subseteq T_{\mathsf{pos}}$. Thus, $T_{\mathsf{pos}} \models \blacklozenge b_i$ for each $b_i \in B$. Lastly, $\{v^{b+}_{K_a}\} \models \bigwedge_{b_j \in B_i^{\perp}} \neg(b_i \wedge b_j)$ for any $b_i \in B$, because K_a is legal knowledge and hence conflict-free. Additionally, $T^c \models \bigwedge_{b_j \in B_i^{\perp}} \neg(b_i \wedge b_j)$ for any $b_i \in B$ because of the fact that T^c is a legal team and Theorem 3.11. Since $\bigwedge_{b_j \in B_i^{\perp}} \neg(b_i \wedge b_j) \in \mathbf{PL}$, $T_{\mathsf{pos}} = T^c \cup \{v^{b+}_{K_a}\}$ and by flatness, it follows that $T_{\mathsf{pos}} \models \bigwedge_{b_j \in B_i^{\perp}} \neg(b_i \wedge b_j)$ for any $b_i \in B$.

Statement (2) follows by analogous reasoning to (1).

To prove (3), suppose that $B \cap K_a = \emptyset$. This means that $v^{c+}_{K_a}(b_i) = v^{c-}_{K_a}(b_i) = \mathbf{0}$ for all $b_i \in B$. This means that $\{v^{c+}_{K_a}\} \not\models \bigvee_{b_i \in B} b_i$ and $\{v^{c-}_{K_a}\} \not\models \bigvee_{b_i \in B} b_i$. Since $\bigvee_{b_i \in B} b_i \in \mathbf{PL}$ and by flatness, it follows that $T_{\mathsf{pos}} \not\models \bigvee_{b_i \in B} b_i$ and $T_{\mathsf{neg}} \not\models \bigvee_{b_i \in B} b_i$. According to Theorem 3.11, it holds that neither T_{pos} nor T_{neg} is a legal team.

For the left-to-right direction of (4), suppose that c is applicable to a relative to T^c. According to Definition 4.2, this means that there exists $v \in T^c$ such that $\{b \in B \mid v(b) = \mathbf{1}\} = B \cap K_a$ and $v(c) = \mathbf{1}$. By Definition 4.4, this implies that $v^{c+}_{K_a} = v$ so that $T^c = T_{\mathsf{pos}}$. Since T^c is a legal team, it follows that T_{pos} is a legal team. Further, it is the case that $v(b) = v^{c-}_{K_a}(b)$ for all $b \in B$ and $v(c) \neq v^{c-}_{K_a}(c)$ due to Definition 4.4. This means that $T_{\mathsf{neg}} \not\models =\!(b_1, \ldots, b_n, c)$ where $\{b_1, \ldots, b_n\} = B$. By Theorem 3.11, it follows that T_{neg} is not a legal team.

For the right-to-left direction of (4), suppose hat T_{pos} is a legal team and T_{neg} is not a legal team. By Definition 4.2, it is sufficient to show that there exists $v \in T^c$ such that $\{b \in B \mid v(b) = \mathbf{1}\} = B \cap K_a$ and $v(c) = \mathbf{1}$. Since T_{pos} is a legal team and by (3), it follows that $B \cap K_a \neq \emptyset$. Since T_{neg} is not a legal team, it follows that $T_{\mathsf{neg}} \not\models =\!(b_1, \ldots, b_n, c)$ where $\{b_1, \ldots, b_n\} = B$ due to Theorem 3.11 and (2). This means that there exists $v \in T^c$ such that $\{b \in B \mid v(b) = \mathbf{1}\} = B \cap K_a$ and $v(c) = \mathbf{1}$, as required.

The statement (5) follows by analogous reasoning to (4).

For the left-to-right direction of (6), suppose that a is not determined by T^c and $B \cap K_a \neq \emptyset$. By the former, it follows that there does not exist $v \in T^c$ such that $\{b \in B \mid v(b) = \mathbf{1}\} = B \cap K_a$. It follows that $T_{\mathsf{pos}} \models =\!(b_1, \ldots, b_n, c)$ and $T_{\mathsf{neg}} \models =\!(b_1, \ldots, b_n, c)$ where $\{b_1, \ldots, b_n\} = B$. By the latter, it follows that $\{v^{c+}_{K_a}\} \models \bigvee_{b_i \in B} b_i$ and $\{v^{c-}_{K_a}\} \models \bigvee_{b_i \in B} b_i$. Since T^c is a legal team, $\bigvee_{b_i \in B} b_i \in \mathbf{PL}$ and by flatness, it follows that $T_{\mathsf{pos}} \models \bigvee_{b_i \in B} b_i$ and $T_{\mathsf{neg}} \models \bigvee_{b_i \in B} b_i$. Due to (1), (2) and by Theorem 3.11, it follows that T_{pos} and T_{neg} are legal teams.

For the right-to-left direction of (6), suppose that T_{pos} and T_{neg} are legal teams. By (3), it follows that $B \cap K_a \neq \emptyset$. Since T_{pos} and T_{neg} are legal teams and by Theorem 3.11, $T_{\mathsf{pos}} \models =\!(b_1, \ldots, b_n, c)$ and $T_{\mathsf{neg}} \models =\!(b_1, \ldots, b_n, c)$ where $\{b_1, \ldots, b_n\} = B$. This implies that there does not exist $v \in T^c$ such that $\{b \in B \mid v(b) = \mathbf{1}\} = B \cap K_a$. By Definition 4.2, it follows that a is not determined by T^c. □

Fact 4.9. Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and T^c be a legal team on $B \cup \{c\}$. Further, let $a \in \mathcal{A}$ be a legal case, K_a the associated legal knowledge and T be a constant extension of T^c by a. For any legal possibility X_i expressed by some $v_i \in T^c$ and the decisive legal possibility X_i, if $T \models \psi$, then $T^c \models \psi$ for $\psi \in \{\varphi^{X_i}_{\mathsf{suf}}, \varphi^{X_i}_{\mathsf{exc}}, \varphi^{X_i}_{\mathsf{nec}+}, \varphi^{X_i}_{\mathsf{nec}-}, \varphi^{X_i}_{\mathsf{con}}\}$.

Proof. Let X_i be an arbitrary legal possibility expressed by some $v_i \in T^c$. It needs to be shown that if $T \models \psi$, then $T^c \models \psi$ for $\psi \in \{\varphi^{X_i}_{\mathsf{suf}}, \varphi^{X_i}_{\mathsf{exc}}, \varphi^{X_i}_{\mathsf{nec}+}, \varphi^{X_i}_{\mathsf{nec}-}, \varphi^{X_i}_{\mathsf{con}}\}$. There are two cases to be considered: (1) $\psi \in \{\varphi^{X_i}_{\mathsf{suf}}, \varphi^{X_i}_{\mathsf{exc}}, \varphi^{X_i}_{\mathsf{nec}+}, \varphi^{X_i}_{\mathsf{nec}-}\}$ and (2) $\psi := \varphi^{X_i}_{\mathsf{con}}$.

First, let $\psi \in \{\varphi^{X_i}_{\mathsf{suf}}, \varphi^{X_i}_{\mathsf{exc}}, \varphi^{X_i}_{\mathsf{nec}+}, \varphi^{X_i}_{\mathsf{nec}-}, \varphi^{X_i}_{\mathsf{con}}\}$. Suppose that $T \models \psi$. Observe that $\psi \in \mathbf{PL}$. Since $T^c \subseteq T$ and by the downward closure property, it follows that $T^c \models \psi$.

Second, $\psi := \varphi^{X_i}_{\mathsf{con}}$. Suppose $T \models \varphi^{X_i}_{\mathsf{con}}$. By Fact 3.15, $T \not\models \varphi^{X_i}_{\mathsf{suf}}$ and $T \not\models \varphi^{X_i}_{\mathsf{exc}}$. Since T is a constant extension of T^c and due to Definition 4.8, it is the case that $T^c \not\models \varphi^{X_i}_{\mathsf{suf}}$ and $T^c \not\models \varphi^{X_i}_{\mathsf{exc}}$. By Fact 3.15, it follows that $T^c \models \varphi^{X_i}_{\mathsf{con}}$. □

Fact 4.11. Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions, T^c be a legal team on $B \cup \{c\}$ and $a \in \mathcal{A}$ be a legal case. If a is determined by T^c, then

(1) c is applicable to a if and only if c is predicted to be applicable to a.

(2) c is not applicable to a if and only if c is predicted to be not applicable to a.

Proof. To prove (1) suppose that a is determined by T^c. For the left-to-right direction of the biimplication, further assume that c is applicable to a. This means that there exists a valuation $v \in T^c$ such that $\{b \in B \mid v(b) = \mathbf{1}\} = B \cap K_a$. By Definition 4.4, $v^{c+}_{K_a} = v$ and hence $T^c \cup \{v^{c+}_{K_a}\} = T^c$. It follows that $T^c \cup \{v^{c+}_{K_a}\}$ is a constant extension due to Definition 4.8. Further, by Fact 4.7, it follows that $T^c \cup \{v^{c-}_{K_a}\}$ is not a legal team. Therefore, $T^c \cup \{v^{c-}_{K_a}\}$ is not a constant extension due to Definition 4.8. According to Definition 4.10, c is predicted to be applicable to a.

To prove the right-to-left direction of the biimplication, suppose that c is predicted to be applicable to a and that a is determined by T^c. Due to the latter and Definition 4.2, it follows that there exists a valuation $v \in T^c$ such that $\{b \in B \mid v(b) = \mathbf{1}\} = B \cap K_a$. It is sufficient to show that $v(c) = \mathbf{1}$. Since c is predicted to be applicable to a and by Definition 4.10, $T^c \cup \{v^{c+}_{K_a}\}$ is a constant extension. According to Definition 4.8, this means that $T^c \cup \{v^{c+}_{K_a}\}$ is a legal team. By Theorem 3.11, it follows that $T^c \cup \{v^{c+}_{K_a}\} \models =\!(b_1, \ldots, b_n, c)$ for $B = \{b_1, \ldots, b_n\}$. Further, by Definition 4.4, it follows that $v^{c+}_{K_a}(c) = \mathbf{1}$, $v^{c+}_{K_a}(b_j) = \mathbf{1}$ for all $b_j \in B \cap K_a$ and $v^{c+}_{K_a}(b_k) = \mathbf{0}$ for all $b_k \in B \setminus K_a$. Therefore, for all

$v_i \in T^c$ with $v_i(b_j) = \mathbf{1}$ for all $b_j \in B \cap K_a$ and $v_i(b_k) = \mathbf{0}$ for all $b_k \in B \setminus K_a$, it holds that $v_i(c) = \mathbf{1}$. Since $\{b \in B \mid v(b) = \mathbf{1}\} = B \cap K_a$ and hence $\{b \in B \mid v(b) = \mathbf{0}\} = B \setminus K_a$, it follows that $v(c) = \mathbf{1}$.

Statement (2) follows by analogous reasoning to (1). □

Proposition 4.13. Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions, T^c be a legal team on $B \cup \{c\}$ and X a legal possibility expressed by some valuation of T^c. Further, let $a \in \mathcal{A}$ be a legal case and K_a the associated legal knowledge. For any $\psi \in \{\varphi^X_{\mathsf{suf}}, \varphi^X_{\mathsf{exc}}, \varphi^X_{\mathsf{nec+}}, \varphi^X_{\mathsf{nec-}}, \varphi^X_{\mathsf{con}}\}$,

(1a) $\underset{\text{pos}}{\checkmark}$ denotes that if $T^c \models \psi$, then $T_{\mathsf{pos}} \models \psi$

(1b) $\underset{\text{neg}}{\checkmark}$ denotes that if $T^c \models \psi$, then $T_{\mathsf{neg}} \models \psi$

(2a) $\underset{\text{pos}}{\text{✗}}$ denotes that if $T^c \models \psi$, then $T_{\mathsf{pos}} \not\models \psi$

(2b) $\underset{\text{neg}}{\text{✗}}$ denotes that if $T^c \models \psi$, then $T_{\mathsf{neg}} \not\models \psi$

where $T_{\mathsf{pos}} = T^c \cup \{v^{c+}_{K_a}\}$ is the positive extension of T^c by a, $T_{\mathsf{neg}} = T^c \cup \{v^{c-}_{K_a}\}$ is the negative extension of T^c by a and X is a legal possibility expressed by some $v \in T^c$.

Table A.1 specifies whether $\psi \in \{\varphi^X_{\mathsf{suf}}, \varphi^X_{\mathsf{exc}}, \varphi^X_{\mathsf{nec+}}, \varphi^X_{\mathsf{nec-}}, \varphi^X_{\mathsf{con}}\}$ remains satisfied by the positive or negative extension of T^c by a, depending on whether $X \subseteq B \cap K_a$ or $X \not\subseteq B \cap K_a$, assuming that $T^c \models \psi$.

	$\psi := \varphi^X_{\mathsf{suf}}$		$\psi := \varphi^X_{\mathsf{exc}}$		$\psi := \varphi^X_{\mathsf{nec+}}$		$\psi := \varphi^X_{\mathsf{nec-}}$		$\psi := \varphi^X_{\mathsf{con}}$	
$X \subseteq B \cap K_a$	✓ pos	✗ neg	✗ pos	✓ neg	✓ pos	✓ neg	✓ pos	✓ neg	✓ pos	✓ neg
$X \not\subseteq B \cap K_a$	✓ pos	✓ neg	✓ pos	✓ neg	✗ pos	✓ neg	✓ pos	✗ neg	✓ pos	✓ neg

Table A.1: Preservation of satisfaction for ψ after extending by a

Proof. The statements are proven separately.

(1) Suppose that $T^c \models \varphi^X_{\mathsf{suf}}$. This means that $T^c \models (\bigwedge_{b \in X} b) \rightarrow c$ and $T^c \models \neg(\bigwedge_{b \in X} b) \vee c$. Therefore, there exist $T_1, T_2 \subseteq T^c$ such that $T_1 \cup T_2 = T^c$, $T_1 \models \neg(\bigwedge_{b \in X} b)$ and $T_2 \models c$. Thus, for all $v \in T_1$, $\{v\} \not\models \bigwedge_{b \in X} b$.

 (1a) Let $X \subseteq B \cap K_a$. To show that $T_{\mathsf{pos}} \models \varphi^X_{\mathsf{suf}}$, it is sufficient to show that $\{v^{c+}_{K_a}\} \not\models \bigwedge_{b \in X} b$ or $\{v^{c+}_{K_a}\} \models c$. By Definition 4.4, $\{v^{c+}_{K_a}\} \models c$.

 To show that $T_{\mathsf{neg}} \not\models \varphi^X_{\mathsf{suf}}$, it is sufficient to show that $\{v^{c-}_{K_a}\} \models \bigwedge_{b \in X} b$ and $\{v^{c-}_{K_a}\} \not\models c$. By Definition 4.4, $\{v^{c-}_{K_a}\} \models \bigwedge_{b \in X} b$ and $\{v^{c-}_{K_a}\} \not\models c$.

 (1b) Let $X \not\subseteq B \cap K_a$. To show that $T_{\mathsf{pos}} \models \varphi^X_{\mathsf{suf}}$ and $T_{\mathsf{neg}} \models \varphi^X_{\mathsf{suf}}$, it is sufficient to show that $\{v^{c+}_{K_a}\} \not\models \bigwedge_{b \in X} b$ or $\{v^{c+}_{K_a}\} \models c$ and $\{v^{c-}_{K_a}\} \not\models \bigwedge_{b \in X} b$ or $\{v^{c-}_{K_a}\} \models c$. Since $X \not\subseteq B \cap K_a$, there exists $b_i \in X$ such that $b_i \notin B \cap K_a$. By Definition 4.4, $\{v^{c+}_{K_a}\} \not\models \bigwedge_{b \in X} b$ and $\{v^{c-}_{K_a}\} \not\models \bigwedge_{b \in X} b$.

(2) Suppose that $T^c \models \varphi^X_{\mathsf{exc}}$. This means that $T^c \models (\bigwedge_{b \in X} b) \rightarrow \neg c$ and $T^c \models \neg(\bigwedge_{b \in X} b) \vee \neg c$. Therefore, there exist $T_1, T_2 \subseteq T^c$ such that $T_1 \cup T_2 = T^c$, $T_1 \models \neg(\bigwedge_{b \in X} b)$ and $T_2 \models \neg c$. Thus, for all $v \in T_1$, $\{v\} \not\models \bigwedge_{b \in X} b$. Further, for all $v \in T_2$, $\{v\} \not\models c$.

 (2a) Let $X \subseteq B \cap K_a$. To show that $T_{\mathsf{pos}} \not\models \varphi^X_{\mathsf{exc}}$, it is sufficient to show that $\{v^{c+}_{K_a}\} \models \bigwedge_{b \in X} b$ and $\{v^{c+}_{K_a}\} \models c$. By Definition 4.4, $\{v^{c+}_{K_a}\} \models c$ and $\{v^{c+}_{K_a}\} \models \bigwedge_{b \in X} b$.

To show that $T_{\mathsf{neg}} \models \varphi^X_{\mathsf{exc}}$, it is sufficient to show that $\{v^{c-}_{K_a}\} \not\models \bigwedge_{b\in X} b$ or $\{v^{c-}_{K_a}\} \not\models c$. By Definition 4.4, $\{v^{c-}_{K_a}\} \not\models c$.

(2b) Let $X \not\subseteq B \cap K_a$. To show that $T_{\mathsf{pos}} \models \varphi^X_{\mathsf{exc}}$ and $T_{\mathsf{neg}} \models \varphi^X_{\mathsf{exc}}$, it is sufficient to show that $\{v^{c+}_{K_a}\} \not\models \bigwedge_{b\in X} b$ or $\{v^{c+}_{K_a}\} \not\models c$ and $\{v^{c-}_{K_a}\} \not\models \bigwedge_{b\in X} b$ or $\{v^{c-}_{K_a}\} \not\models c$. Since $X \not\subseteq B\cap K_a$, there exists $b_i \in X$ such that $b_i \notin B \cap K_a$. Consequently, $\{v^{c+}_{K_a}\} \not\models \bigwedge_{b\in X} b$ and $\{v^{c-}_{K_a}\} \not\models \bigwedge_{b\in X} b$.

(3) Suppose that $T^c \models \varphi^X_{\mathsf{nec+}}$. This means that $T^c \models c \to (\bigwedge_{b\in X} b)$ and $T^c \models \neg c \vee (\bigwedge_{b\in X} b)$. Therefore, there exist $T_1, T_2 \subseteq T^c$ such that $T_1 \cup T_2 = T^c$, $T_1 \models \neg c$ and $T_2 \models \bigwedge_{b\in X} b$. Further, for all $v \in T_1$, $\{v\} \not\models c$.

(3a) Let $X \subseteq B \cap K_a$. To show that $T_{\mathsf{pos}} \models \varphi^X_{\mathsf{nec+}}$ and $T_{\mathsf{neg}} \models \varphi^X_{\mathsf{nec+}}$, it is sufficient to show that $\{v^{c+}_{K_a}\} \models \bigwedge_{b\in X} b$ or $\{v^{c+}_{K_a}\} \not\models c$ and $\{v^{c-}_{K_a}\} \models \bigwedge_{b\in X} b$ or $\{v^{c-}_{K_a}\} \not\models c$. By Definition 4.4, $\{v^{c+}_{K_a}\} \models \bigwedge_{b\in X} b$ and $\{v^{c-}_{K_a}\} \models \bigwedge_{b\in X} b$.

(3b) Let $X \not\subseteq B\cap K_a$. This means that there exists $b_i \in X$ such that $b_i \notin B \cap K_a$. To show that $T_{\mathsf{pos}} \not\models \varphi^X_{\mathsf{nec+}}$, it is sufficient to show that $\{v^{c+}_{K_a}\} \not\models \bigwedge_{b\in X} b$ and $\{v^{c+}_{K_a}\} \models c$. By Definition 4.4, $\{v^{c+}_{K_a}\} \models c$. Since there exists $b_1 \in X$ such that $b_i \notin B \cap K_a$, $\{v^{c+}_{K_a}\} \not\models \bigwedge_{b\in X} b$ due to Definition 4.4.

To show that $T_{\mathsf{neg}} \models \varphi^X_{\mathsf{nec+}}$, it is sufficient to show that $\{v^{c-}_{K_a}\} \not\models c$ or $\{v^{c-}_{K_a}\} \models \bigwedge_{b\in X} b$. By Definition 4.4, $\{v^{c-}_{K_a}\} \not\models c$.

(4) Suppose that $T^c \models \varphi^X_{\mathsf{nec-}}$. This means that $T^c \models \neg c \to (\bigwedge_{b\in X} b)$ and $T^c \models \neg\neg c \vee (\bigwedge_{b\in X} b)$. Therefore, there exist $T_1, T_2 \subseteq T^c$ such that $T_1 \cup T_2 = T^c$, $T_1 \models \neg\neg c$ and $T_2 \models \bigwedge_{b\in X} b$. For any $v \in T_1$, it holds that $\{v\} \not\models \neg c$ and therefore $\{v\} \models c$.

(4a) Let $X \subseteq B \cap K_a$. To show that $T_{\mathsf{pos}} \models \varphi^X_{\mathsf{nec-}}$ and $T_{\mathsf{neg}} \models \varphi^X_{\mathsf{nec-}}$, it is sufficient to show that $\{v^{c+}_{K_a}\} \models \bigwedge_{b\in X} b$ or $\{v^{c+}_{K_a}\} \models c$ and $\{v^{c-}_{K_a}\} \models \bigwedge_{b\in X} b$ or $\{v^{c-}_{K_a}\} \models c$. By Definition 4.4, $\{v^{c+}_{K_a}\} \models \bigwedge_{b\in X} b$ and $\{v^{c-}_{K_a}\} \models \bigwedge_{b\in X} b$.

(4b) Let $X \not\subseteq B\cap K_a$. This means that there exists $b_i \in X$ such that $b_i \notin B \cap K_a$. To show that $T_{\mathsf{pos}} \models \varphi^X_{\mathsf{nec-}}$, it is sufficient to show that $\{v^{c+}_{K_a}\} \models c$ or $\{v^{c+}_{K_a}\} \models \bigwedge_{b\in X} b$. By Definition 4.4, $\{v^{c+}_{K_a}\} \models c$.

To show that $T_{\mathsf{neg}} \not\models \varphi^X_{\mathsf{nec-}}$, it is sufficient to show that $\{v^{c-}_{K_a}\} \not\models c$ and $\{v^{c-}_{K_a}\} \not\models \bigwedge_{b\in X} b$. By Definition 4.4, $\{v^{c-}_{K_a}\} \not\models c$. Since there exists $b_i \in X$ such that $b_i \notin B\cap K_a$, $\{v^{c-}_{K_a}\} \not\models \bigwedge_{b\in X} b$ due to Definition 4.4.

(5) Suppose that $T \models \varphi^X_{\mathsf{con}}$. This means that there exist $v_i, v_j \in T^c$ such that $v_i(b) = v_j(b) = \mathbf{1}$ for all $b \in X$, $v_i(c) = \mathbf{1}$ and $v_j(c) = \mathbf{0}$. Since $T^c \subseteq T_{\mathsf{pos}}$ and $T^c \subseteq T_{\mathsf{neg}}$, it follows that $v_i, v_j \in T_{\mathsf{pos}}$ and $v_i, v_j \in T_{\mathsf{neg}}$. Therefore, $T_{\mathsf{pos}} \models \varphi^X_{\mathsf{con}}$ and $T_{\mathsf{neg}} \models \varphi^X_{\mathsf{con}}$.

□

Fact A.1. Let $c \in \mathcal{C}$ be a concept and $B \subseteq \mathcal{B}$ be a non-empty set of conditions.

(1) There does not exist T^c on $B \cup \{c\}$ such that $T^c \models \varphi^{X_i}_{\mathsf{suf}}$ and $T^c \models \varphi^{X_j}_{\mathsf{nec+}}$ where X_i and X_j are legal possibilities expressed by some valuations $v_i, v_j \in T^c$ and $X_j \not\subseteq X_i$.

(2) There does not exist T^c on $B \cup \{c\}$ such that $T^c \models \varphi^{X_i}_{\mathsf{exc}}$ and $T^c \models \varphi^{X_j}_{\mathsf{nec-}}$ where X_i and X_j are legal possibilities expressed by some valuations $v_i, v_j \in T^c$ and $X_j \not\subseteq X_i$.

Proof. To prove (1), suppose for reductio that there exists a concept $c \in \mathcal{C}$, non-empty set of conditions

$B \subseteq \mathcal{B}$ and a legal team T^c on $B \cup \{c\}$ such that $T^c \models \varphi_{\mathsf{suf}}^{X_i}$ and $T^c \models \varphi_{\mathsf{nec+}}^{X_j}$ where X_i and X_j are legal possibilities expressed by some valuations $v_i, v_j \in T^c$. Further, assume that $X_j \not\subseteq X_i$. Since $T^c \models \varphi_{\mathsf{nec+}}^{X_j}$, $T^c \models c \rightarrow (\bigwedge_{b \in X_j} b)$ and $T^c \models \neg c \vee (\bigwedge_{b \in X_j} b)$. This means that there exist $T_1, T_2 \subseteq T^c$ such that $T_1 \cup T_2 = T^c$, $T_1 \models \neg c$ and $T_2 \models \bigwedge_{b \in X_j} b$. Thus, for all $v \in T_1$, $\{v\} \not\models c$. The fact that $T^c \models \varphi_{\mathsf{suf}}^{X_i}$ implies that $T^c \models (\bigwedge_{b \in X_i} b) \rightarrow c$ and $T^c \models \neg(\bigwedge_{b \in X_i} b) \vee c$. Therefore, there exist $T_3, T_4 \subseteq T^c$ such that $T_3 \cup T_4 = T^c$, $T_3 \models \neg(\bigwedge_{b \in X_i} b)$ and $T_4 \models c$. Therefore, for all $v \in T_3$, $\{v\} \not\models \bigwedge_{b \in X_i} b$.

Since $X_j \not\subseteq X_i$, there exists some $b \in B$ such that $b \in X_j$ and $b \notin X_i$. It follows that $\{v_i\} \not\models \bigwedge_{b \in X_j} b$. Thus, $v_i \notin T_2$. As $T^c = T_1 \cup T_2$, it holds that $v_i \in T_1$. Therefore, $\{v_i\} \not\models c$ and $v_i \notin T_4$. Since $T^c = T_3 \cup T_4$, this implies that $v_i \in T_3$. This means that $\{v_i\} \not\models \bigwedge_{b \in X_i} b$ so that the legal possibility X_i is not expressed by v_i. This is a contradiction.

Statement (2) follows by analogous reasoning to (1). $\square$

Fact A.2. Let $c \in \mathcal{C}$ be arbitrary. There exists some non-empty $B \subseteq \mathcal{B}$, legal team T^c on $B \cup \{c\}$ and legal possibility $Z \subseteq B$ such that

(1) $T^c \models \varphi_{\mathsf{suf}}^{X_i}$ and $T^c \models \varphi_{\mathsf{exc}}^{X_j}$ where X_i and X_j are legal possibilities expressed by some valuations $v_i, v_j \in T^c$ and $X_i \subseteq Z$ and $X_j \subseteq Z$.

(2) $T^c \models \varphi_{\mathsf{nec+}}^{X_i}$ and $T^c \models \varphi_{\mathsf{nec-}}^{X_j}$ where X_i and X_j are legal possibilities expressed by some valuations $v_i, v_j \in T^c$ and $X_i \not\subseteq Z$ and $X_j \not\subseteq Z$.

Proof. Let $c \in \mathcal{C}$ be arbitrary. To prove (1) and (2), it is sufficient to provide examples.

For (1), let $B = \{b_1, b_2\}$, $B_1^{\perp} = B_2^{\perp} = \emptyset$ and $Z = \{b_1, b_2\}$. Further, let T^c on $B \cup \{c\}$ be defined as in Table A.2. It is easy to verify that T^c is a legal team. Further, $X_1 \subseteq Z$ and $X_2 \subseteq Z$. To see

valuation v_i	legal possibility X_i	b_1	b_2	concept c
v_1	$\{b_1\}$	**1**	**0**	**1**
v_2	$\{b_2\}$	**0**	**1**	**0**

Table A.2: Legal team for Fact A.2

that $T^c \models \varphi_{\mathsf{suf}}^{X_1}$, let $T_1 = \{v_1\}$ and $T_2 = \{v_2\}$. Clearly, $T_1 \models c$, $T_2 \models \neg(\bigwedge_{b \in X_1} b)$ and $T^c = T_1 \cup T_2$. This means that $T^c \models \neg(\bigwedge_{b \in X_1} b) \vee c$ and hence $T^c \models (\bigwedge_{b \in X_1} b) \rightarrow c$. Consequently, $T^c \models \varphi_{\mathsf{suf}}^{X_1}$. By analogous reasoning, it follows that $T^c \models \varphi_{\mathsf{exc}}^{X_2}$.

For (2), let $B = \{b_1, b_2, b_3\}$, $B_1^{\perp} = B_2^{\perp} = B_3^{\perp} = \emptyset$ and $Z = \{b_3\}$. Further, T^c on $B \cup \{c\}$ be defined as in Table A.3. It is easy to verify that T^c is a legal team. Further, $X_1 \not\subseteq Z$ and $X_2 \not\subseteq Z$. To see

valuation v_i	legal possibility X_i	b_1	b_2	b_3	concept c
v_1	$\{b_1\}$	**1**	**0**	**0**	**1**
v_2	$\{b_2, b_3\}$	**0**	**1**	**1**	**0**

Table A.3: Legal team for Fact A.2

that $T^c \models \varphi_{\mathsf{nec}}^{X_1}$, let $T_1 = \{v_2\}$ and $T_2 = \{v_1\}$. Clearly, $T_1 \models \neg c$, $T_2 \models (\bigwedge_{b \in X_1} b)$ and $T^c = T_1 \cup T_2$. This means that $T^c \models \neg c \vee (\bigwedge_{b \in X_1} b)$ and hence $T^c \models c \rightarrow (\bigwedge_{b \in X_1} b)$. Consequently, $T^c \models \varphi_{\mathsf{nec+}}^{X_1}$. By analogous reasoning, it follows that $T^c \models \varphi_{\mathsf{nec-}}^{X_2}$. $\square$

A.3 Proofs of Chapter 5

Fact 5.3. Let $c \in \mathcal{C}$ be a legal concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and T^c be a legal team on $B \cup \{c\}$. Further, let $a \in \mathcal{A}$ be a case such that K_a is legal knowledge and c is predictable for a with respect to T^c.

(1) If the prediction that c is applicable to a is strongly unstable, then there does not exist a legal possibility X expressed by some $v \in T^c$ such that $T^c \models \varphi^X_{\mathsf{suf}}$ and $X \subseteq B \cap K_a$.

(2) If the prediction that c is not applicable to a is strongly unstable, then there does not exist a legal possibility X expressed by some $v \in T^c$ such that $T^c \models \varphi^X_{\mathsf{exc}}$ for c and $X \subseteq B \cap K_a$.

Proof. Statement (1) is proven by contraposition. Suppose that there does exist a legal possibility X expressed by some valuation $v \in T^c$ such that $T^c \models \varphi^X_{\mathsf{suf}}$ and $X \subseteq B \cap K_a$. Let $a_i \in \mathcal{A}$ be an arbitrary case such that K_{a_i} is legal knowledge and $B \cap K_a \subseteq B \cap K_{a_i}$. Therefore, $X \subseteq B \cap K_{a_i}$. Due to Proposition 4.13, it holds that $T^c \cup \{v^{c-}_{K_{a_i}}\} \not\models \varphi^X_{\mathsf{suf}}$. By Definition 4.10, it follows that c is predicted to be applicable to a_i or c is unpredictable for a_i. Since a_i was arbitrary, it follows that for any $a_j \in \mathcal{A}$ such that K_{a_j} is legal knowledge and $B \cap K_a \subseteq B \cap K_{a_j}$, c is predicted to be applicable to a_j or c is unpredictable for a_j. By Definition 5.1, the prediction of c for a with respect to T^c is not strongly unstable.

Statement (2) follows by analogous reasoning to (1). □

Proposition 5.4. Let $c \in \mathcal{C}$ be a legal concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and T^c be a legal team on $B \cup \{c\}$. Further, let $a \in \mathcal{A}$ be a case such that K_a is legal knowledge and c is predictable for a with respect to T^c.

(1) If $T^c \models \blacklozenge \left(\bigwedge_{b \in B \cap K_a} b \wedge \neg c\right)$, then the prediction that c is applicable to a is strongly unstable.

(2) If the prediction that c is applicable to a is strongly unstable and $T^c \not\models \blacklozenge \left(\bigwedge_{b \in B \cap K_a} b \wedge \neg c\right)$, then legal possibilities Y_i and X_j exist, expressed by some valuations $v_i, v_j \in T^c$, such that $T^c \models \varphi^{Y_i}_{\mathsf{nec}-} \wedge \varphi^{X_j}_{\mathsf{exc}}$, $Y_i \not\subseteq (B \cap K_a)$ and $Y_i \cup X_j \cup (B \cap K_a)$ is conflict-free.

(3) If $T^c \models \blacklozenge \left(\bigwedge_{b \in B \cap K_a} b \wedge c\right)$, then the prediction that c is not applicable to a is strongly unstable.

(4) If the prediction that c is not applicable to a is strongly unstable and $T^c \not\models \blacklozenge \left(\bigwedge_{b \in B \cap K_a} b \wedge c\right)$, then legal possibilities Y_i and X_j exist, expressed by some valuations $v_i, v_j \in T^c$, such that $T^c \models \varphi^{Y_i}_{\mathsf{nec}+} \wedge \varphi^{X_j}_{\mathsf{suf}}$, $Y_i \not\subseteq (B \cap K_a)$, and $Y_i \cup X_j \cup (B \cap K_a)$ is conflict-free.

Proof. To prove statement (1), assume that $T^c \models \blacklozenge(\bigwedge_{b \in B \cap K_a} b \wedge \neg c)$. Hence, there exists $v \in T^c$ such that $\{v\} \models (\bigwedge_{B \cap K_a} b \wedge \neg c)$, which means that $v(c) = \mathbf{0}$. Note that $X = \{b \in B \mid v(b) = \mathbf{1}\}$ is a legal possibility due to Theorem 3.11. Let $a_i \in \mathcal{A}$ be a case such that $K_{a_i} = X$. It follows that K_{a_i} is legal knowledge. By Definition 4.2, a_i is determined by T^c and c is not applicable to a_i. By Fact 4.11, c is predicted to be not applicable to a_i. As $\{v\} \models (\bigwedge_{B \cap K_a} b \wedge \neg c)$, it holds that $B \cap K_a \subseteq X$ and hence $B \cap K_a \subseteq B \cap K_{a_i}$. By Definition 5.1, the prediction that c is applicable to a is strongly unstable.

To prove statement (2), suppose that c is predicted to be applicable to a, this prediction is strongly unstable, and $T^c \not\models \blacklozenge(\bigwedge_{b \in B \cap K_a} b \wedge \neg c)$.

Let us first show that that there exists $v_i \in T^c$ such that $T^c \models \varphi^{Y_i}_{\mathsf{nec}-}$ and $Y_i \not\subseteq B \cap K_a$, where Y_i is the legal possibility expressed by some $v_i \in T^c$.

For this purpose, let us show that a is not determined by T^c. To show this, suppose for reductio that a is determined by T^c. This means that there exists $v \in T^c$ such that $\{b \in B \mid v(b) = \mathbf{1}\} = B \cap K_a$. Since c is predicted to be applicable, it is the case that c is applicable to a due to Fact 4.11. By Definition 4.2, it follows that $v(c) = \mathbf{1}$. Further, as c is predicted to be applicable to a and this prediction is strongly unstable, it follows by Fact 5.3 that there does not exist a legal possibility X_s expressed by some $v_s \in T^c$ such that $T^c \models \varphi^{X_s}_{\mathsf{suf}}$ and $X_s \subseteq B \cap K_a$. It follows that $T^c \not\models \varphi^{B \cap K_a}_{\mathsf{suf}}$. By Proposition 3.14, it follows that there exists some valuation $v_k \in T^c$ such that $\{v_k\} \models (\bigwedge_{b \in B \cap K_a} b \wedge \neg c)$. This means that $T^c \models \blacklozenge(\bigwedge_{b \in B \cap K_a} b \wedge \neg c)$. This is a contradiction. Therefore, a is not determined by T^c.

Due to the fact that c is predicted to be applicable to a, a is not determined by T^c, and due to Definition 4.2, Proposition 4.13, and Theorem 4.16, it follows that there exists $v_i \in T^c$ such that $T^c \models \varphi^{Y_i}_{\mathsf{nec}-}$ and $Y_i \not\subseteq B \cap K_a$, where Y_i is the legal possibility expressed by v_i, or there exists $v_s \in T^c$ such that $T^c \models \varphi^{X_s}_{\mathsf{suf}}$ and $X_s \subseteq B \cap K_a$, where X_s is the legal possibility expressed by $v_s \in T^c$. Recall that, since c is predicted to be applicable to a and this prediction is strongly unstable, it follows by Fact 5.3 that the latter is not the case. It follows that there exists $v_i \in T^c$ such that $T^c \models \varphi^{Y_i}_{\mathsf{nec}-}$ and $Y_i \not\subseteq B \cap K_a$, where Y_i is the legal possibility expressed by $v_i \in T^c$.

Since the prediction that c is applicable to a is strongly unstable, there exists $a_i \in \mathcal{A}$ such that K_{a_i} is legal knowledge, $B \cap K_a \subseteq B \cap K_{a_i}$, and c is predicted to be not applicable to a_i. By Definition 4.2 and Proposition 4.13, it follows that $Y_i \subseteq B \cap K_{a_i}$. Furthermore, since $T^c \not\models \blacklozenge \left(\bigwedge_{b \in B \cap K_a} b \wedge \neg c\right)$, it follows by Fact 4.11 that a_i is not determined by T^c.

Now let us show that that there exists $v_j \in T^c$ such that $T^c \models \varphi^{X_j}_{\mathsf{exc}}$ and $X_j \subseteq B \cap K_{a_i}$, where X_j is the legal possibility expressed by $v_j \in T^c$.

For this purpose, let us show that there does not exist $v_l \in T^c$ such that $T^c \models \varphi^{Y_l}_{\mathsf{nec}+}$ and $Y_l \not\subseteq B \cap K_{a_i}$, where Y_l is the legal possibility expressed by v_l. To prove this, suppose for reductio that such a v_l does exist. Since $B \cap K_a \subseteq B \cap K_{a_i}$, it follows that $Y_l \not\subseteq B \cap K_a$. By Definition 4.2 and Proposition 4.13, it is the case that c is not predicted to be applicable to a. This is a contradiction. Thus, there does not exist $v_l \in T^c$ such that $T^c \models \varphi^{Y_l}_{\mathsf{nec}+}$ and $Y_l \not\subseteq B \cap K_{a_i}$ where Y_l is the legal possibility expressed by v_l.

Due to the fact that c is predicted to be applicable to a_i, a_i is not predicted by T^c, and due to Definition 4.2, Proposition 4.13, and Theorem 4.16, it follows that there exists $v_l \in T^c$ such that $T^c \models \varphi^{Y_l}_{\mathsf{nec}+}$ and $Y_l \not\subseteq B \cap K_a$, where Y_l is the legal possibility expressed by $v_l \in T^c$, or there exists $v_j \in T^c$ such that $T^c \models \varphi^{X_j}_{\mathsf{exc}}$ and $X_j \subseteq B \cap K_{a_i}$, where X_j is the legal possibility expressed by $v_j \in T^c$. Since the former is not the case, it follows that there exists $v_j \in T^c$ such that $T^c \models \varphi^{X_j}_{\mathsf{exc}}$ and $X_j \subseteq B \cap K_{a_i}$, where X_j is the legal possibility expressed by $v_j \in T^c$.

This means that there exist legal possibilities Y_i and X_j expressed by some valuations v_i and v_j such that $T^c \models \varphi^{Y_i}_{\mathsf{nec}-} \wedge \varphi^{X_j}_{\mathsf{exc}}$. As previously stated, $Y_i \not\subseteq B \cap K_a$. By the fact that $Y_i \subseteq B \cap K_{a_i}$, $X_j \subseteq B \cap K_{a_i}$, $B \cap K_a \subseteq B \cap K_{a_i}$, K_{a_i} is legal knowledge, and hence K_{a_i} is conflict-free, it follows that $Y_i \cup X_j \cup (B \cap K_a)$ is conflict-free.

Statements (3) and (4) follow by reasoning analogous to the proof of statements (1) and (2). □

Proposition 5.6. Let $c \in \mathcal{C}$ be a legal concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and T^c be a legal team on $B \cup \{c\}$. Further, let $a \in \mathcal{A}$ be a case such that K_a is legal knowledge and c is predictable for a with respect to T^c. The prediction that a is applicable (not applicable) to a is weakly unstable if and only if there exists $a_i \in \mathcal{A}$ such that K_{a_i} is legal knowledge and $B \cap K_a \subseteq B \cap K_{a_i}$ and further that one of the following conditions is satisfied:

(1) c is indeterminably unpredictable for a_i

(2) $T^c \models \varphi_{\mathsf{suf}}^{X_i} \wedge \varphi_{\mathsf{exc}}^{X_j}$ and $X_i, X_j \subseteq B \cap K_{a_i}$ where X_i and X_j are legal possibilities expressed by some $v_i, v_j \in T^c$

(3) $T^c \models \varphi_{\mathsf{nec}+}^{Y_i} \wedge \varphi_{\mathsf{nec}-}^{Y_j}$ and $T^c \models \varphi_{\mathsf{exc}}^{X_j}$ and $Y_i, Y_j \not\subseteq B \cap K_{a_i}$ where Y_i and Y_j are legal possibilities expressed by some $v_i, v_j \in T^c$

Proof. For the left-to-right direction, suppose that the prediction of c for a with respect to T^c is weakly unstable. By Definition 5.1, this means that there exists $a_i \in \mathcal{A}$ such that K_{a_i} is legal knowledge and $B \cap K_a \subseteq B \cap K_{a_i}$, and c is not predictable for a_i with respect to T^c. By Theorem 4.16, it is sufficient to show that $B \cap K_{a_i} \neq \emptyset$. If a is determined by T^c, then by Definition 4.2 and Theorem 3.11, it follows that $B \cap K_a \neq \emptyset$. If a is not determined by T^c, it follows by Theorem 4.16 that $B \cap K_a \neq \emptyset$ because c is predictable for a. Since these cases are exhaustive, it follows that $B \cap K_a \neq \emptyset$. Since $B \cap K_a \subseteq B \cap K_{a_i}$, it is the case that $B \cap K_{a_i} \neq \emptyset$.

The right-to-left direction follows immediately from Theorem 4.16. □

Proposition 5.13. Let $c \in \mathcal{C}$ be a concept.

(1) Static sequential predictions are order-independent for any non-empty set of conditions $B \subseteq \mathcal{B}$, any legal team on $B \cup \{c\}$, and any non-empty set of legal cases $A \subseteq \mathcal{A}$.

(2) Dynamic sequential predictions are not order-independent, for some non-empty set of conditions $B \subseteq \mathcal{B}$, a legal team on $B \cup \{c\}$, and some non-empty set of legal cases $A \subseteq \mathcal{A}$.

Proof. Statement (1) follows immediately from Definition 5.8. To prove (2), it is sufficient to provide a counterexample such that the predictions differ due the different order of the finite sequence of legal cases. Consider the following example.

Let $c \in \mathcal{C}$ be a concept, let $B = \{b_1, b_2, b_3, b_4\}$ and such $B_i^{\perp} = \emptyset$ for all $b_i \in B$. Further, let T^c on $B \cup \{c\}$ be defined as in Table A.4.

valuation v_i	legal possibility X_i	b_1	b_2	b_3	b_4	concept c
v_1	$\{b_1\}$	**1**	**0**	**0**	**0**	**1**
v_2	$\{b_3\}$	**0**	**0**	**1**	**0**	**0**
v_3	$\{b_4\}$	**0**	**0**	**0**	**1**	**0**
v_4	$\{b_2, b_3\}$	**0**	**1**	**1**	**0**	**0**

Table A.4: Legal team T^c

It is easy to verify that T^c is a legal team, $T^c \models \varphi_{\mathsf{suf}}^{X_1} \wedge \varphi_{\mathsf{nec}+}^{X_1} \wedge \varphi_{\mathsf{exc}}^{X_2} \wedge \varphi_{\mathsf{exc}}^{X_3} \wedge \varphi_{\mathsf{exc}}^{X_4}$ and that no legal possibility expressed by T^c satisfies any other property.

Consider the following two legal cases a_1 and a_2 with associated legal knowledge $K_{a_1} = \{b_2\}$ and $K_{a_2} = \{b_1, b_2\}$. Additionally, let $seq_1 : \{1,2\} \to \{a_1, a_2\}$ be defined by $seq_1(1) = a_1$ and $seq_1(2) = a_2$ and let $seq_2 : \{1,2\} \to \{a_1, a_2\}$ be defined by $seq_2(1) = a_2$ and $seq_2(2) = a_1$.

Let us consider seq_1. Since there does not exist $v \in T^c$ such that $\{b \in B \mid v(b) = \mathbf{1}\} = B \cap K_{a_1}$ and due to Fact 4.7, it follows that $T^c \cup \{v_{K_{a_1}}^{c+}\}$ and $T^c \cup \{v_{K_{a_1}}^{c-}\}$ are legal teams. Further, since $X_1 \not\subseteq B \cap K_{a_1}$ and by Proposition 4.13, it follows that $T^c \cup \{v_{K_{a_1}}^{c+}\} \not\models \varphi_{\mathsf{nec}+}^{X_1}$. Furthermore, due to the fact $X_1 \not\subseteq B \cap K_{a_1}$

and by Proposition 4.13, it follows that $T^c \cup \{v^{c-}_{K_{a_1}}\} \models \varphi^{X_1}_{\mathsf{suf}} \wedge \varphi^{X_1}_{\mathsf{nec+}} \wedge \varphi^{X_2}_{\mathsf{exc}} \wedge \varphi^{X_3}_{\mathsf{exc}} \wedge \varphi^{X_4}_{\mathsf{exc}}$. This means that $T^c \cup \{v^{c+}_{K_{a_1}}\}$ is not a constant extension and $T^c \cup \{v^{c-}_{K_{a_1}}\}$ is a constant extension. According to Definition 4.10, c is predicted to be not applicable to a_1 with respect T^c. By Definition 5.11, this means that c is sequentially dynamically predicted to be not applicable to a_1 for the sequence seq_1 with respect T^c. The sequential update results in $T^c_1 = T^c \cup \{v^{c-}_{K_{a_1}}\}$ which is specified by Table A.5.

valuation v_i	legal possibility X_i	b_1	b_2	b_3	b_4	concept c
v_1	$\{b_1\}$	**1**	**0**	**0**	**0**	**1**
v_2	$\{b_3\}$	**0**	**0**	**1**	**0**	**0**
v_3	$\{b_4\}$	**0**	**0**	**0**	**1**	**0**
v_4	$\{b_2, b_3\}$	**0**	**1**	**1**	**0**	**0**
$v_5 = v^{c-}_{K_{a_1}}$	$\{b_2\}$	**0**	**1**	**0**	**0**	**0**

Table A.5: Sequential update of T^c according to $seq_1(1) = a_1$

Observe that $T^c_1 \models \varphi^{X_1}_{\mathsf{suf}} \wedge \varphi^{X_1}_{\mathsf{nec+}} \wedge \varphi^{X_2}_{\mathsf{exc}} \wedge \varphi^{X_3}_{\mathsf{exc}} \wedge \varphi^{X_4}_{\mathsf{exc}}$ and it is easy to verify that $T^c_1 \models \varphi^{X_5}_{\mathsf{exc}}$. Further, no legal possibility expressed by T^c_1 satisfies any other property. According to Definition 5.11, the applicability of c to a_2 is then dynamically sequentially predicted based on T^c_1. Since $X_1 \subseteq B \cap K_{a_2}$ and by Proposition 4.13, $T^c_1 \cup \{v^{c-}_{K_{a_2}}\} \not\models \varphi^{X_1}_{\mathsf{suf}}$. Further, $T^c_1 \cup \{v^{c+}_{K_{a_2}}\} \not\models \varphi^{X_5}_{\mathsf{exc}}$ because of Proposition 4.13 and of the fact that $X_4 \subseteq B \cap K_{a_2}$. This means that $T^c_1 \cup \{v^{c+}_{K_{a_2}}\}$ and $T^c_1 \cup \{v^{c-}_{K_{a_2}}\}$ are not constant extensions. According to Definition 4.10, c is not predictable to a_2 with respect to T^c_1. By Definition 5.11, c is not dynamically sequentially predictable to a_2 for seq_1 with respect to T^c.

Next, let us consider the sequence seq_2. Since there does not exist $v \in T^c$ such that $\{b \in B \mid v(b) = \mathbf{1}\} = B \cap K_{a_2}$ and due to Fact 4.7, it follows that $T^c \cup \{v^{c+}_{K_{a_2}}\}$ and $T^c \cup \{v^{c-}_{K_{a_2}}\}$ are legal teams. Further, since $X_1 \subseteq B \cap K_{a_2}$ and by Proposition 4.13, it follows that $T^c \cup \{v^{c-}_{K_{a_2}}\} \not\models \varphi^{X_1}_{\mathsf{suf}}$. Further $T^c \cup \{v^{c+}_{K_{a_2}}\} \models \varphi^{X_1}_{\mathsf{suf}} \wedge \varphi^{X_1}_{\mathsf{nec+}} \wedge \varphi^{X_2}_{\mathsf{exc}} \wedge \varphi^{X_3}_{\mathsf{exc}} \wedge \varphi^{X_4}_{\mathsf{exc}}$ because of $X_1 \subseteq B \cap K_{a_2}$, $X_2, X_3, X_4 \not\subseteq B \cap K_{a_2}$ and Proposition 4.13. This means that $T^c \cup \{v^{c+}_{K_{a_2}}\}$ is a constant extension and $T^c \cup \{v^{c-}_{K_{a_2}}\}$ is not a constant extension. According to Definition 4.10, c is predicted to be applicable to a_2 with respect T^c. By Definition 5.11, this means that c is sequentially dynamically predicted to be applicable to a_2 for the sequence seq_2 with respect T^c.

The sequential update results in $T^c_1 = T^c \cup \{v^{c-}_{K_{a_2}}\}$ which is specified by Table A.6. According to Definition 5.11, the applicability of c to a_2 is then dynamically sequentially predicted based on T^c_1.

valuation v_i	legal possibility X_i	b_1	b_2	b_3	b_4	concept c
v_1	$\{b_1\}$	**1**	**0**	**0**	**0**	**1**
v_2	$\{b_3\}$	**0**	**0**	**1**	**0**	**0**
v_3	$\{b_4\}$	**0**	**0**	**0**	**1**	**0**
v_4	$\{b_2, b_3\}$	**0**	**1**	**1**	**0**	**0**
$v_5 = v^{c+}_{K_{a_2}}$	$\{b_1, b_2\}$	**1**	**1**	**0**	**0**	**1**

Table A.6: Sequential update of T^c according to $seq_2(1) = a_2$

Observe that $T^c_1 \models \varphi^{X_1}_{\mathsf{suf}} \wedge \varphi^{X_1}_{\mathsf{nec+}} \wedge \varphi^{X_2}_{\mathsf{exc}} \wedge \varphi^{X_3}_{\mathsf{exc}} \wedge \varphi^{X_4}_{\mathsf{exc}}$ and it is easy to verify that $T^c_1 \models \varphi^{X_5}_{\mathsf{suf}}$. Further, no legal possibility expressed by T^c_1 satisfies any other property.

Since there does not exist $v \in T^c_1$ such that $\{b \in B \mid v(b) = \mathbf{1}\} = B \cap K_{a_1}$ and due to Fact 4.7, it follows that $T^c_1 \cup \{v^{c+}_{K_{a_1}}\}$ and $T^c_1 \cup \{v^{c-}_{K_{a_1}}\}$ are legal teams. Further, since $X_1 \not\subseteq B \cap K_{a_1}$ and by Proposition 4.13, it follows that $T^c_1 \cup \{v^{c+}_{K_{a_2}}\} \not\models \varphi^{X_1}_{\mathsf{nec+}}$. Furthermore, due to the fact that $X_1, X_5 \not\subseteq B \cap K_{a_2}$ and by Proposition 4.13, it follows that $T^c_1 \cup \{v^{c-}_{K_{a_1}}\} \models \varphi^{X_1}_{\mathsf{suf}} \wedge \varphi^{X_1}_{\mathsf{nec+}} \wedge \varphi^{X_2}_{\mathsf{exc}} \wedge \varphi^{X_3}_{\mathsf{exc}} \wedge \varphi^{X_4}_{\mathsf{exc}} \wedge \varphi^{X_5}_{\mathsf{suf}}$. This means that $T^c_1 \cup \{v^{c-}_{K_{a_2}}\}$ is a constant extension and $T^c_1 \cup \{v^{c+}_{K_{a_1}}\}$ is not a constant extension. According to Definition 4.10, c is predicted to be not applicable to a_1 with respect T^c_1. By Definition 5.11, this means that c is sequentially dynamically predicted to be not applicable to a_1 for the sequence seq_2 with respect T^c.

Since c is dynamically sequentially predicted to be not applicable to a_2 for the sequence seq_2 of A with respect to T^c and c is *not* dynamically sequentially predicted to be not applicable to a_2 for the sequence seq_2 of A with respect to T^c, it follows due to Definition 5.12 that dynamic sequential predictions are not order-independent for some non-empty set of conditions $B \subseteq \mathcal{B}$, a legal team on $B \cup \{c\}$, and some non-empty set of legal cases $A \subseteq \mathcal{A}$. □

Fact 5.19. Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and T^c be a legal team on $B \cup \{c\}$.

(1) If X is totally applicability-relevant for c relative to T^c, then X is applicability-relevant for c and is not applicability-irrelevant for c relative to T^c.

(2) $\emptyset$ is applicability-relevant for c and totally applicability-relevant for c relative to T^c.

(3) c is not constant if and only if there exists some non-empty $X \subseteq B$ such that X is applicability-relevant for c relative to T^c.

Proof. Proposition (1) immediately follow from Definition 5.17. Proposition (2) follows immediately follows from Definition 5.17 and the fact that $\emptyset \subseteq X$ for any $X \subseteq B$.

For the left-to-right direction of (3), assume that c is not constant. Thus, $T^c \not\models =\!(c)$ and $\emptyset$ is not a dependency set for c. Further, $T^c \models =\!(b_1, \ldots, b_n, c)$ where $B = \{b_1, \ldots, b_n\}$ due to T^c being a legal team and Theorem 3.11. This means that B is a dependency set. Therefore, B is a minimal dependency set for c with respect to set inclusion, or there exists a non-empty $X \subseteq B$ which is a minimal dependency set for c with respect to set inclusion. It follows in both cases that there exists a non-empty set of conditions which is applicability-relevant for c relative to T^c.

The right-to-left direction of (3) is proven by contraposition. Suppose that c is constant. Thus, $T^c \models =\!(c)$. Therefore, $\emptyset$ is a dependency set. Note that for any non-empty set of conditions $X \subseteq B$, $X \not\subseteq \emptyset$ and $\emptyset \subseteq X$. Thus, $\emptyset$ is the only minimal dependency set for c. This means that there does not exist a non-empty $X \subseteq B$ such that X is applicability-relevant relative to T^c. □

Fact 5.20. Let $c \in \mathcal{C}$ be a concept, $B \subseteq \mathcal{B}$ be a non-empty set of conditions and T^c be a legal team on $B \cup \{c\}$. Suppose that c is not constant with respect to T^c, then the following facts hold:

(1) If a legal possibility $X \subseteq B$ expressed by some valuation of T^c is sufficient and positively necessary relative to T^c, then some non-empty set of conditions $Y \subseteq X$ is applicability-relevant for c relative to T^c.

(2) If a legal possibility $X \subseteq B$ expressed by some valuation of T^c is an exception and negatively necessary relative to T^c, then some non-empty set of conditions $Y \subseteq X$ is applicability-relevant for c relative to T^c.

Proof. Suppose that c is not constant. Hence, $T^c \not\models =\!(c)$ so that $\emptyset$ is not a dependency set for c.

To prove proposition (1), suppose that $X \subseteq B$ is a legal possibility expressed by some valuation of T^c. Assume that X is sufficient and positively necessary relative to T^c. Let us first show that X is a dependency set for c.

Since X is sufficient relative to T^c, $T^c \models \varphi^X_{\mathsf{suf}}$. This means that $T^c \models (\bigwedge_{b \in X} b) \to c$ and hence $T^c \models \neg(\bigwedge_{b \in X} b) \vee c$. Therefore, there exist $T_1, T_2 \subseteq T^c$ such that $T_1 \cup T_2 = T^c$, $T_1 \models \neg(\bigwedge_{b \in X} b)$ and $T_2 \models c$. Thus, for all $v \in T_1$, $\{v\} \not\models \bigwedge_{b \in X} b$. Note that for any $v_i \in T^c$ such that $v_i(b) = \mathbf{1}$ for all $b \in X$, $\{v_i\} \models \bigwedge_{b \in X} b$. Consequently, $v_i \notin T_1$ and hence $v_i \in T_2$. Therefore, (a) for any valuation $v_i \in T^c$ with $v_i(b) = \mathbf{1}$ for all $b \in X$, $v_i(c) = \mathbf{1}$.

Additionally, since X is positively necessary relative to T^c, $T^c \models \varphi^X_{\mathsf{nec+}}$. This means that $T^c \models c \to (\bigwedge_{b \in X} b)$ and hence $T^c \models \neg c \vee \bigwedge_{b \in X} b$. Therefore, there exist $T_3, T_4 \subseteq T^c$ such that $T_3 \cup T_4 = T^c$, $T_3 \models \neg c$ and $T_4 \models \bigwedge_{b \in X} b$. Thus, for all $v \in T_3$, $\{v\} \not\models c$. Note that for any $v_j \in T^c$ with $v_j(c) = \mathbf{1}$, $v_j \notin T_3$ and hence $v_j \in T_4$. Therefore, (b) for any valuation $v_j \in T^c$ with $v_j(c) = \mathbf{1}$, $v_j(b) = \mathbf{1}$ for all $b \in X$.

By (a) and (b) it follows that for any valuation $v \in T^c$, $v(c) = \mathbf{1}$ if and only if $v(b) = \mathbf{1}$ for all $b \in X$. Consequently, for any valuations $v_i, v_j \in T^c$ with $v_i(b) = v_j(b)$ for all $b \in X$, $v_i(c) = v_j(c)$. This implies that $T^c \models =\!(b_1, \ldots, b_n, c)$ where $X = \{b_1, \ldots, b_n\}$. Thus, X is a dependency set. Since $\emptyset$ is not a dependency set, either X is a minimal dependency set or there exists a non-empty minimal dependency set $Y \subseteq X$. It follows in both cases that there exists a non-empty set of conditions $Y \subseteq X$ which is applicability-relevant for c.

By analogous reasoning, it follows that proposition (2) is the case. $\square$