The Source Conflict Cost Criterion: A Pragmatic Approach to Floating Conclusions

By

Jorge Andrés Morales Delgado

A thesis submitted to the Victoria University of Wellington in fulfilment of the requirements for the degree of Doctor of Philosophy

Victoria University of Wellington

2020

Acknowledgements

This research was possible to the effort provided by several institutions and numerous persons, to which I want to express my gratitude.

I would like to acknowledge the comprehensive research grant facilitated by the *Universidad de Costa Rica* through the *Oficina de Asuntos Internacionales* in coordination with *Sede del Caribe*. Their continuing support has had a positive impact on my ability to focus on my research. In the same vein, I would like to thank *Victoria University of Wellington* for their research funding through the *Victoria Doctoral Scholarship*.

I would also like to thank Prof Dr Edwin Mares, who very kindly worked with me throughout my time as a PhD student, providing prompt and insightful feedback on an ongoing basis. Prof Dr Mares is not only an outstanding scholar but also a kind person who always sees the best in his students, and with optimism, care and compassion urge us to keep moving forward. You are truly a role model in academia where excellence and kindness are often seen as conflicting traits. It has been an honour to conduct my PhD studies under your supervision.

Although not involved directly in the research, I want to acknowledge the mentorship I received throughout my undergraduate studies from the following outstanding philosophers from the Universidad de Costa Rica: Dr Max Freund, Dr Mauricio Molina, Dr Mario Solis, Dr Juan Diego Moya, Dr Jethro Masis and Dr Jorge Jimenez. I will always treasure your support and encouragement.

Finally, but more importantly, I want to thank my wife, Andrea Mejia, for her understanding and support throughout this process. This PhD has been a long journey we have taken since the day we met in an ethics course as undergraduates ten years ago. I could not have done it without you, and it would not have been as meaningful as it has been. This is for you with love.

Abstract

Our research examines the problem of multiple lines of reasoning reaching the same conclusion, but only through different and unrelated arguments. In the context of non-monotonic logic, these types of conclusions are referred to as *floating conclusions*. The field of defeasible reasoning is divided between those who claim that floating conclusions ought not to be accepted through a prudent or skeptical point of view, whereas others argue that they are good enough conclusions to be admitted even from a conservative or skeptical standard. We approach the problem of floating conclusions through the formal framework of *Inheritance Networks*. These networks provide the simplest and most straightforward gateway into the technical aspects surrounding floating conclusions in the context of non-monotonic logic and defeasible reasoning.

To address the problem of floating conclusions, we construct a unifying framework of analysis, namely, the *Source Conflict Cost Criterion* (SCCC), that contains two basic elements: source conflict and cost. Both elements are simplified through a binary model, through which we provide a comprehensive understanding of the floating conclusions as well as the problematic nature of the debate surrounding this type of inferences. The SCCC addresses three key objectives: (a) the assessment of floating conclusions and the debate surrounding its epistemological dimension, (b) the construction of a general and unified framework of analysis for floating conclusions, and (c) the specification of the normative conditions for the admission of floating conclusions as skeptically acceptable information.

Table of Contents

1. Introduction	1
1.1 Defeasible Reasoning	1
1.2 Inheritance Networks	3
1.3 Problems and Objectives	6
1.4 Structure	11
2. Inheritance Networks	15
2.1 Defeasible Reasoning	15
2.2 Basic Elements	27
2.3 Paths	37
2.4 Extracting Information	43
2.5 Chapter Summary	55
3. Features and Delimitation	57
3.1 Weaknesses	57
3.2 Strengths	64
3.3 Delimitation	71
3.4 Chapter Summary	80
4. Ambiguity	81
4.1 The Problem	81
4.2 Credulous Approach	88
4.3 Skeptical Approaches	91
4.4 Chapter Summary	103
5. Floating Conclusions	104
5.1 Producing Floating Conclusions	104
5.2 Approaches	115
5.3 Discussion	123
5.4 Chapter Summary	129

6. The Debate over Floating Conclusions	130
6.1 Defending Indirect Skepticism	130
6.2 Defending Direct Skepticism	137
6.3 Overview of the Debate	143
6.4 Chapter Summary	150
7. The Source Conflict Cost Criterion	151
7.1 The Notion of Source	151
7.2 The Notion of Conflict	157
7.3 The Notion of Cost	162
7.4 The Criterion	169
7.5 Chapter Summary	177
8. Analysing Floating Conclusions through the SCCC	179
8.1 Conflicting Sources with High Cost	179
8.2 Conflicting Sources with Low Cost	186
8.3 Non-Conflicting Sources with High Cost	195
8.4 Non-Conflicting Sources with Low Cost	200
8.5 SCCC and Inheritance Networks	206
8.6 Chapter Summary	210
9. Synthesis	211
9.1 SCCC and Cognition	211
9.2 Ethics of Belief	217
9.3 Empirical Basis	225
9.4 Chapter Summary	231
10. Conclusions	232

List of Figures

1. Introduction	
Figure 1.2.1 Basic Inheritance Network	4
Figure 1.3.1 Ambiguous Inheritance Network	6
2. Inheritance Networks	
Figure 2.2.4-a Directed and Undirected Graphs	35
Figure 2.2.4-b Cyclic and Acyclic Graphs	35
Figure 2.3.1 Paths with Positive Links	38
Figure 2.3.2-a Path with a Negative Link	40
Figure 2.3.2-b Interrupted Negative Path	41
Figure 2.3.2-c Route Starting with a Negative Link	42
Figure 2.4.1 Basic Inheritance Network	46
Figure 2.4.2-a General Single Inheritance Network	47
Figure 2.4.2-b Basic Single Inheritance Network	48
Figure 2.4.2-c General Inheritance Network with a Negative Link	49
Figure 2.4.2-d Basic Inheritance Network with a Negative Link	51
Figure 2.4.3-a General Multiple Inheritance Network	53
Figure 2.4.3-b Basic Multiple Inheritance Network	54
3. Features and Delimitation	

Figure 3.3.2-a Basic Redundancy	74
Figure 3.3.2-b Basic Redundancy with Negative Links	75
Figure 3.3.3 Multiple Inheritance Network	78

4. Ambiguity

Figure 4.1.1 Nixon Diamond	82
Figure 4.1.3-a Nixon Diamond Extension 1	86
Figure 4.1.3-b Nixon Diamond Extension 2	87

5. Floating Conclusions

Figure 5.1.2-a Seedless-Grape-Vine network (sgv-net)	106
Figure 5.1.2-b SGV-Net First Path	107
Figure 5.1.2-c SGV-Net Second Path	108
Figure 5.1.2-d SGV-Net Third Path	108
Figure 5.1.2-e SGV-Net Fourth Path	109
Figure 5.1.3-a SGV-Net First Conflict	111
Figure 5.1.3-b SGV-Net Second Conflict	112
Figure 5.1.3-c SGV-Net First Extension	113
Figure 5.1.3-d SGV-Net Second Extension	114

6. The Debate over Floating Conclusions

Figure 6.3.1 The Floating Conclusions Arguments for Indirect Skepticism	145
Figure 6.3.2 The Floating Conclusions Arguments for Direct Skepticism	148
Figure 6.3.3 The Floating Conclusions Debate	149

7. The Source Conflict Cost Criterion

Figure 7.4.1-a Binary Source Conflict Criterion Model	171
Figure 7.4.1-b Binary Cost Criterion Model	171
Figure 7.4.2-a High Source Conflict Criterion	172
Figure 7.4.2-b Low Source Conflict Criterion	173
Figure 7.4.3-a High-Cost Criterion	174
Figure 7.4.3-b Low-Cost Criterion	175
Figure 7.4.4 Complete Source Conflict Cost Criterion Binary Model	176

8. Analysing Floating Conclusions through the SCCC

Figure 8.1.1 Floating Conclusions Debate and the +SC +C condition	181
Figure 8.1.3 Floating Conclusions Stance and the +SC +C condition	186
Figure 8.2.1 Floating Conclusions Debate and the +SC -C condition	188
Figure 8.2.3 Floating Conclusions Stance and the +SC -C condition	194
Figure 8.3.3 Floating Conclusions Stance and the -SC +C condition	200
Figure 8.4.3 Floating Conclusions Stance and the -SC -C condition	206

Chapter 1 Introduction

This work is the result of pondering a very simple yet profound problem for defeasible reasoning. Our thesis examines the problem known as floating conclusions. This problem occurs when lines of reasoning, converge on one and the same conclusion, without having a commonly shared argument. We examine floating conclusions within the formalism known as inheritance networks, which, despite its simplicity, helps us focus on the problem itself rather than on the technicalities through which the problem is brought about. Our research proposes a comprehensive framework to understand the problem of floating conclusions for defeasible reasoning. As such, our work includes topics and themes of multiple threads of research as opposed to a monochromatic line of thought. Thus, this is not a thesis on logic, epistemology or cognitive science, but rather an attempt to use the tools from all these fields to try to understand the problem under scrutiny. In the following sections, we go over the main topics and highlight the narrative to be found in the forthcoming chapters.

1.1 Defeasible Reasoning

As we mentioned before, the problem of floating conclusions emerges within the general field of defeasible reasoning. Thus, before diving into the problem that lies at the heart of our research, we need to take a step back to look at the big picture.

Defeasible reasoning is, in a very general sense, the field that examines certain modes of human common-sense reasoning in which our inferences and reasoning processes are not taken as final, but stand to be altered as new information comes to light. This type of reasoning process is contrasted with classical deductive reasoning, in which once a given inference is made, it does not stand to be retracted. Thus, in the classical approach to deduction, the acquisition of information is a linear and additive process. This trait is also usually referred to as monotonicity. For example, the basic Euclidean geometry has a set of axioms and the information acquired from those axioms builds in an increasingly additive fashion. Given the specific properties of such a domain, no further inference can undermine a previous one. This picture of classical mathematical reasoning is quite attractive and stable; nevertheless, it is far from what actual human common-sense reasoning seems to behave like (Stein, 1991).

Since our reasoning processes can fail due to a plethora of reasons (for example the various ways in which our cognitive architecture or perceptual modalities may trick us), this alone does not suffice to explain what defeasible reasoning is about. Abductive reasoning and probabilistic reasoning are both instances of fallible non-monotonic reasoning, which generally speaking, are not considered within the standard account of defeasible reasoning. Thus, defeasible reasoning restricts itself to those forms of reasoning in which the connection between sets of premises and a given conclusion can cease to hold, by the addition of further information or by the fallibility of the claims involved in the reasoning. For example, the claim that Latin Americans speak Spanish is true for the most part, but can nonetheless stand to be corrected. If the Latin American under consideration is Brazilian, then she might speak Portuguese and not Spanish, but that does not change the legitimate typicality claim of Latin Americans as Spanish speakers, despite some cases' not complying with it.

Given that defeasible reasoning involves epistemic notions (such as fallibility and retractability) as well as logical ones (such as inferences, reasoning and monotonicity), the research programme itself has various disciplines at its core. On the one hand, philosophy has provided the epistemological insights into the extent and limitations of the main concepts involved. In this sense, philosophy serves as the theoretical foundation for the main ideas underlying defeasible reasoning. On the other hand, logic has provided a set of precise formal tools through which the philosophical intuitions can be examined. At the same time, these formal tools provide a way to gain further insights into the way defeasible reasoning operates. Finally, as a hybrid approach, Artificial Intelligence has also provided input into the field. In particular, it not only takes the insights from logical research and empirically tests them, but the experimental results of Artificial Intelligence provide further insights into both, the logical frameworks and the philosophical concepts. All of the above is a brief characterisation the field, that nonetheless provides a first step understanding of defeasible reasoning as a research programme.

In line with the above, defeasible reasoning can be approached as an area of epistemology or as a chapter of logic. In this vein, there are various ways to undertake any given debate within the field. Our research is an instance of this. We start with a formal framework and examine the technical problem associated with this framework (*ambiguity*). We further discuss how this initial problem produces a more complex problem (*floating conclusions*). Nevertheless, after the technical

assessment, our research moves into the conceptual and epistemological dimension.

With this general picture regarding the theoretical background that underlies our research, in the next section, we briefly discuss the formal framework through which we conduct our work. This formal framework is simple yet advantageous for our specific purposes.

1.2 Inheritance Networks

The concept of an inheritance or semantic network is usually attributed to Ross Quillian in his 1967 publication in *Behavioral Science* entitled "Word Concepts: A Theory and Simulation of some Basic Semantic Capabilities". Such work had as its core purpose to "[...] develop a theory of the structure of human long-term memory, and to embody this theory in a computer model such that the machine can utilize it to perform complex memory-dependent tasks." (Quillian, 1967, 410). The fundamental intuition underlying Quillian's framework and which all of the upcoming generation of network-like frameworks would follow (despite the ample array of variation) revolved around of a model that "[...] consists, basically of a mass of nodes, interconnected by different kinds of associative links" (Quillian, 1967, 411). Furthermore, Quillian's original enterprise was to create a computer model of human memory that could depict how human agents process and compute semantic information. As Brachman (1977) points out, the work of Quillian aimed to construct "a general associative mechanism for encoding the 'meaning' of words" (127).

The extent and limitations of Quillian's initial motivation and the success of such research programme need not to directly concern us in the context of our current work. Rather, we focus on how such theoretical intuition was later undertaken by logicians to ground a similar representational framework. More precisely, the underlying idea of semantic networks was taken to the field of defeasible reasoning, in which the connection of nodes through arcs stand for the claim that a given node typically or defeasibly has a particular property.



Figure 1.2.1 Basic Inheritance Network

For example, Figure 1.2.1 depicts a basic inheritance network, in which we have the individual Carlos and the attributes: South American, Latin American, Spanish Speaker, Communist and Religious. This network is arranged such that Carlos is linked to the attribute of being South American that is linked to being Latin American which is then linked to the attribute of being a Spanish Speaker. Likewise, the network links the same individual Carlos to the Communist attribute which is negatively linked to the property of being Religious. Thus, we can take each route as a line of reasoning or argument that allows us to infer implicit information from a given node based on how it is linked to other nodes within network.

In the context of our specific inheritance network we can infer that Carlos is a Spanish Speaker as well as him not being Religious. Both attributes are made from the two different routes contained within the network. This very simple arrangement of nodes and links illustrates the way in which we can represent defeasible knowledge in a simple but efficient fashion.

Furthermore, the links between attributes are only defeasible in nature. Should we acquire further information that Carlos is from Brazil (e.g. in the form of an additional node inserted after the

5

Latin American attribute), then the conclusion that he speaks Spanish might not be warranted. In the same vein, Carlos might be a communist militant associated with the Liberation Theology, and thus the conclusion that he is not religious might not hold. Despite such atypical scenarios, the prototypical inference of Carlos being a Spanish Speaker and not being Religious is not wrong but rather a defeasible inference, in the context of the defeasible information provided by these kinds of networks.

The use of Quillian's framework in the context of defeasible reasoning is far from the original conception of such type of networks. Nevertheless, after extensive work and research on such endeavour, the route has been adopted as a legitimate line of research within the field, and it is this specific thread that we inquire in our research. In such a vein, we adopt a terminological distinction that highlights the two different endeavours. On the one hand, and when we refer to Quillian's original idea of a network we use of the term *semantic network*. On the other hand, when we refer to networks as applied to the field of defeasible reasoning, we use the term *inheritance networks* or simply *networks*. Despite such differentiation, the true extent and limitation of how far apart or rather how close both research endeavours are from each other is something that we note, yet do not pursue.

It is equally important to emphasise that the origins of the network-like models for defeasible reasoning shape the representational framework even if such network-like framework is taken out of its initial context. In this sense, we mention two important methodological aspects throughout our research. First, not all the concepts and problems concerning semantic networks apply directly to inheritance networks. That is, in the process of using semantic networks for defeasible reasoning some debates could cease to have a place. Second, and serving as the converse of the previous aspect, concepts and ideas that arise in the context of defeasible reasoning might not have a place in the initial conception of Quillian's network. For the purposes of our work, we follow the idea of networks as a representational framework consisting of nodes and arcs connecting such nodes, but only to the extent that it concerns the field of defeasible reasoning.

The above serves as a brief sketch of the formal framework that grounds our assessment of defeasible reasoning. We are now better equipped to discuss the fundamental problem that our dissertation explores, i.e. the problem of *floating conclusions*.

1.3 Problem and Objectives

Floating conclusions is the problem that lies at the centre of our research. This is, to a certain extent, a simple problem. Nevertheless, to have a sound understanding of this problem, we need to progressively consider the intermediate steps that help produce this puzzle.

The first step towards producing floating conclusions is the problem of ambiguity within defeasible reasoning. Ambiguity within inheritance networks is a situation in which two or more streams of information arrive at opposite conclusions. To provide a preliminary idea of the nature of this problem we can consider the following network:



Figure 1.3.1 Ambiguous Inheritance Network

Figure 1.3.1 depicts a network where we have two lines of reasoning regarding the individual Manuel. On the one hand, we have a route that indicates that Manuel has the attribute of being Latin American, which in turn is defeasibly associated with the attribute of being religious. On the other hand, we have a route that indicates that Manuel has the attribute of being a Communist, which is then negatively associated with the attribute of being religious. The attribute of being religious is said to be ambiguous with respect to Manuel because the network contains two

unrelated routes pointing to opposite information. Thus, it is not clear what should be said of Manuel with respect to the religious attribute, as we have two lines of reasoning supporting each conclusion, i.e. him being religious and him not being religious.

The salient feature of ambiguous networks is that there is no mechanism through which the dispute can be arbitrated. Thus, one is left with a situation to believe one thing and its opposite without having the resources to distinguish the best alternative.

To address the problem of ambiguity, there are at least two different approaches. First, we have a liberal or *credulous* approach, which seeks to extract as much information as it can, even at the expense of conflict. The credulous approach is a mechanism that operates by accepting conclusions as long as such conclusions are contained in any of the involved lines of reasoning of a given network. Second, we have a more conservative or *skeptical* approach that seeks to restrain itself as much as possible from endorsing such conflicting information. The skeptical approach achieves such conservative attitude towards extracting information from inheritance networks by only accepting conclusions that are contained in the *intersection* of the different lines of reasoning associated to a given inheritance network (as opposed to the credulous approach that accepts conclusions that are contained in any line of reasoning). The credulous strategy is portrayed as a more prudent strategy.

Both mechanisms set themselves to tackle the problem of ambiguity, and they do so from different and opposing stances to endorsing information from inheritance networks. Nevertheless, in the process of addressing the problem of ambiguity, these mechanisms open a set of new problems and questions regarding the topic of information processing within such representational structures. More precisely, there is not one, but two different ways to instantiate the operation of intersecting the different lines of reasoning within a given inheritance network: indirect and direct skepticism.

On the one hand, *indirect skepticism* tells us that the conclusions that we ought to endorse from a given network are the attributes endorsed by every single route or line of reasoning within a given network. Such privileged conclusions are referred to as *skeptically acceptable attributes*. On the other hand, *direct skepticism* tells us that the conclusions we ought to endorse from a given network are those conclusions which are contained in an argument that is common to every single line of

reasoning. Such privileged argument that is common to every single line of reasoning is referred to as a *skeptically acceptable argument*. The direct version of skeptical inference, instead of endorsing *skeptically acceptable attributes*, it only supports the conclusions that follow from a *skeptically acceptable argument*.

In the case of Manuel's religiosity (Figure 1.3.1), both versions of skeptical inference would come to the same conclusion that nothing can be inferred regarding him being religious or not. Nevertheless, they do so for different reasons. On the one hand, indirect skepticism states that since the attribute of being religious or not being religious is not contained in every single line of reasoning (one line of reasoning contains one and the other contains the opposite) then nothing can be inferred regarding Manuel's religiosity. That is, there are no *skeptically acceptable attributes*. On the other hand, direct skepticism states that there is no common argument shared by all lines of reasoning (i.e. a skeptically acceptable argument). That is, there is no *skeptically acceptable acceptable argument*. Thus, in the absence of such an argument nothing can be inferred regarding Manuel's religiosity. In the case of inheritance networks such as the one in Figure 1.3.1 both versions of skeptical inference support the same conclusions.

As we previously mentioned, the problem of floating conclusions relies on the fact that there is not one but rather two different, and *prima facie*, equally legitimate skeptical approaches. In particular, a natural concern revolves around the extent to which both approaches to skeptical inference are just different ways to carry out the same idea or whether they are ultimately irreconcilable. The existence of two skeptical reasoning processes to ambiguous networks opens a myriad of challenges to information-processing approaches within inheritance networks. The problem emerges because both mechanisms that attempt to instantiate the same theoretical intuition are not only different in nature, but they also provide, under specific circumstances, different and opposing results. One of such cases is what is commonly referred to as floating conclusions.

Floating conclusions refer to information that is endorsed by every thread in a given network, but it does not have a unique argument that is shared by all threads of the network. These conclusions are said to "float" among all lines of reasoning of the network without having a common argument shared by every single line of reasoning. To better understand the inner workings of this abstract problem, we consider a simple but illustrative example due to Matthew Ginsberg in 1987¹.

¹We attribute this example to Ginsberg in 1987 although the case does not appears in print until a 1993

Let us consider the case of two different streams of information concerning a given individual Nixon. One route state that Nixon is a Republican and him being a Republican makes him a hawk with respect to international relations, which in turn renders him an extremist in political terms. Another route tells us that Nixon is a Quaker, and as such he is a dove in the context of international relations, making him an extremist in political terms. Both lines of reasoning reach the same conclusion from different arguments. It is in this sense, that the conclusion of Nixon being an extremist is said to "float" among the various lines of reasoning without having on unique argument shared by all the involved streams of information. The question is whether the conclusion that Nixon is an extremist should be inferred.

On the one hand, one could argue that despite the different lines of reasoning, one certainly ends up arriving at the conclusion that Nixon is an extremist. The fact that one argument relies on him being a Republican whereas the other relies on him being a Quaker should not distract us from the end result that is one and the same. This is the reasoning that motivates indirect skepticism, since it endorses those conclusions supported by every single line of reasoning, which in this case is the conclusion of Nixon being an extremist. That is, the attribute extremist is a *skeptically acceptable attribute*. From this perspective, the conclusion seems warranted.

On the other hand, one could argue that despite all lines of reasoning reaching the same conclusion, we should restrain from endorsing that particular conclusion because such lines of reasoning reach the same conclusion but only through different arguments. After all, Nixon being a dove is at odds with Nixon being a hawk, and despite both positions being extreme one could argue that the combination of traits for this individual might mitigate his political extremism. That is, both venues can be right in that he is a Republican and a Quaker, but such combination might lighten his political stance in the realm of international relations. This is the motivation behind the direct version of skeptical inference, since it endorses those conclusions supported by a route or argument shared by every single line of reasoning. In this case, there is no such argument. That is, there is no *skeptically acceptable argument*. In this sense, the conclusion that Nixon is an extremist with respect to his political stance does not seem like an appropriate conclusion to draw.

This problem reflects the situation in which a given conclusion is defeasibly justified by every line

publication with a slightly different arrangement. Nevertheless, we follow Horty and others who attribute this example to Ginsberg "during the question session after the AAAI-87 presentation" (Horty, 2002, 61).

of reasoning, among a given array of arguments in support of such conclusion, but none of the involved lines of reasoning coincides in a commonly shared argument to endorse such conclusion. That is, the conclusion is justified but only through different and non-shared arguments. A prevalent view within the field is that floating conclusions are, in general, less than reliable pieces of information and as such, we should restrain from admitting them as acceptable from a skeptical standpoint.

On a first impression, floating conclusions seem like a highly localised problem associated with inheritance networks with little or no bearing at all outside the technicalities of the field. Nevertheless, the problem itself underpins some fundamental puzzles associated with justification and decision-making. The problem is important if we take into account that it can emerge in various domains such as legal and medical reasoning. One of the recurring examples we work through in the context of our research is the hypothetical case in which different witnesses all coincide in the culpability of a person with respect to a crime, but all of them disagree among themselves on the extent to which such person committed the crime. Should we dispense the life or freedom of the involved defendant or should we let all the weight of the law fall upon him disregarding that the witness's testimonies are at odds among themselves? Such crucial questions that strike us as fascinating from a purely epistemological point of view, can have real-world applications.

In light of the above, one would think that this problem has received a systematic examination, if not yet resolved to some degree of satisfaction. Nevertheless, the literature suggests a more discouraging state of affairs. In particular, there are two fundamental reasons as to why the phenomenon of floating conclusions and its subsequent treatment in the literature is less than satisfying, which we call the 'epistemological' and the 'technical' problems.

First, the debate surrounding floating conclusions, and both approaches to skeptical inference is dispersed throughout the literature without any attempt to systematise such discussion. The epistemic aspects are far from well understood. Thus, not only the various insights into floating conclusions are scattered, but also the debate itself suffers a similar fate. Such state of affairs undermines a proper philosophical understanding of the nature and conditions of acceptance of this type of inferences. Furthermore, there is no general and unified approach to the problem itself. That is, different researchers have tackled floating conclusions (to the small extent they have done so) from various approaches and have put forward their stances without satisfactorily addressing

the critiques from the other end of the debate. All of the above renders the discussion, and any possible attempt to solve the debate, as a non-converging dialogue to understanding the problem of floating conclusions.

Second, and as a corollary of the above, the lack of a proper understanding of the epistemic aspects of floating conclusions undermines the possibility of determining which type of skeptical inference ought to be set in place under which type of circumstances. That is, the lack of a proper understanding of the epistemic nature of the problem preempts the very possibility to make an informed decision of the advantages and disadvantages provided by each version of skeptical inference and place their technical features in the context of the larger epistemological framework. The aforesaid shortcoming renders itself as the conditions upon which there is a gap between our epistemological and technical understanding of the problem at hand.

The above is the discouraging state of affairs that surrounds this simple, yet unsolved problem within the field of defeasible reasoning. This is the challenge that motivates this research, and the background upon which our work is constructed.

This research has set itself four key objectives, through which our thesis unfolds. The first objective is to examine the problem of floating conclusions. We provide a thorough review of the problem within the context of inheritance networks, and discuss the epistemological aspects surrounding the technical problem. Here the objective is to systematise the various views of floating conclusions. The second objective is to construct a framework that can account for floating conclusions as well as the epistemic aspects associated with such inferences. The third objective is to apply our framework of analysis to the problem of floating conclusions, with a particular interest in the epistemological foundations of skeptical inferences. The fourth objective is to provide the general conditions to determine the acceptance/rejection conditions for floating conclusions and each type of skeptical inference accordingly. That is, to present a general normative stance, of the circumstances under which floating conclusions should be regarded as skeptically acceptable information and when such inferences are best dismissed.

1.4 Structure

To tackle the problem and address the objectives set for our research, we are going to organise our

work in a four-component structure, which serves as a logical decomposition of the thesis. Each of the four logical components contains one or more chapters to address the task at hand for each stage.

Part I aims to set the theoretical background. To address this component, we devote three chapters. First, chapter 2 (Inheritance Networks) discusses the basic components of the formal framework of inheritance networks. Here we also briefly review defeasible reasoning as a research programme and the specific approach we have for the purposes of our research. Moreover, we inquire into defeasible reasoning as the theoretical background that underpins our assessment of inheritance networks. Second, chapter 3 (Features and Delimitation) undertakes a more critical examination of inheritance networks. This chapter looks into the strengths and limitations of this formal framework for defeasible reasoning. Moreover, it provides an argument not only as to how the strengths of this formal framework make it an optimal choice for our assessment of floating conclusions but also how its shortcomings actually play a positive role for the specific aims of our research. Third, chapter 4 (Ambiguity) builds on top of the previous chapters and looks into one of the most basic but troublesome problems emerging from inheritance networks, namely the problem of ambiguity. We examine the conditions upon which the problem of ambiguity emerges, and more importantly, the different approaches to solve it. We devote special attention to the two-fold approach of the skeptical solution to ambiguity: indirect and direct skepticism. This chapter sets the fundamental aspects that anchor the rest of the investigation.

Part II examines the core problem that emerges from the theoretical background set in Part I. This part contains two chapters. First, chapter 5 (*Floating Conclusions*) examines the details surrounding the problem of floating conclusions. It considers the technical conditions that produce this problem, which is a result of two different ways to instantiate the skeptical approach to ambiguous networks. Moreover, we assess how both approaches handle networks involving floating conclusions. Lastly, we discuss the technical aspects surrounding floating conclusions, and the extent to which they play a role within inheritance networks. Second, chapter 6 (*The Debate over Floating Conclusions*) follows the thread from the previous chapter, but sets aside most of the technical aspects concerning floating conclusions and discusses in detail the epistemological dimension of the problem. Here, we examine the arguments presented in favour and against each stance of the debate and highlight the underlying intuitions in defence of each stance. This chapter systematises the different stances surrounding this debate and summarises the most important arguments for each approach.

Departing from the above, Part III sets itself as the central component of our approach to address the problem of floating conclusions. Chapter 7 (*The Source Conflict Cost Criterion*) discusses the notions of a source, conflict between sources and the costs associated with a given conclusion. Through these concepts, we construct the *source conflict and cost criterion* (SCCC), which is the central tool underlying our approach to floating conclusions. We also provide a binary simplification of the SCCC for the purposes of our research. At this stage, we limit ourselves to the construction of our framework, as the application is left to a separate chapter.

Finally, in Part IV we apply the framework of analysis constructed in Part III to tackle the problem presented in Part II. This last stage of our research has two chapters. Chapter 8 (*Analysing Floating Conclusions through the SCCC*) approaches the problem of floating conclusions (presented in chapter 5) and use the framework of analysis to systematise and assess the debate surrounding floating conclusions (discussed in chapter 6). Based on the aforementioned, we present the general conditions of acceptance and rejection of floating conclusions as skeptically acceptable information according to the various specifications in our framework of analysis. Finally, Chapter 9 (*Synthesis*) offers a brief but necessary reflection and analysis on the meaning of our framework of analysis and how it ties with some of the more salient debates and topics touched upon the first of part of thesis as well as alternative lines of inquiry, such as looking into the empricial basis of our approach to floating conclusions.

All of the above has its own particular variances in style and pace, which we consider important to emphasise. Part I relies heavily on the literature and research done so far on inheritance networks and defeasible reasoning. As such, we include direct quotes when we deem relevant. The caveat of our frequent usage of direct quotations lies in that our work emerges in the context of this research programme, yet the proposed contribution departs qualitatively from what has been done so far. Thus, we substantially rely on the aspects we wish to bring attention through quotations which otherwise would not be relevant had we follow a different methodological approach to the problem at hand. Nevertheless, since our aim is not to overload the text, we have placed the quotations that we deem important but not as vital as footnotes. The above contrasts directly with Part II, in which our use of a direct quotation is sparser and almost absent in Parts III and IV, where we concentrate in developing with full detail our framework of analysis and its respective application. All of the above we hope is for the optimal depiction not only of our work and proposed contribution but also in the pursuit for a genuine understanding of the theoretical framework in which the problem we inquire emerges.

In such a vein, we put forward this research first, as an attempt to tackle the problem of floating conclusions and the two types of skeptical inferences and second, as the groundwork to settle the methodological and theoretical conditions upon which further progress can be made. Throughout our research, we aim to provide a qualitative break from what has been done so far in the context of the aforesaid problem. We are confident that the answers constructed in the context of this research are a viable approach to tackle floating conclusions. More importantly, we hope that our work can set the conditions upon which a more fruitful debate can be undertaken that supersede the palpable shortcomings previous attempts have had, in a way that tangible progress can be made.

Chapter 2 Inheritance Networks

Inheritance networks are conceived as a particular formal framework for knowledge representation and defeasible reasoning. The network-like structure emerged from research into *Semantics Networks* (Quillian, 1967), consisting of nodes and links among them, to model defeasible reasoning (Horty, 1994). In this chapter, we discuss and analyse the main components of inheritance networks, and the theoretical intuitions behind such framework. This chapter aims to provide the general but thorough overview of the formal framework upon which our whole research is constructed. In particular, we examine the theoretical foundations and the adjacent fields of study surrounding our work.

2.1 Defeasible Reasoning

Defeasible reasoning as a research programme has as one of its central objectives to better understand the nature of various forms of human common-sense reasoning and, more importantly, provide the tools and mechanisms to understand such phenomena. In this section, we look into defeasible reasoning to better understand one of the pillars underlying our research. This section aims to provide a picture (brief and succinct as possible) of the complex field of defeasible reasoning. In particular, we aim to highlight the main ideas underlying defeasible reasoning such that the discussion, in further sections, of inheritance networks can be better understood.

2.1.1 Fallibility

In contemporary epistemology, we acknowledge fallibility as a fundamental trait of human reasoning. The fallible feature of our everyday common-sense reasoning makes our conclusion at any given point tentative and open to correction or retraction, i.e. defeasible. As Pollock states:

[...] defeasible reasoning is not just common; it is thoroughly pervasive and absolutely essential. Almost everything we believe is believed at least indirectly on the basis of defeasible reasoning, and things could not have been any other way. (1995, 41)

Furthermore, this tentative or defeasible trait of common sense reasoning can be seen as a suboptimal or skewed version of better forms of drawing inferences. Nevertheless, defeasible reasoning is itself a standard common practice inherent to human agents in the context of their successful operation within a given environment. As Pollock states:

What is overlooked is that defeasible reasoning is normal reasoning. Its use is not just a matter of computational efficiency. It is logically impossible to reason successfully about the world around us using only deductive reasoning. All interesting reasoning outside mathematics involves defeasible steps. (1995, 41)

In contrast with the depiction of human reasoning as a malleable inference process, most formal models of inference rely on quite a contrary assumption, namely, monotonicity. Monotonicity expresses the idea contained in classical logic that if some conclusion is inferred from a given set of premises, the addition of more premises to the initial set does not interfere with previous inferences (Costa, 1999)². To put it in another way, a belief cannot be withdrawn if it was derived at some earlier point in the inference process, regardless of new evidence or sources of information appended to the knowledge base of the agent or system. The aforesaid is known as the monotonicity principle. For example, let us consider the following argument:

- 1. Birds Fly
- 2. Tweedy is a bird
- 3. Tweedy Flies

According to the aforesaid premises, the conclusion that Tweedy flies seems to follow according to classical logic. Nevertheless, we can add a premise to the argument as follows:

² Standard logic is monotonic. Increasing the amount of information available as premises never leads to losses of correctly drawn conclusions. (Costa, 1999, 33)

- 1. Birds Fly
- 2. Tweedy is a bird
- 2b. Tweedy is a Penguin
- 3. Tweedy Flies

In this modified version of the argument, the conclusion we once held now seems unwarranted. The addition of further information in the form of an additional premise now obstructs the conclusion that Tweedy flies. This very simple situation illustrates the way in which the monotony principle fails to account very ordinary common sense reasoning.

The retraction of beliefs can occur due to a myriad of reasons, ranging from cognitive psychological biases that interfere with optimal reasoning to practical constraints that might shape the depth and precision of our inferences. For example, Strasser & Antonelli (2018)³ identify at least three broad ways in which conclusions drawn from our everyday reasoning may fail. First, our inferences may follow a general typicality rule, e.g. birds fly, but a specific bird may fail to possess such property. Second, an explanation for a given observation may be corrected in light of new information, which better explains the same phenomenon. Third, there are cases in which our conclusions or inferences are probabilistic in nature.

Despite the truism of pointing out the fallibility of human reasoning, historically, the analysis of such elusive reasoning modalities, such as fallible, retractable and, the often-skewed character of human commonsense reasoning were seen as inferior forms of reasoning. These "defective" forms of reasoning processes were held as undesirable. Following this line, Koons claims:

According to Aristotle, deductive logic (especially in the form of the syllogism) plays a central role in the articulation of scientific understanding, deducing observable phenomena from definitions of

³ Defeasible reasoning is dynamic in that it allows for a retraction of inferences. Take, for instance, reasoning on the basis of normality or typicality assumptions. [...] Another example is abductive reasoning. Given the observation that the streets are wet we may infer the explanation that it has been raining recently. However, recalling that this very day the streets are cleaned and that the roof tops are dry, we will retract this inference. [...] As a last example take probabilistic reasoning where we infer "*X* is a *B*" from "*X* is an *A* and most *A*s are *B*s". Clearly, we may learn that *X* is an exceptional *A* with respect to being a *B*. (Strasse & Antonelli, 2018, 2)

natures that hold universally and without exception. However, in the practical matters of everyday life, we rely upon generalizations that hold only "for the most part", under normal circumstances, and the application of such common sense generalizations involves merely dialectical reasoning, reasoning that is defeasible and falls short of deductive validity. Aristotle lays out a large number and great variety of examples of such reasoning in his work entitled the Topics. Investigations in logic after Aristotle (from later antiquity through the twentieth century) seem to have focused exclusively on deductive logic. This continued to be true as the predicate logic was developed by Peirce, Frege, Russell, Whitehead, and others in the late nineteenth and early twentieth centuries. (2017, 3)

Thus, the fallible mode of reasoning was largely dismissed and unattended for infallible reasoning, i.e. the type of reasoning in which conclusions and inferences cannot be diminished once they are established. There are several reasons that can explain such philosophical attitude towards human common sense reasoning. For example, Aristotle's view of science (which was the paradigmatic example of knowledge *par excellence*) was of an endeavour that constructs universal laws that hold no matter what (Koons, 2017)⁴. Such a view portrays the conviction that knowledge can be obtained only through infallible and demonstrative reasoning processes, such as those involved in the knowledge realm of geometry and other disciplines alike. In the case of Aristotle, such reasoning patterns that do not comply with the strictly infallible forms of reasoning or are not subject to logical analysis. As such, these forms are categorized as *rhetorical demonstrations* or *rhetorical syllogisms*. Nevertheless, Aristotle points out the elusive nature of such reasoning patterns from the perspective of a logical structure and how its value lies between the notions of truth and plausibility, which is far apart from the exact nature of the syllogism of the *Analytics*.

Nevertheless, the monotonic stance was soon found to be too narrow to have a real use in the

⁴ According to Aristotle, deductive logic (especially in the form of the syllogism) plays a central role in the articulation of scientific understanding, deducing observable phenomena from definitions of natures that hold universally and without exception. (Koons, 2017, 3)

construction of computational models of human reasoning (Costa, 1999)⁵. Thus, epistemic certainty offered by the monotonic account is far from being the norm in human reasoning. Adhering to the monotonic view puts high epistemic standard by restraining the inferences an agent can pursue despite their fallible nature (Nute, 2003). The problematic side effect of such an approach is that there is a significant fragment of our knowledge that does not comply with the monotonic assumption of the inference process.

The stable picture depicted by the monotonic restriction concerning knowledge and justification fails to account for the fact that human reasoning involves situations in which the addition of new and possibly conflicting information makes us retract the previous held beliefs and conclusions (Kraus, Lehman & Magidor, 1990)⁶. Hence, such property is too restrictive, and despite the advantage of guaranteeing the exclusion of faulty inference, by means of a strict restriction on the deducibility relation, it turns out to be highly counterproductive for understanding common sense reasoning.

So far, we have provided a brief overview of the historical antecedents of defeasible reasoning. In the next subsection, we narrow down the meaning and scope of defeasible reasoning as an object of philosophical and logical inquiry.

2.1.2 Delimitation

One first and very rough depiction of nonmonotonic reasoning is described as non-deductive inferences. According to this very broad and general view, the existence of failing conditions such as the addition of new information provides the defeasible nature to this reasoning pattern, and render these inferences as not deductive (Pollock, 1987)⁷. Nevertheless, the plain failure to adhere

⁵ [...] computer scientists have been interested (at least during the last two decades) in studying reasoning systems where the law of monotonicity fails [...] These systems can infer B when A is taken as a premise, but they might fail to infer B when additional premises supplement A. (Costa, 1999, 33) ⁶ In everyday life, however, it seems clear that we, human beings, draw sensible conclusions from what we know and that, on the face of new information, we often have to take back previous conclusions, even when the new information we gathered in no way made us want to take back our previous assumptions (Kraus et al, 1990).

⁷ For instance, inductive reasoning is not deductive, and in perception, when one judges the color of something on the basis of how it looks to him, he is not reasoning deductively. Such reasoning is defeasible, in the sense that the premises taken by themselves may justify us in accepting the conclusion,

to a deductive reasoning pattern still seems very wide and general, but it does provide the starting point of a variety of definitions and depictions of nonmonotonic reasoning (Lehman & Magidor, $1990, 1)^8$.

The need to provide a delimited account and specification of defeasible reasoning is largely due to the loose use of the term to embrace a wide range of the various so-called "elusive" human reasoning patterns. Such reasoning patterns have a great variety and need not refer to the same reasoning modality. Hence, there is a legitimate discussion as to the extent of what can be considered as defeasible reasoning (Rankin, 1988). For example, various reasoning patterns do not comply with standard classical logic, but the sole reason for failing to comply with classical logic needs not to imply such reasoning pattern is to be categorised as defeasible. Despite the commonalities that one can locate among most approaches to defeasible reasoning, there are nuances regarding the precise extent and limitations of such a concept. In line with the elusive nature of defeasible reasoning, Koons (2017) suggests a distinction between a wide and a narrow sense.

On the one hand, a narrow sense of the term defeasible reasoning would describe the inferences in which one engages with the presumption that such inference could be further retracted in light of new information, i.e. inference rules that involve some clause of what is normally the case. As Koons puts it: "*defeasible reasoning* has typically been limited to inferences involving rough-and-ready, exception-permitting generalizations, that is, inferring what has or will happen on the basis of what *normally* happens." (2017, 1). This account of defeasible reasoning would only include what Strasser & Antonelli referred to as a dynamic reasoning form involving problems of typicality (2018, 2)⁹. Adhering to this thread and in line with Pollock's (1995) approach to defeasible reasoning, Nute (1988)¹⁰ defines defeasible reasoning as a reasoning pattern according to which,

but when additional information is added, that conclusion may no longer be justified. (481)

⁸ Nonmonotonic logic is the study of those ways of inferring additional information from given information that do not satisfy the monotonicity property satisfied by all methods based on classical (mathematical) logic. (Lehman & Magidor, 1990, 1)

⁹ Defeasible reasoning is dynamic in that it allows for a retraction of inferences. Take, for instance, reasoning on the basis of normality or typicality assumptions. (Strasse & Antonelli, 2019, 2)

¹⁰ We reason defeasibly when we reach conclusions that we might be forced to retract when faced with additional information. [...] This reasoning is defeasible, but its defeasibility is not because of

inferences and conclusions that at some point or by some specific set of reasons were held as valid, could be retracted later on. In particular, Nute explains:

We reason defeasibly when we reach conclusions that we might be forced to retract when faced with additional information. [...] This reasoning is defeasible, but its defeasibility is not because of incorrectness. [...] It is the kind of "other things being equal" reasoning that proceeds from the assumption that we are dealing with the usual or normal case. Conclusions based on this kind of reasoning may be defeated if we find that the situation is not usual or normal. (1988, 251)

According to defeasible reasoning, conclusions are held tentatively, and although they might hold in a given moment, they do not escape the possibility of revision or dismissal. In line with Nute's view of defeasible reasoning, Kraus et al. (1990) depict this reasoning pattern as one in which making novel inferences involve the ability to retract the derived information in light of new data or evidence.

On the other hand, a wide sense of defeasible reasoning can refer to various forms of non-deductive reasoning pattern in which a non-deductive inference process is conducted in such a way that inferences are fallible, such as abductive and analogical reasoning or some forms of inductive reasoning. As Koon states:

This narrower sense of defeasible reasoning, [...] excludes from the topic the study of other forms of non-deductive reasoning, including inference to the best explanation, abduction, analogical reasoning, and scientific induction. This exclusion is to some extent artificial, but it reflects the fact that the formal study of these other forms of non-deductive reasoning remains quite rudimentary. (2017, 1)

This wide account of defeasible reasoning would include what Strasser & Antonelli referred to as a dynamic reasoning form involving typicality, abductive and probabilistic reasoning (2018).

incorrectness. [...] It is the kind of "other things being equal" reasoning that proceeds from the assumption that we are dealing with the usual or normal case. Conclusions based on this kind of reasoning may be defeated if we find that the situation is not usual or normal. (Nute, 1988, 251)

One problem with this approach lies in that it includes many instances of fallible reasoning which, although defeasible they have their own research programmes. For purposes of this research, we use defeasible and nonmonotonic reasoning interchangeably when a conclusion is strongly or at least sufficiently supported by its premises, but reserve the right to revise and dismiss the connection between premises and conclusions since there is not an infallible connection between them.

The revision process of the correctability of inferences is a rational and rule-following procedure. The defeasible nature of an inference does not imply that the conclusion is to some extent false, it just means that the particular inference *could* be further defeated in light of new information or new rules connecting previously held information. This depiction of defeasible reasoning as reasoning pattern provides an epistemic and formal framework for human common sense reasoning such that it does not follow a deductive and infallible pattern. Hence, we can be said to opt for the narrow sense of the concept of nonmonotonic reasoning.

So far, we have provided a delimitation of defeasible reasoning and how is understood in the context of our research. In the next subsection, we sketch an outlook of the methods and approaches pertaining to this field.

2.1.3 Methods

Defeasible reasoning has been a major concern in contemporary epistemology. Nevertheless, in the pursuit to provide a precise characterisation of this type of reasoning, one encounters the existence of a twofold approach. The bifurcation revolves around defeasible reasoning being a chapter of logic or a chapter of epistemology (Koons, 2017). Based on this two-fold approach to defeasible reasoning, one could opt for an epistemological or for a logical approach towards the subject, with equal legitimacy since both encompass different sets of motivations, methods, and objectives.

In the context of the previous dilemma, of whether one should consider defeasible reasoning as a chapter of philosophy or as a chapter of logic, it is also crucial to consider its strong connection with other interdisciplinary endeavours (for example, Artificial Intelligence has substantially helped shape defeasible reasoning). Within such interdisciplinary endeavour, it has been upheld

the view that human reasoning and the epistemic basis of such mechanism are essentially defeasible, or that at least, the best chances of constructing computational models of such reasoning, is through nonmonotonic frameworks. It is such convergence between logic, philosophy and computer science that grounds defeasible reasoning. Hence, a clear-cut demarcation between one and the other seems available if one looks at the various approaches, but not if one looks at their underlying motivations.

On the one hand, the epistemological approach is concerned with the conditions under which one can be said to possess sufficient warrant regarding some particular belief, and how such warrant is transferred under various epistemic circumstances to other beliefs. More precisely:

In the epistemological approach, defeasible reasoning is studied as a form of inference, that is, as a process by which we add to our stock of knowledge. The epistemological approach is concerned with the transmission of warrant, with the question of when an inference, starting with justified or warranted beliefs, produces a new belief that is also warranted. (Koons, 2017, 10)

This approach to defeasible reasoning focuses on the epistemological foundations of this type of inference as well as the various circumstances which can provide the basis for some form of knowledge.

On the other hand, the logical approach is concerned with the analysis of nonmonotonic consequences relations, and study their structure and formal semantics, or as Koons explains:

Logical approaches to defeasible reasoning treat the subject as a part of logic: the study of nonmonotonic consequence relations (in contrast to the monotonicity of classical logic). These relations are defined on propositions, not on the beliefs of an agent, so the focus is not on epistemology per se, [...]. (2017, 10)

This approach to defeasible inference focuses on the formal properties a given consequence relation ought to satisfy in order to both: (a) provide the formal logical properties that serve the basis of a formal framework and, (b) provide a satisfactory account of the intuitive aspects

associated to defeasible reasoning. In such a vein, Koons (2017)¹¹ points out that theoretical commitments one takes on the epistemological dimension invariably affect the logical framework; conversely, the specific logical features have a direct implication in the epistemological foundations of the framework. Thus, whether one should consider defeasible reasoning as a chapter of philosophy or as a chapter of logic, it seems they ultimately influence one another.

The above summarises some of the main aspects associated with the tools and methods by which the topic of defeasible reasoning can be approached, and how those various approaches emphasise different angles of the same phenomenon.

2.1.4 Outlook

There are various kinds of problems associated with defeasible reasoning, and overall schema to classify them following a set of common themes. The distinctions revolve around a central feature of defeasible reasoning, which is the problem of conflicting information, i.e. multiple threads pointing towards information or conclusions that cannot stand together in a non-contradictory fashion.

First, we have a distinction that revolves around where the conflict within the body of information occurs. According to this criterion, we have an external and an internal account of defeasible reasoning, which is referred to as external and internal *dynamics*. Second, we have a distinction that revolves around the extent to which the conflict could or could not be resolved. According to this criterion, we have resolvable and non-resolvable conflicts Strasser & Antonelli (2018).

According to the location distinction, we have two ways to understand the various types of problems for defeasible reasoning. On the one hand, the *external dynamics* refers to instances in which the process of retraction or correction occurs as a result of gaining access to new information. On the other hand, the *internal dynamics* refers to instances in which the process of retraction occurs not as a result of new information but rather in light of processing

¹¹ In particular, a logical theory of defeasible consequence will have epistemological consequences. [...] However, a logical theory of defeasible consequence would have a wider scope of application than a merely epistemological theory of inference. [...] a theory of nonmonotonic logic will certainly have implications for epistemology. (Koons, 2017, 10)

information already contained within a given body of knowledge (Strasser & Antonelli, 2018)¹². To better understand both types of correction or retraction processes, we are going to consider two cases.

First, let us suppose we have a belief set Γ , from which we can make inferences about a new acquaintance named Ricardo Rodriguez of whom we have limited knowledge. Let us assume that within Γ we have that Ricardo Rodriguez has a Latin American name, and all sorts claims associated with Latin American people. We want to make an inference about the native language of this new acquaintance of ours, which we do not have in our initial set Γ (for illustrative purposes we treat the feature of a native language as allowing a single native language). From Γ we might defeasibly infer that Ricardo Rodriguez native language is Spanish. Nevertheless, if we add to our initial belief set Γ a new token of information stating that Ricardo Rodriguez is Brazilian, then our initial inference of Ricardo's native language being Spanish seems to be unwarranted in contrast with Ricardo's native language being Portuguese. In this example, the retraction of a conclusion happens in light of the addition of new and previously unknown information. Thus, this is the type of case which Batens (2009) characterises as an instance of an *external* dynamic as motivating the retraction or correction process.

Now, let us suppose we have a belief set Γ ', from which we can make inferences about a new acquaintance named Carlos Alvares of whom we also have limited knowledge. As before, we assume that within Γ ' we have that Carlos Alvares was born in Latin America, to Brazilian parents (living in Brazil at the time of his birth) who gave him a Spanish name, but immediately moved to Haiti and lived all his life in Haiti, and we assume other claims associated to Latin American people. Let us suppose we want to make an inference about the native language of this new acquaintance of ours (information we do not have in our initial set Γ '). Furthermore, we are going to have the working assumption that our inference is derived from the original belief set Γ ' and no further information is added. In this example, depending on a variety of factors and how the information is processed in the context of our defeasible inference process, we might have one or

¹² Most of scholarly attention has been paid to what has been called the *synchronic* (Pollock (2008)) or the *external dynamics* (Batens (2004)) of defeasible reasoning. For this, inferences are retracted as a result of gaining new information. [...] It has been noted, that beside the external there is also a *diachronic* (Pollock (2008)) or *internal* dynamics (Batens (2004)). It is the result of retracting inferences without adding new information in terms of new premises, but due to finding out more about the given premises by means of further analyzing them. (Strasser & Antonelli, 2018)
more retractions. For example, from the fact that Carlos Alvares is a Brazilian, born in Brazil to Brazilian parents, we might defeasibly infer that his native language is Portuguese. Nevertheless, from the fact that he moved with his parents to Haiti and lived there all his life, we may come to infer that his native language is French. In this case, the number of claims and information are quite manageable, but it illustrates that with sufficient complexity, the inference process might be a retraction-prone endeavour. According to Batens (2009) what this type of situation points out is that the retraction process emerges not from the addition of new and unknown information to our initial belief set, but rather in light of the *internal* reasoning process. That is, unlike the previous situation in which the addition of further information is what elicited the retraction of the initial inference (i.e. an *external* occurrence), here the retraction process happens in light of the information that is already contained in the belief set (i.e. an *internal* occurrence).

Moving away from the source of the conflict or retraction process, there lies the question as to the extent to which the conflict itself can be resolved. In the face of conflicting information, a mechanism needs to be set in place to address the conflict. This is the *quid* of research into defeasible reasoning. According to this way to classify problems within defeasible reasoning, we would have two major types of problems: resolvable and non-resolvable conflicts (Strasser & Antonelli, 2018)¹³.

First, there is a class of problems within defeasible reasoning which, although operating on conflicting information, there are well-known conflict resolution mechanisms or principles that can address the conflict. For example, when two lines of reasoning conflict with one another, one can operate on the principle that the thread that involves the least number of inference steps ought to be preferred over longer chains of reasoning (this is commonly known as the *inferential distance ordering* algorithm). Also, there are cases in which, in light of conflicting information, we ought to prefer lines of reasoning which provide more specific information about the conclusion under dispute (this is known as the *specificity* principle) (Strasser & Antonelli, 2018). Regardless where does one stand in light of these type of problems, the overall agreement is that some solution by means of a set of principles and procedures within the framework is at hand. Consequently,

¹³ A separate issue from the formal properties of a non-monotonic consequence relation, although one that is strictly intertwined with it, is the issue of how *conflicts* between potential defeasible conclusions are to be handled. We can distinguish two types of conflict handling in defeasible reasoning [...] (Strasser & Antonelli, 2018)

although interesting by themselves these problems tend to be less controversial and disputed.

Second, and unlike the previously mentioned class of problems, we have another set of problems which are also concerned with conflicting situations, but for which a well-defined procedure to address the conflict in a satisfactory fashion is not at hand. For these types of problems, and in the absence of a well-defined mechanism or solution, various strategies are put forward. For example, in cases involving ambiguity¹⁴, there are credulous and skeptical approaches (Strasser & Antonelli, 2018). Nevertheless, the agreement on these general principles is lesser than for cases involving resolvable conflicts. Furthermore, the extent to which these strategies address the problem is less agreed upon.

It is important to mention that the previous aspects will become much clearer as it is explained later in the thesis. Nevertheless, it is important that this subsection and overall section provides a context to better understand the motives behind our research effort. Furthermore, it is paramount to mention briefly where our research stands within the discussion. In particular, our investigation can be located as an inquiry into a *non-resolvable* problem concerning the *internal dynamics*. As we stated, the full extent of this stance is developed in detail and with concrete examples throughout the rest of our research.

This concludes our brief sketch of defeasible reasoning, its extent and limitations, as well as the particular approach we take in light of the aforesaid. Moreover, we also discussed the various methods attached to the topic. All of the above serves the purpose of placing the nature of our inquiry in context. In such a vein, we proceed in the next section with our account of inheritance networks.

2.2 Basic Elements

In this section, we provide an overview of the basic notions of inheritance networks and discuss how they can be arranged to form more complex structures. In particular, we discuss the historical background of these networks as well the elementary components such as nodes and links. This

¹⁴ This is a specific and technical problem for nonmonotonic logics, which we restrain from presenting or discussing in detail at this stage (although we briefly describe it in the introduction of this thesis), as this will be one of the key elements of our research which will be discussed and addressed later on.

section aims to introduce the most basic concepts surrounding inheritance networks. The work done in this section serves as the basic input on which the remaining of our research is grounded.

2.2.1 Background

Gabbay & Schlechta (2016)¹⁵ point out that inheritance networks were introduced in the 1980s. More precisely, it is usually agreed upon that with Touretzky's work of 1984 inheritance networks as a formal framework for defeasible reasoning are born. Nevertheless, Touretzky's work follows from preceding efforts in exploring the extent and limitations of using network-like formalisms for various representational tasks. The idea of network-like formalisms in the context of cognitive science and knowledge representation in general is attributed to Quillian (Brachman, 1977). Quillian tried to "[...] to develop a theory of the structure of human long-term memory, and to embody this theory in a computer model such that the machine can utilise it to perform complex memory-dependent tasks." (1967, 410). Quillian's representational framework consisted of nodes standing for words and concepts as well as a different type of links such as conceptual inclusion or argument validity or plausible inference. This simple arrangement would, in turn, serve the purpose of reflecting human semantic memory. In particular, Brachman states:

This simple representational format would presumably allow the encoding of large bodies of factual information, and would facilitate the interpretation of new material in terms of large quantities of previously stored information ("assimilation"). In this way, it might be used to help simulate the reading and understanding of text. (1977, 127)

This type of formalism would enable us to understand more of various cognitive tasks such as storing and processing large amounts of information. Quillian's work inaugurated what would be a broad range of network-like formalisms. In this context, Touretzky's work is the first and one of the most salient in the field of non-monotonic reasoning.

The above leads to an important distinction between Quillian's notion of semantic network and other types of network-like formalisms, such as Touretzky's. The distinction relies more on the

¹⁵ Inheritance networks were introduced in the 1980s (see, e.g. Touretzky et al. 1984, 1986, 1987), and exist in a multitude of more or less differing formalisms [...] (Gabbay, & Schlechta, 2016, 78)

theoretical purposes and contexts in which the various network-like formalisms are used and how they are understood. On the one hand, networks as in *Semantic Networks* refer to the family of network-like formalisms that aim to describe and model one or more features of human semantic memory and its operations. As Carpenter & Thomason point out:

The term "semantic networks" points to a variety of graph-based formalisms that are widely used for the representation of knowledge in computational systems. These network formalisms were introduced in this computational context by Quillian (see [Quillian67]), who used them to model human associative memory and hierarchical or taxonomic reasoning. (1990, 309)

On the other hand, Carpenter & Thomason highlight that such node-plus-link network structure is not reserved to the understanding of concepts and human semantic memory, but rather can be extended to other endeavours and fields in which the focus is not entirely human semantic memory. These other fields can be Artificial Intelligence and any of its subfields, such as expert systems, non-monotonic logics, etc. To this end, Carpenter & Thomason propose the convention of talking about networks when the term is used in the narrow sense (1990)¹⁶. Although initiated through Quillian's work, the term *network* refers to a class or family of frameworks that share the same motivation. Besides having the common nodes plus links in conjunction with inheritance as the "backbone", these networks are mostly thought as forms of representing and operating with some form of knowledge. In such a vein, Brachman states:

Many systems for representing knowledge can be considered semantic networks largely because they feature the notion of an explicit taxonomic hierarchy, a Tree or lattice-like structure for categorizing classes of things in the world being represented. The backbone of the hierarchy is provided by some sort of "inheritance" link between the

¹⁶ Semantic networks have many uses and associations, including the cognitive ones that first suggested them. Here, however, we are mainly interested in their use in relatively small "frame-based" knowledge representation systems for applications in artificial intelligence, primarily in expert systems; we will use the less pretentious term "network" to indicate this focus of what we have to say. (Carpenter & Thomason, 1990, 309)

representational objects, known as "nodes" in some systems and as "frames" in others. (1983, 30)

Given that our research is constrained to defeasible reasoning, we use the term inheritance networks to refer to this type of network-like formalisms as applied to the field of non-monotonic reasoning (in some cases we just refer to them as networks). We reserve the term "semantic network" to refer to Quillian's type of network-like formalisms. This convention follows the thread of the most active work (mainly throughout the '80s) that was initially done with this type of frameworks for this particular aim.

Moving on from the previous discussion, it is worth mentioning some of the reasons as to why this type of network-like structures had some significant level of success within various fields. As Gabbay & Schlechta pointed out, one of the advantages of these representational structures, lies in that, besides sharing many similarities with logic-based approaches, they seem to provide an intuitive and straightforward bridge to the cognitive aspects of information processing in human agents (2010)¹⁷.

Another attractive feature associated with these types of networks is that they have a "natural correspondence" with the mathematical structure known as graphs. This natural correspondence is useful inasmuch as there is already ample study in how information can be represented and extracted within graphs through mathematical procedures (Thomason, Horty & Touretzky, 1987)¹⁸.

So far, we have provided a succinct description of some of the key ideas and aspects that provide a better picture and background of the formal framework we use throughout our research. This is far from all that can be said about the subject matter, as extensive research and discussion has expanded the field in such a way that a whole dissertation can be devoted to such aim. Nevertheless, our object of inquiry *can be expressed in the framework* due to its simplicity, but

¹⁷ Inheritance systems or diagrams have an intuitive appeal. They seem close to human reasoning, natural, and are also implemented (see [Mor98]). (Gabbay & Schlechta, 2010, 251)

¹⁸ One traditional attraction of inheritance networks has always been their natural correspondence with graphs, which makes them particularly appropriate as vehicles of knowledge representation for concurrent computing architectures, where graph-searches can be very fast. (Thomason, Horty & Touretzky, 1987, 1)

our research is not *about the framework* itself. Thus, a compromise is made, and a more detailed theoretical discussion of the concepts and ideas related to inheritance networks is undertaken when necessary or when its inclusion enhances the discussion at hand. In the following sections, we examine the various elements that compose this type of networks.

2.2.2 Nodes

Nodes can stand for a particular object or entity, or they can stand for attributes, classes or properties. Although a distinction is often made between individuals and properties, we consider them part of the same class of elements. It is also important to mention that we follow the convention used by Thomason et al. (1987), which despite some minor variations is to a very large extent the default notational convention; when we make substantial deviation, we will point such divergence.

If we are dealing with the domain of philosophers, we can have a set of nodes to stand for each of the philosophers. In this vein, we would have a node that stands for Peter Singer, Susan Haack, Simon de Beauvoir, and so on. To refer to this use of nodes, we use lower letters from the beginning of the alphabet, such as *a*, *b*, *c*, *etc*. If we are dealing with the domain of German philosophers of the 19th century, we can have a node standing for the property of having a moustache, and we can have another node standing for the attribute of being bald, or a node standing for the class of being European. To refer to this use of nodes, we use letters from the middle of the alphabet such as *p*, *q*, *r*, etc. Given that we want for our language a minimal level of generality, we use letters from the end of the alphabet (*u*, *v*, *w*, *x*, *y*, *z*) as variables for both objects and properties.

The exact nature of nodes in the context of inheritance networks has been an object of debate. In particular, it has been discussed the extent and limitations to which the various interpretations can be conflated (Woods, 1975). Nevertheless, we restrain ourselves from inquiring into the full extent of this aspect and side of the discussion, as it is not, the aim nor purpose of our research.

The brief account of the notion of nodes serves as an overview of an intuitive concept that is of extensive use throughout our research. It is important to emphasise that, although there is ample room for a thorough discussion of the epistemological, methodological and metaphysical aspects surrounding the notion of nodes, we do not undertake such debate.

2.2.3 Links

Positive and negative links provide the two fundamental ways in which a given set of nodes can be connected. These links provide the basic expressive power within this type of frameworks. The notions of links and nodes are the two components that enable, with some minor caveats, the construction of the whole edifice of inheritance networks. For example, we can imagine a situation describing philosophers. In such context, we can have the node *a* standing for 'Susan Haack', and the node *p* standing for the 'philosopher' attribute. In this case, we would naturally want to express that Susan Haack is linked to the philosopher property. The aforesaid is expressed as $a \rightarrow p$. Thus, the positive link denoted by the \rightarrow (arrow) is referred to as the *isa* relation. The *isa* link stands for any positive connection between any two nodes within a given context. Furthermore, we can imagine a situation describing philosophers. In such context, we can have the node *b* standing for Suzana Herculano-Houzel, and still having *p* standing for the 'philosopher' attribute. In the aforesaid context, we would want to express that Suzana Herculano-Houzel is not a philosopher (she is a neuroscientist). The aforesaid would be expressed as $b \not\rightarrow p$. Thus, the negative link denoted by the ($\not\rightarrow$) (arrow with a stroke) is referred to as the *not-isa* relation. The *not-isa* link stands for a negative connection between any two nodes within a given context.

These networks are used as a representational framework involving statements, more precisely, what Thomason et al. (1987) have referred to as *assertions*. In this sense, we are taking into account both positive (*isa*) and negative (*not-isa*) links $x \rightarrow y$ and $x \not\rightarrow y$ both stand for the claims of x being or not a y. Positive and negative links between nodes are both interpreted as positive and negative assertions.

Given these formal frameworks aim to model a wide range of human common sense reasoning, they need to incorporate some mechanism to handle defeasible features of human reasoning. The defeasible trait of human reasoning stands in opposition with what occurs in very restricted settings, such as mathematical reasoning in which the stable feature of monotonicity is inherent (Thomason & Horty, 1989)¹⁹. Consequently, for inheritance networks to be taken as a serious

¹⁹ Familiar logical theories, including intuitionistic, modal, and higher-order logics, were inspired by the need to construct theories of mathematical reasoning. And in mathematical reasoning, assumptions must be made explicit, and theorems must be reusable in proving later results. Monotonicity runs deep in these theories, and incremental modifications to standard logics are not likely, we believe, to provide

framework for defeasible reasoning they need to handle such defeasible traits, like the typicality or normality of certain attributes or properties that may belong or can be used to describe a certain individual, but for various reasons, may fail on some particular cases, i.e. the exception conditions.

Thus, inheritance networks need to include a mechanism that can handle standard features such as nodes, properties and attributes, as well as being able to operate with these components in a defeasible framework. To account for the above, defeasibility is built-in inheritance networks at the most basic level, i.e. through the *isa* and *not-isa* associations. In such a vein, we make a consistent reading of both *isa* and *not-isa* links as defeasible or retractable links. Under this account, the direct connection between any two given nodes stands for an association between nodes that *typically* hold, but it is open to exceptions *without* the typicality being lost. As Carpenter & Thomason point out:

These sentences may be maintained even in the face of exceptions. That is, the statement that birds fly expressed above has all the earmarks of truth, and remains useful for reasoning, when *ad hoc* exceptions are known - particular birds that don't fly - and even when general exceptions are admitted - atypical birds, such as ostriches and penguins, that are naturally flightless. (1990, 323)

In such a vein, we are interested in modelling and representing *defeasible assertions* (also referred to as generic statements). This type of statements holds, even when an exception to a given claim can be found (Gabbay & Schlechta, 2010). For example, the claim 'Birds fly' is a generic statement that holds even a specific type of birds, i.e. penguins, do not fly (Thomason et al. 1987). The above stands in direct opposition to statements that hold without exception, such as definitions (Carpenter & Thomason, 1990). Thus, the kinds of assertions we consider are defeasible assertions. We do not take into account strict or indefeasible assertions. This, of course, is a discretionary decision, and it is worth mentioning that research has been done for strict networks (Thomason et al. 1987) as well as networks which include both strict and defeasible assertions (Horty, 1990). Nevertheless, this distinction has no relevant impact for the purposes of our research objectives.

tractable characterizations even of apparently simple applied concepts such as nonmonotonic "IS-A." (Thomason & Horty, 1989, 221)

Furthermore, we can point out to one further important aspect, associated with both *isa* and *not-isa* links. For example, if we have that $x \rightarrow y$, we will say that x is the *starting node* and y the *terminal node*. Likewise, the above holds for *not-isa* links. As a corollary of the aforesaid, we say that both *isa* and *not-isa* have a sense of directionality. For example, if we have $w \neq z$, then w is not a z, but not the other way around, i.e. $z \neq w$ does not hold. Thus, we say that links have a sense of directionality since cycles are not considered. The aforesaid is a corollary of our previous depiction of both *isa* and *not-isa* links that includes the notion of a starting and a terminating node. Although cycles can be considered, their inclusion add a considerable amount of complexity. Moreover, almost all research has focused on directed acyclic networks.

This concludes our brief exposition of the notion of a node within the framework of inheritance networks. In the next subsection, we discuss how the concepts of nodes and links build the basic and elementary unit known as a network or simply a *net*.

2.2.4 Nets

The notion of a *network* or simply a *net*, which encompasses all of the components discussed so far. A net Γ is a triplet $\Gamma = \langle I, P, N \rangle$, such that *I* is a set of individuals and properties, together with a set *P* of positive *isa* links (\rightarrow) and a set *N* of negative *not-isa* links (\rightarrow) (both *P* and *N* are subsets of (*I x I*)). In the literature it is customary to identify a network Γ with just the set of positive and negative links contained in it (Carpenter & Thomason, 1990). For most purposes, the above should not represent major confusion or problem; nevertheless, where there is room for ambiguity, we will make the distinction explicit. Otherwise, we will be licentious.

As we mentioned earlier, an inheritance network can be equated with a *graph*. A graph is a mathematical structure that is composed of nodes (vertices in graph terminology) and links (edges in graph terminology). Naturally, graphs as a mathematical structure are more general than graphs as a representational framework for defeasible reasoning. We consider two elementary variations that we do not take into account in the context of inheritance networks. In the previous subsection, when we considered links (edges in graph terminology), we stressed the idea that the connection between any two given nodes presupposes a starting and terminal node. Nevertheless, this is just a restriction of a more general case of what is known as *undirected graphs*.



Figure 2.2.4-a Directed and Undirected Graphs

For example, the network Γ in Figure 2.2.4-a presents a graph inasmuch as x is connected to y, and the former is the starting node, whereas the latter is the terminal node. Even though Γ ' depicts a situation in which x' and y' are connected, in this case, there is no sense of a starting or terminal node but rather just a connection between both nodes. The former graph is a directed graph whereas the latter is an undirected graph. As we previously mentioned, inheritance networks are taken as directed rather than undirected graphs. Nevertheless, this is just a constraint within the larger set of structures known as graphs.

A second significant distinction consists in the absence of loops (also referred to as cycles) within the type of graphs we have considered so far. To illustrate the idea of a loop within a graph, we can take the following networks:



Figure 2.2.4-b Cyclic and Acyclic Graphs

As portrayed in Figure 2.2.4-b, both the graph Γ and Γ ' containing the nodes x, y, z and x', y', z' respectively are instances of directed graphs. Nevertheless, unlike Γ , the graph Γ ' contains a cycle. In particular, we have that x' is connected with z' through the $x' \rightarrow z'$ edge, but we also have that z' is connected to x' through the $z' \rightarrow x'$ edge. To put it another way, the acyclic feature refers to the fact that the $(I \times I)$ set has an antisymmetric ordering relation. This feature is what is known as a loop or a cycle within a graph. Unlike the feature of being directed, the presence of cycles is not precluded in the context of inheritance networks.

There are inheritance networks that allow the existence of cycles. The admission of cycles opens the door for problems and debates of its own, which adds another level of complexity to graphs as representational tools for defeasible reasoning. The problem that lies at the centre of our research does not need cycles to be instantiated. That is, the problem we are examining emerges in networks without this additional layer of complexity that needs not to worry us. Such variations illustrate how inheritance networks rely on more or less restricted forms of graphs as a mathematical structure, but in the end, they are just a particular specification of a more general mathematical structure. With the addition or restriction of features, we get different and idiosyncratic versions of inheritance networks. For simplicity purposes and to focus our attention on the main objectives of our research, we work with directed acyclic graphs only.

To summarise, inheritance networks, as used throughout our research are a specific type of graphs, namely, *directed acyclic graphs* (Brewka, 1991), for three specific reasons. First, inheritance networks are graphs because they are a collection of nodes and links that can be translated to the notions of vertices and edges for graphs. Second, inheritance networks are directed graphs since any association between any two nodes (or vertices) by means of a given link (or edges) has a form of directionality, i.e. there is a starting node (vertex) and a terminal node (vertex). Third, inheritance networks are acyclic because cycles within the network are not allowed. Having stated the above, it needs to be said that there are variations within inheritance networks, which allow one or more features we stated as restricted.

Before proceeding to the next subsection, we recall that, in the first section of this chapter, we highlighted defeasible reasoning as one of the key fields which have taken inheritance networks as a research tool. We stressed the imprint placed by defeasible reasoning to these network-like formalisms apart from Quillian's semantic networks. In such a vein, formal frameworks that aim to model and describe a wide range of human common-sense reasoning need to account for some

mechanism to handle defeasible features of human reasoning. Consequently, for inheritance networks to be taken as a satisfactory framework to address this task, they need to handle defeasible traits, such as typicality or normality of certain attributes or properties that may belong or can be used to describe a certain individual, but for various reasons, may fail on some particular cases, i.e. the exception conditions.

Inheritance networks include not only a mechanism that can handle standard features such as nodes, properties and attributes, but also operate with such components in a defeasible framework. To account for the above, defeasibility is included in inheritance networks at the most basic level through the *isa* and *not-isa* associations. In this sense, we have made a consistent reading of both *isa* and *not-isa* links as defeasible or retractable associations. Thus, the direct connection between any two given nodes stands for an association between nodes that *typically* hold, but it is open to exceptions *without* the typicality being lost.

The extent and limitation of the relation between inheritance networks and defeasible reasoning is far deeper, and it would demand thorough research by itself. Here we limit ourselves to highlight the connection between both endeavours as underlying motivation throughout our work. This resurfaces in the final chapters on the thesis when we undertake the epistemological aspects of inheritance networks.

Having provided a picture of both positive and negative links, as well as the core concept of a net, in the following section we discuss the concept of a path, which builds from the notions of links. In particular, we look into the notion of paths as a proxy to subsequent features that will enable us to set in place some information extraction mechanisms within this type of structures.

2.3 Paths

In this section, we extend the notion of links into the idea of a path. This concept offers us a more complex, but still an elementary structure that allows us to extract information from inheritance networks. The aim of this section is to depict a basic concept that serves as a building block for the forthcoming sections and chapters.

2.3.1 Positive Paths

So far, we have that a given arrangement of nodes (i.e. the set I of individuals and properties), are connected through *isa* (\rightarrow) and *not-isa* ($\not\rightarrow$) links. More precisely, we have that (\rightarrow) and ($\not\rightarrow$) are subsets of (*I x I*). A path is just an extension of this idea applied to a set of links. More precisely, a path is a sequence of links, within the vast arrangement of connections between nodes, such that we can follow the trace from a given initial node to a specific terminal node.

To be more precise, a path can be defined as a sequence of positive or negative links that can be traced from an initial node to a terminal node, within some particular structure. Thus, if we have the nodes $\{w, x, y, z\}$ and the relations $\{w \rightarrow x, x \rightarrow y, y \rightarrow z\}$ for a given network Γ , we can state that there is a *path* that links *w* to the node *z*. Starting with *w* we can trace a sequence of links up to *z*. Likewise, we can state that there is a *path* that links *w* to the node *x*.



Figure 2.3.1 Paths with Positive Links

As portrayed in Figure 2.3.1 we simply state that $w \to x \to y \to z$, as well as $w \to x \to y$ and $w \to x$ are all paths of Γ as defined by the *isa* (\rightarrow) or *not-isa* (\rightarrow) relation accordingly. With respect to the path $w \to x \to y \to z$, we have that w is the starting node and z is the terminal node. Furthermore, x and y are intermediary nodes. With respect to the path $w \to x \to y$, we have that w is the starting

node, y is the terminal node, and x the only intermediary node. Finally, with respect to the path $w \rightarrow x$, we have that w is the starting node, x is the terminal node. In this last case, there are no intermediary nodes, as this is the simplest form a path can take, i.e., a sequence of one link. Following Carpenter & Thomason (1990), we use lower case Greek letter σ , τ , ... etc., to stand for paths within a given network. Thus, in the previous example, we can refer to each of the paths by and indexed lower case σ , as follows:

$$\sigma_1: w \to x \to y \to z,$$

$$\sigma_2: w \to x \to y$$

$$\sigma_3: w \to x$$

Before proceeding with negative paths, it is important to notice that paths, by definition, are understood as a sequence of links. In such vein, and strictly speaking, from the previous network, we could have selected x as the starting node. From such choice, we would have the paths $x \rightarrow y$ and $x \rightarrow y \rightarrow z$, as they comply with our definition of a path. Nevertheless, and for the purposes of our discussion, we follow the convention that the starting node would be w rather than x. Again, rather than a technical feature, the above is just a convention that makes our assessment of the networks of our interest easier to examine and process, with any loss of generality.

Having stated the above, we conclude our discussion of the simplest form of a path for inheritance networks. As we have depicted, a path is an extension of the concept of a link. Nevertheless, links can contain not only positive assertions but also negative assertions. The above is the topic of the following subsection.

2.3.2 Negative Paths

Unlike paths that contain only positive *isa* links, paths containing negative links vary depending on where the negative link is located within a larger series of links. In particular, we have three cases, (a) the negative link is located between an intermediary and terminal node of the route, (b) the negative link is located between two intermediary nodes within the route, or (c) the negative link is located between the starting and an intermediary node of the route.

In the case of (a), since the negative link is allocated just at the end of the route between the starting

and terminal node, we state that we have a legitimate path. In particular, such path would represent that the starting node is negatively connected to the terminal node.



Figure 2.3.2-a Path with a Negative Link

For example, let assume we have the set of nodes $\{w, x, y, z\}$ and the relations $\{w \to x, x \to y, y \neq z\}$ for a given network Γ as depicted in Figure 2.3.2-a. In this context, and with respect to the *z* property, i.e. the $w \to x \to y \neq z$ path, we have that *w* is the starting node, *z* the terminal node and both *x* and *y* are intermediary nodes. Furthermore, the nodes are connected through *isa* links and a *not-isa* link, such that the negative link is at the end of the route, i.e. between the last intermediary node and the terminal node of the path. This is a case in which we have a path involving a negative link, which negatively connects *w* with *z*. In the context of this network, we also have the positive paths $w \to x \to y$ and $w \to x$.

In the case of (b), the negative link would be located between two intermediary nodes. Unlike (a), the negative link is placed before the terminal node is reached. Thus, we say that the route between the starting and terminal node does not provide any information *with respect to the terminal node* z, since it is discontinued by the existence of a negative link before it reaches the terminal node.



Figure 2.3.2-b Interrupted Negative Path

For example, let assume we have the set of nodes $\{w, x, y, z\}$ and the relations $\{w \rightarrow x, x \not\rightarrow y, y \rightarrow z\}$ for a given network Γ as described in Figure 2.3.2-b. As in the previous example, and with respect to property *z*, we would have that *w* is the starting node, *z* the terminal node and both *x* and *y* intermediary nodes. Nevertheless, and unlike the previous example, we have that, with respect to the *z* attribute, the *not-isa* link is allocated before the end of the route. Hence, the *not-isa* link is located between two intermediary nodes, namely, between *x* and *y*. In this case, we are said not to have a path that connects the nodes *w* and *z*. Furthermore, the negative link does not provide any information *with respect to the terminal node z* — the reason being that the potential path is *interrupted* before it reaches the terminal node. The above being said, with respect to the *y* property we do have a negative path, namely the path that has the starting node *w* and goes through the intermediary node *x* and from such intermediary node it finds a *not-isa* connection to the *y* attribute. Following this last thread, we would also have the positive path that has the starting node *w* and the terminal node *x*.

In the case of (c), the negative link would be located between the starting node and an intermediary node. Following the same dynamic as in (b), we have that unlike (a), the negative link is placed

before the terminal node is reached. Thus, we say that the route between the starting and terminal node does not provide any information *with respect to the terminal node z* since the existence of a negative link discontinues it before it reaches the terminating node.



Figure 2.3.2-c Route Starting with a Negative Link

For example, let us assume we have the set of nodes $\{w, x, y, z\}$ and the relations $\{w \nleftrightarrow x, x \rightarrow y, y \rightarrow z\}$ in a given network Γ , as described in Figure 2.3.2-c. Like in the previous example, with respect to the property z we have that w is the starting node, z the terminal node and both x and y intermediary nodes. Hence, we have that the *not-isa* link is allocated before the end of the route. In this case, the *not-isa* link is located between the starting node and an intermediary node, namely, between w and x. Under these circumstances, we are said not to have a path, *with respect to the terminal node z*, that connects the nodes w and z. The reason being that, the potential path is *interrupted* before it reaches the terminal node. Moreover, we would also have that the path with w as the starting node and y as the terminal node would also be blocked for the same reasons. The above being said, we do have the negative path that has w as the starting node and x as the terminal node. In this last case, we would have the negative path $w \nleftrightarrow x$.

The above summarises the notion of a path that takes into account negative links. In particular, a path such that it includes not only positive but also negative links can be defined as a sequence of positive links that can be traced from an initial to a terminal node, within some particular structure.

On the one hand, if the path contains only positive links, i.e. we only have positive links from start to finish, then there is a path between the starting node and terminal node. On the other hand, if we have a route in which all of the links are positive, except the last link which is negative, then we have a path that negatively connects the starting node to the terminal node. Any other case is not counted as a path with respect to the terminal node.

This concludes the discussion of the notion of a path such that it accounts for negative associations. This concept, whether it involves only positive or negative links, is an elementary component upon which we extract further information. This last task is the aim of the next section.

2.4 Extracting Information

In this section, we examine one of the fundamental mechanisms through which nodes can be attributed non-immediate properties. We present the notion of *inheritance* as the process for extracting information in a given network. This section aims to put forward a key notion from which more complicated concepts follow. The main purpose being the discussion of the concept of inheritance in a section of its own to avoid burying its importance in the presence of other notions.

2.4.1 Inheritance

So far, we have seen the arrangement of nodes standing for individuals on the one hand and properties on the other. Furthermore, we have noted the two types of connections available for all nodes of a given network, namely positive *isa* and negative *not-isa* type of links. Following such thread, we extend our understanding of inheritance networks by providing a simple and mechanical inference process to extract relevant information. This process is known as *inheritance*.

Inheritance is the procedure that extracts knowledge from a given network structure by ascribing properties to individuals based on the positive or negative links that can be traced between two nodes. The main objective of any account of inheritance for network-like formalisms revolves around providing the criteria under which a given assertion is to be inferred from a given network. As Horty, Thomason & Touretzky state:

Since we identify the links in a net with assertions, a net can be viewed as a set of hypotheses, or axioms. Let us say, informally for a moment, that an assertion A is *supported* by a net Γ if we can reasonably conclude that A is true whenever all the links in Γ are true--if the information contained in Γ would naturally lead to the conclusion that A. Our object here is to explicate this informal idea. We want to know exactly what we can conclude from a given net; so we need to provide a formal account of the conditions under which a net Γ supports an assertion A. (1990, 314)

These authors equate this problem to what occurs in deductive logic and use what they see as the "roundabout strategy" to characterise the inference relation. In particular, they appeal to the process that first determines the arguments or lines of reasoning and then characterises as deducible the information that follows from such lines of reasoning. This indirect account of the inference relation, as they see it in logic, is what lies behind the idea of inheritance for networks. Horty et al. (1990) explain this as follows:

In the context of ordinary deductive logic, we often find ourselves in a similar situation, when we want to know what statements are *deducible* from a given set of hypotheses. In that context, it is a common practice to approach the question in a roundabout way. Instead of defining the relation of deducibility directly, one first characterizes the deductions--sequences of statements representing certain kinds of arguments, or chains of reasoning--and then defines a statement as deducible from a set of hypotheses if those hypotheses permit a deduction of that statement. (314)

In the context of inheritance networks, the above is applied by first determining the paths that are available from a given network. Such paths are seen as the arguments or "chains of reasoning". Thus, the concept of inheritance relies on the notion of a path to determine which attributes can be inferred from a given network. In particular, the network supports the attributes that are endorsed by the paths associated with such network, or as Horty et al. (1990) put it:

Instead of trying to specify directly the statements supported by a given

net, we first characterize the arguments or chains of reasoning-represented, now, by paths--that are *permitted* by a net. As in the case of ordinary deducibility, this relation between sets of hypotheses and the chains of reasoning they permit is really the central idea; and it will be the primary focus of our attention. Once we have identified the paths that a net permits, it is natural to define the statements supported by a net by stipulating that a net *supports* a statement just in case it *permits* a path *enabling* that statement. (314)

Unlike previous notions, the idea of inheritance is relatively less problematic and debated. That is not to say that the notion of inheritance is exempt from discussion, but such discussion is lesser due to the very straightforward idea it carries. In particular, the concept of inheritance only makes explicit a simple intuition that is already contained in the notion of a path. In particular, we have two fundamental instances of inheritance: single and multiple inheritance.

Before proceeding with the rest of the subsection, we point out to an important feature that we employ from now on and throughout our research. Adhering to what we have stated in the previous sections, we have that the set I of individuals and properties is composed of lower-case letters. Nevertheless, to make our examples more intuitive and easier to place in context, we relax the aforesaid restriction and allow a string of characters to stand for individuals and properties.



Figure 2.4.1 Basic Inheritance Network

Thus, under this convention, and following the network described in Figure 2.4.1, we have the nodes $\{Cat, Mammal, Animal\}$ and the relations $\{Cat \rightarrow Mammal, Mammal \rightarrow Animal\}$. This very simple network serves the purpose of showing how following a more liberal notation, the representational structures can gain in explanatory and intuitive appeal. Furthermore, this convention is the norm in the literature.

The above concludes our introduction to the notion of inheritance within the family of networklike formalisms under scrutiny. Nevertheless, as it stands, it is expressed as an abstract idea. In the following two subsections we draw from the aforesaid intuitions to ground this inference mechanism. In particular, we review each form of inheritance (single and multiple inheritance), provide some examples and discuss some of their salient features. Single inheritance is the case in which there is a unique route that connects a starting node to a terminal node, by passing through a series of intermediary nodes. In the case of single inheritance, the main feature is that a given starting node or intermediary node is linked to a unique node at most and no more than one other node. Thus, in this type of structures, any node can be connected to one and only one other node. Hence, this type of structures offers a very linear pattern that starts from an individual and moves towards to the properties or attributes to which such node can be linked.



Figure 2.4.2-a General Single Inheritance Network

To address single inheritance, let us assume a given network Γ with the nodes $\{w, x, y, z\}$ and the relations $\{w \to x, x \to y, y \to z\}$, as described in Figure 2.4.2-a. In this context, we have the following paths:

$$\sigma_1: w \to x$$

$$\sigma_2: w \to x \to y$$

$$\sigma_3: w \to x \to y \to z$$

In line with the above, we have that, with respect to *z*, *w* is the starting node, and *z* is the terminal node, and there is a path that links *w* with the property *z*, namely the path σ_3 . Thus, in this context, we have that *w* inherits attribute *z*. Furthermore, with respect to the attribute *y*, we have that *w* also inherits such attribute, namely through the σ_2 path.



Figure 2.4.2-b Basic Single Inheritance Network

The network Γ described in Figure 2.4.2-b portrays such a linear structure in which we have the nodes {*Socrates, Greek, Human, Mortal*} and the relations {*Socrates \rightarrow Greek, Greek \rightarrow Human, Human \rightarrow Mortal}. Thus, we have the following paths:*

 $\sigma_1: Socrates \to Greek$ $\sigma_2: Socrates \to Greek \to Human$ $\sigma_3: Socrates \to Greek \to Human \to Mortal$

In this example, with *Socrates* as starting node, and with respect to the *Mortal* node, we can reach the aforesaid attribute through a series of *isa* links. Based on the notion of inheritance, and in light of the previous structural arrangement, we can state that *Socrates* inherits or has the attribute of

being a *Mortal*. We also have that concerning the same starting node *Socrates*, and with respect to the human attribute, we have the σ_2 path. Based on such path, we state that the *Socrates* node inherits the *human* attribute. The same process applies for the direct link between the *Socrates* node and the *Greek* attribute, which is supported through the σ_1 path.

Throughout our research, we are going to frequently refer not only to a specific network, but more generally to the particular way in which the set of nodes and links associations are instantiated in a given network, i.e. how the network is organised. To this end, we employ the notion of a *structural arrangement* or simply a *structure* to refer not just to a specific network, but to the specific arrangement instantiated by such network.

Having stated the above, we want to draw the attention towards the fact that, any given node of the previous structures is connected to one and only one other node within the same network (through an *isa* link). For example, the *Socrates* node is only connected to one other node (the *Greek* node), which in turn is connected to only one other node (the *Human* node), and so on. This specific feature is what distinguishes these networks as only involving single inheritance. Furthermore, single inheritance equally applies to structures involving negative links.



Figure 2.4.2-c General Inheritance Network with a Negative Link

To inspect the notion of inheritance of properties in networks involving negative links, let us assume a network Γ with the nodes $\{w, x, y, z\}$ and the relations $\{w \to x, x \to y, y \neq z\}$, as described in Figure 2.4.2-c. In this context, we have the following paths:

$$\sigma_1: w \to x$$

$$\sigma_2: w \to x \to y$$

$$\sigma_3: w \to x \to y \not\to z$$

In line with the above, we have that *w* is the starting node and *z* is the terminal node, and there is a path that negatively links *w* with the property *z*, namely the path σ_3 . Thus, we say that *w* explicitly fails to inherit attribute *z*. Again, it is worth noting that Γ provides grounds to endorse the attribute *y* to the starting node *w*, namely path σ_2 . By the same token, Γ also provides reasons to endorse the attribute *x* to the starting node *w*.

It is crucial to point out that this notion of failing to inherit an attribute, in the context of negative links, is different from what is known as *Negation as Failure* in computational logic, which is a form of negation that occurs in the absence of any information to suggest otherwise. To the contrary, here we operate with what is referred to as a *strong negation* (Gabbay & Shclechta, 2016). Thus, in the context of inheritance networks, what we mean by failure is the explicit claim of the negative connection between two links.



Figure 2.4.2-d Basic Inheritance Network with a Negative Link

Figure 2.4.2-d depicts the network Γ with the nodes {*Nixon*, *Republican*, *Pacifist*} and the relations {*Nixon* \rightarrow *Republican*, *Republican* \neq *Pacifist*}. Thus, we have the following paths:

 $\sigma_1: Nixon \rightarrow Republican$ $\sigma_2: Nixon \rightarrow Republican
\leftrightarrow Pacifist$

In this example, we start with the individual node *Nixon* and follow through an *isa* link that connects it to the *Republican* node, and from such node, it connects it to the *Pacifist* node through a *not-isa* link. Thus, based on the notion of inheritance, we can state that *Nixon* explicitly fails to inherit the *Pacifist* attribute, namely through the σ_2 path. In addition, we have that in virtue of the σ_1 path, the starting node *Nixon* does inherit the *Republican* property. Like in the previous examples, we have that any given node of the previous structure is connected to one and only one other node within the same network, through an *isa* link. For example, the *Nixon* node is only connected to one other node, namely, the *Republican* node, which in turn is connected to only one other node. Again, the inclusion of negative links does not change the structural feature of single inheritance.

Although single inheritance structures capture an important class of expressions and inferences, they are not expressive enough to implement a larger class of scenarios. Hence, such structures are restricted frameworks that can only represent the most basic of situations. The next section examines inheritance networks with further expressive power.

2.4.3 Multiple Inheritance

As a more ample and extended class of inheritance structures, which builds on top of single inheritance, we have multiple inheritance. This structural arrangement offers a much more expressive capability to inheritance networks.

Multiple inheritance refers to cases in which a given node can have either an *isa* or a *not-isa* link with not only one, but with various other nodes. In this type of situations, there is a connection between a starting node and a terminal node but, unlike cases of single inheritance, any starting or intermediary node can be linked to more than one other node, through an *isa* or a *not-isa* connection. This stands in contrast to cases of single inheritance in which any given node cannot be linked to more than one other node. Thus, multiple inheritance follows the same linear pattern of single inheritance structures, with the caveat that such linearity can be shared among nodes. As Carpenter and Thomason observe:

Both the definition of inheritance given here [...] are cases of multiple inheritance. What this means is that there may be links from one node to more than one other node. For instance, we might want to say that whales are mammals and that whales are ocean dwellers. Note that nothing in the definitions prohibits such uses of multiple classification. One use of such multiple inheritance is to simultaneously classify objects along a number of different dimensions. For instance, consider the simple classification of animals in terms of their biological class and habitat [...]. (1990, 319)

This feature provides greater expressive power in cases where items need to be classified among a variety of different properties. As the previously mentioned authors state, a concept is more often than not, related not to one but multiple other concepts in parallel at any given time. In this situation, both the concept itself as well as its relationship with other items are needed to capture the phenomenon under consideration.



Figure 2.4.3-a General Multiple Inheritance Network

To illustrate multiple inheritance, let us assume a network Γ with the nodes {*w*, *x*₁, *x*₂, *y*₁, *y*₂, *z*₁, *z*₂} and the relations { $w \rightarrow x_1, x_1 \rightarrow y_1, y_1 \rightarrow z_1, w \rightarrow x_2, x_2 \rightarrow y_2, y_2 \rightarrow z_2$ }, as described in Figure 2.4.3-a. In this context, with respect to the starting node *w* we have the following paths:

$\sigma_1: w \to x_1$	$\sigma_4: w \to x_2$
$\sigma_2: w \to x_1 \to y_1$	$\sigma_5: w \to x_2 \to y_2$
$\sigma_3: w \to x_1 \to y_1 \to z_1$	$\sigma_6: w \to x_2 \to y_2 \to z_2$

What is important to notice is that, unlike the networks depicted in the previous subsection, here the starting node *w* is directly connected not to one, but to two other nodes. On the one hand, *w* is connected to the x_1 node. From this connection, we have paths σ_1 , σ_2 , and σ_3 , which grounds the endorsement of the attributes x_1 , y_1 , and z_1 , respectively. On the other hand, *w* is connected to the x_2 node. From this connection, we have paths σ_4 , σ_5 , and σ_6 , which supports the endorsement of the attributes x_2 , y_2 , and z_2 respectively. The salient feature of this network is that it "branches" to two different nodes. Based on the above, we can say that the network enables the endorsement of both sets of attributes, namely x_1 , y_1 , z_1 and x_2 , y_2 , z_2 .



Figure 2.4.3-b Basic Multiple Inheritance Network

Figure 2.4.3-b depicts network Γ in which we have the nodes {*Nixon*, *Republican*, *Right.Wing*, *Conservative*, *Quaker*, *Pacifist*} and the relations {*Nixon* \rightarrow *Republican*, *Republican* \rightarrow *Right.Wing*, *Right.Wing* \rightarrow *Conservative*, *Nixon* \rightarrow *Quaker*, *Quaker* \rightarrow *Pacifist*}. Thus, we have the following paths:

> $\sigma_1: Nixon \rightarrow Republican$ $\sigma_2: Nixon \rightarrow Republican \rightarrow Right. Wing$ $\sigma_3: Nixon \rightarrow Republican \rightarrow Right. Wing \rightarrow Conservative$ $\sigma_4: Nixon \rightarrow Quaker$ $\sigma_5: Nixon \rightarrow Quaker \rightarrow Pacifist$

Like the previous network, here the starting *Nixon* node has a direct connection with not one, but two different nodes. Thus, we say that, with respect to the *Nixon* node, the network "branches".

55

On the one hand, the *Nixon* node goes through the *Republican* node. From this connection, we have paths σ_1 , σ_2 , and σ_3 , which supports the endorsement of the attributes *Republican*, *Right.Wing* and *Conservative*, respectively. On the other hand, the *Nixon* node goes through the *Quaker* node. From this connection, we have paths σ_4 and σ_5 , which supports the endorsement of the attributes *Quaker* and *Pacifist* respectively.

Based on both branches, the *Nixon* node inherits *both* the *Republican*, *Right.Wing* and *Conservative* attributes as well as the *Quaker* and *Pacifist* properties. The starting *Nixon* node is connected not only to one but two different outgoing nodes, namely, to the *Republican* node and the *Quaker* node. Thus, either outgoing node by itself, i.e. either the *Republican* or the *Quaker* node will render us a single inheritance structure, but the fact that the *Nixon* node is connected to both nodes is what gives us a case of multiple inheritance.

Multiple inheritance provides us with a more complex structural arrangement that significantly varies in expressive power. On the one hand, we have single inheritance, which operates on the basis that nodes are linearly connected through a series of *isa* or *not-isa* link, but any given node is connected to one and only one other node. On the other hand, we have the case of multiple inheritance, which basically is an extension of single inheritance, such that nodes can be connected to more than one other *outgoing* node within the structure. The difference lies in the number of outgoing nodes to which any given node can be connected. Both types of inheritance structures share the underlying motivation of representing knowledge through the elements of properties, classes or attributes that can be shared among individuals, through the *isa* and *not-isa* type of links. Multiple inheritance allows a given attribute or property to be shared among several preceding nodes efficiently. Single inheritance networks can be seen as a restricted class of the more general multiple inheritance networks. Whenever we refer to networks or inheritance networks throughout our research, we will be referring to networks involving multiple inheritance.

The above, concludes the exposition of multiple inheritance networks and a preliminary overview of how we can go about extracting information from these networks.

2.5 Chapter Summary

In this chapter, we examined the basic elements of inheritance networks. In particular, we reviewed

the different components such as nodes, links and paths. We looked into the main mechanism that allows extracting information, i.e. inheritance of properties. Furthermore, we stressed how the inheritance networks research programme can be seen as an adaptation of Quillian's *semantic networks*, that aims to model and aid the understanding of defeasible reasoning. Despite the above, our depiction of inheritance networks restricted itself to the various components and their inner workings. We have not engaged in a discussion of the extent and limitation of these representational structures in the context of defeasible reasoning. That is, we have not discussed the caveats, adaptations and assumptions need to be made to understand these kinds of networks in the context of defeasible reasoning. The aforesaid exceeds the objectives of this chapter.

In the next chapter, we undertake the discussion of the strength and weaknesses of inheritance networks. Furthermore, we provide a delimitation to the type of networks we employ as well as to how such limitation shapes and fall in line with our research.

Chapter 3 Features and Delimitation

In this chapter, we reexamine several of the concepts from the previous chapter. Given that we have a better understanding of inheritance networks and its components, we devote this chapter to discuss the extent and limitation of some of their most important elements. Furthermore, we examine how all of the above fits in the context of our research. This chapter aims to examine the theoretical foundations upon which the remainder of our research develops and the adjacent fields of study surrounding our work. We aim to assess the extent and limitations of inheritance networks as representational structures, discuss some of its shortcomings and the role, if any, such shortcomings have in our overall research.

3.1 Weaknesses

In this section, we discuss some of the limitations and criticism directed towards inheritance networks. The criticisms, along with some of the inherent limitations are, in a sense dispersed, but taken as a whole, follow a recurring theme. The recurring theme refers to the limited nature of inheritance networks. This section aims to revisit some of the most important concepts of inheritance networks, as well as some of the main assumptions, and assess how they fit the more general scheme of our research.

3.1.1 Expressiveness

In this section, we assess various aspects associated with inheritance networks, which are of theoretical relevance, yet somewhat escaped the scope of the previous chapter. These topics converge on the very limited nature of these representational structures. We describe such limitations, and approach them as something we are going to use to our advantage.

First, we briefly address the topic of connectives. As Gabbay & Schlechta (2016) note, in the context of inheritance networks, there are no connectives. That is, conjunction, disjunction, etc. do not play a role within inheritance networks. Furthermore, all relations between nodes are restricted

to the positive or negative defeasible links. This is one of the first limitations of inheritance networks. In line with the lack of logical connectives, Horty et al. (1990) state:

[...] the process of drawing conclusions from a set of defeasible hypotheses through inheritance reasoning is quite different from the process of drawing conclusions from through deduction. Inheritance reasoning doesn't depend on the interplay of connectives, for example, since there aren't really any connectives, to speak of, in our semantic nets (the symbols \rightarrow and \rightarrow are not connectives, since they apply to individuals and kind terms rather than sentences, and they do not nest); [...]. (314)

This relates to the way inferences are drawn within these networks. These authors remark that, strictly speaking, neither \rightarrow nor \rightarrow are to be considered as connectives. For example, both *isa* and *not-isa* links cannot nest, and they also apply to individuals and properties, rather than whole statements.

Second, inheritance networks, as we have discussed so far, are restricted to encoding 1-place predicates. As such, any binary or *n*-ary relation cannot be encoded into inheritance networks. There have been attempts to account for *n*-place relations, but they have been limited to the binary case. As Carpenter & Thomason explain:

Our own work on relations (much of it unpublished) suggests that relational reasoning leads to computational complexity, and that it is not trivial to find good compromises in the relational case between tractability and expressive power. We are looking for systems that require more computational complexity than in the simple nonrelational cases, but that are still in the low polynomial time-complexity range. (1990, 338)

Such restriction imposed on esearch on inheritance networks is connected to the computational complexity associated with these structures when more than the simple 1-place predicates theories are taken into consideration. The aforesaid authors state that:

Most research in inheritance theory has ignored relations altogether, and published works that do consider relations, such as [Touretzky86]' only consider the binary case. This contrasts with logics in general and in particular with nonmonotonic logics, which at the outset consider relations taking arbitrary numbers of arguments. (1990, 338)

Thus, a compromise between tractability and expressiveness is achieved by only considering the monadic case. Moreover, Carpenter & Thomason highlighted that some of the complexity intricacies associated with *n*-ary relations arise not exclusively from their non-monotonic paradigmatic trait, as the more straightforward monotonic case offers a set of challenges of its own $(1990)^{20}$.

In line with the above, we also need to emphasise another important restriction of this class of representational structures. Inheritance networks do not have an explicit notation to represent quantifiers. Thus, strictly speaking, inheritance networks do not include any form of quantification. To understand this, we need to recall the connection among nodes is defeasible. That is, a claim stating the link between any two given nodes stands to account situations in which the association fails. For example, the expression *Quaker* \rightarrow *Pacifist* within a given network standing for the *Quaker* attribute positively associated with the *Pacifist* property does not entail that all Quakers have the attribute of being Pacifists. We could very well imagine a given individual Nixon having the attribute of being a *Quaker* and a direct negative connection with the property of being a *Pacifist*.²¹ The universality feature is what would follow from first-order logic, but does not hold for non-monotonic inheritance networks, as Horty et al. (1990) explain:

There is nothing in ordinary logic very close in meaning to generic statements like these. In particular, "Birds fly" cannot be interpreted through a universally quantified formula of the form $Vx[Px \supset Qx]$, and "Mammals don't fly" does not mean anything like $Vx[Rx \supset \neg Qx]$, since

²⁰ Theoretical problems are often compounded by nonmonotonicity, but are not due to nonmonotonicity alone. Even in the monotonic case, we have encountered subtle research problems in developing the theory of relations. (Carpenter & Thomason, 1990, 340)

²¹ The specific structural arrangement of how this can be achieved need not worry us here, and we will not go into further detail about this example, as it will be addressed in the following chapter in full detail.

the generic statements can be true even in the face of exceptions to the universal claims--a bird that cannot fly, for example. We describe a pair of assertions having the form $x \rightarrow y$ and $x \not\rightarrow y$ as *conflicting assertions*. (313)

The above only holds for non-monotonic inheritance networks, i.e. networks that can account for defeasible connections among its nodes, which as it stands is what we have described so far. Nevertheless, there are classes of strictly monotonic inheritance networks which can be reduced to first-order logic (although this has also been contested, e.g. Thomason et al. (1987). This is the topic of the last subsection within this chapter. In the next subsection, we consider the problem of semantics for inheritance networks.

3.1.2 Semantics

Artificial Intelligence was one of the primary fields to use inheritance networks as representational structures. These networks were seen as an appropriate mechanism to process and operate on taxonomic information. Nevertheless, work on the semantic dimension of this family of networks did not keep up to speed with progress in other areas, as Horty states:

Inheritance systems were originally developed within artificial intelligence in response to the practical need for an efficient way of representing and accessing taxonomic information. These systems, along with network representation more generally, were first presented without any semantic analysis at all, or else only with a procedural semantics, according to which the meanings of the representations was supposed to be specified implicitly by the programs operating on them. (1994, 111)

This lead work on semantics for inheritance networks to branch into two different camps: the *translational* or *indirect* approach and the *non-translational* or *direct* approach.

Throughout the initial research on inheritance networks, the view was held that they were simply a variation on how information could be depicted. In such a vein, these networks could be regarded

as nothing more than an alternative to standard first-order logic. Following this account, inheritance networks are just a notional variation to already established logical languages. Thus, the problem of the semantics for inheritance networks is reduced to classical logic, and as such, the semantics would be defined by the semantics of first-order logic. This approach to the semantics of inheritance networks is said to be a *translational* or *indirect* view (Horty, 1994)²². The *translational* approach, although contested in regard to its feasibility, not only serves as a bridge between inheritance networks. Nevertheless, this approach finds its first challenges when inheritance networks are used for defeasible reasoning (Horty, 1994)²³. The translational view extends the approach to inheritance networks that makes a defeasible interpretation of both *isa* and *not-isa* links by means of a mapping to another nonmonotonic logic formalism. Altogether, this translational approach not only serves as a way to understand inheritance networks, but it also provides the semantics to whichever logical formalisms it is translated, as noted by Horty:

Because their [nonmontonic networks] informal representation required a nonmonotonic consequence relation, it was plain that these representational formalisms could not naturally be translated into classical logic. Nevertheless, it seemed to many that the indirect approach could be extended also to this case by translating the networks into one or another of the non monotonic logics. In analogy with the earlier work of Hayes and others, this research was taken to support the conclusion that defeasible inheritance networks could be viewed as syntactically restricted default theories. (1994, 112)

Moving away from one or various tenets contained in the translational approach towards inheritance networks, we have a *non-translational* approach (Horty, 1994). The main assumption

²² Because of these mappings, it was generally concluded that the networks could be regarded simply as notational variants of syntactically restricted first order theories —distinguished, perhaps [...] only by their attractive appearance on the printed page. [...] [an indirect strategy] ... specifying the meaning of a network formalism through a mapping into an ordinary logical language, usually classical first order logic. (Horty, 1994)

²³ With the attempt to incorporate defeasible information into inheritance hierarchies—in systems such as FRL [Roberts and Goldstein, 1977], KRL [Bobrow and Winograd, 1977], and NETL [Fahlman, 1979]—questions concerning the precise meaning of these networks representations arose once again, [...] (Horty, 1994).
of the non-translational approach to inheritance networks is the idea that these networks are not just a variant of some other logical formalism, but rather a formal framework on their own. This approach is explained by Horty as follows:

At approximately the same time however, a very different kind of technique for analysing the meaning of defeasible inheritance networks was being developed—initially, in a 1984 dissertation by Touretzky, published two years laters [...]. The point of providing a semantic theory for some representational formalism is to allow us delineate the consequences of a set of facts expressed in that formalisms, to explore the characteristics of these consequence sets, and to test the original facts for properties such as consistency. One way to do this, of course, is by mapping the representational formalism into some logical language for which ideas like consequence and consistency have already been defined. But as Touretzky noticed, a theory that achieves the same ends can also be developed, at least in the case of inheritance networks, entirely in terms of the network language itself, without going through the intermediate step of translation into a separate logic. (1994, 112)

In such a vein, inheritance networks stand to have their own set of semantic features without the need for a mapping to another formalism. One corollary of this view is the need to develop a self-standing semantical account, i.e. not a mapping to an already known logical formalism for which the semantic is already established.

The criticism raised against non-translational or direct approaches lies on whether they are carrying out the task of providing a model-theoretic account of inheritance networks or whether they are providing a procedural approach. Some critics of direct approaches state that the subject matter is ambiguous with respect to the declarative/procedural distinction (Horty, 1994). The point of contention between both accounts is the non-translational approaches reliance on the paths of the specific networks which displaces the problem of providing a semantics to the already well-known notion of a path within a network, as noted by Horty:

For theories of this kind [non-translational], the paths through a network often form the main focus of attention (because of this, direct theories

are sometimes described as *path-based*). These paths, which are analogous to proofs in ordinary logic, are supposed to represent arguments or inference procedures; and they do tend to correspond in a loose way to the procedures actually carried out in implications of inheritance systems. To that extent, then, the approach is explicitly procedural. The [procedural] account of meaning provided in these direct theories for a statement belonging to the network language tells us what can be derived from a network containing that statement, as well as the conditions under which that statement itself can be derived from some network. (1994, 113)

This aspect restricts the dependence on a syntactic construction rather than on some defined notion of truth. As Horty states "The treatment of derivability is purely syntactic (or proof-theoretic); it does not rely on a prior notion of truth in a structure for items belonging to the network language. In fact, there is no attempt at all to define truth conditions for defeasible links occurring in inheritance networks [...]" (1994, 113).

The above produces a twofold problem. On the one hand, *translational* approaches seem to operate on a simple and clear-cut theoretical presumption which benefits inheritance networks inasmuch as they can carry any well-known notions to the target logic in the mapping process. Nevertheless, this can also be seen as its drawback inasmuch as these networks seem to offer no real self-standing advantage other than their typographical appeal. Thus, the translational approach can be seen as trivialising inheritance networks. On the other hand, *non-translational* approaches, although a very attractive account of inheritance networks as self-standing and proper formalisms, they seem to rely on a procedural semantics. Unlike a declarative semantic approach, the direct approach seems not to comply with various model-theoretic notions to which well-known logical formalisms do comply, such as having some basic notion of truth. Nevertheless, this might be a misguided criticism which need not apply to inheritance networks as representational structures, as alluded by Horty:

Another way in which the direct theories differ from the standard declarative paradigm is that they tend not to rely on model-theoretic notions, which were taken by some researchers [...] to form the core of any rigorous semantics for a representational formalism. [...] the direct

approach has sometimes been criticized for this [no attempt at all to define truth conditions for defeasible links] reason. This kind of criticism seems to be misguided; or at least, it relies on standards very different from those at work in other areas of logic, where proof theory and model theory stand on a more even footing. (1994, 113)

In particular, it has been pointed out that inheritance networks have a "dual conceptual role" in the following sense:

Semantic nets play a dual conceptual role: we can think of them either as models or as theories (i.e., structured sets of assumptions, together with procedures for inferring conclusions from these assumptions). This may seem surprising to those who are used to the crisp division in traditional logic between semantic and proof-theoretic ideas. Semantic nets, however, come from a different tradition, in which the distinction is not so clear: it is harder to draw the line, in knowledge representation, between what is represented and what is doing the representing. (Thomason et al. 1987, 7)

Following this thread, inheritance networks can be seen either as models on the one hand or as theories on the other hand. As such, a clear distinction between "what is represented and what is doing the representing" is not straightforward (Thomason et al. 1987).

This concludes our assessment of the problem of a formal semantics for inheritance networks. The aforesaid problems render inheritance networks as a somewhat limited tool. Like we stated previously, although an interesting problem by itself, this is not the focus of our research nor the problem we wish to address. In the next section, we revisit such shortcomings and explain how we are going to use it to our advantage, in the context of our research.

3.2 Strengths

In this section, and in direct opposition to what we did previously, we assess the features of inheritance networks that render them as an attractive tool for defeasible reasoning. Moreover, we

point out how the traits that frame inheritance networks as a limited tool are going to work to our advantage. This section aims to discuss the impact of the limited expressive power of the networks of our choosing has for our research, and the problem we set ourselves to examine. That is, we aim to assess how such limitations fit in our research.

3.2.1 Attractive Features

In this subsection, we discuss some of the attractive features inherent to the class of networks we are using throughout our research. We focus on some of the unique traits that make inheritance networks an attractive tool for the study of defeasible reasoning.

First, and as a byproduct of the simplicity and expressive limitations of inheritance networks, these networks are *efficient*. In particular, inheritance networks are more tractable and efficient to operate with and about compared to other logical formalisms. As Carpenter and Thomason remark:

The primary reason for employing inheritance mechanisms is that it is usually more computationally tractable, in terms of speed and efficiency, to perform reasoning using inheritance algorithms than to directly implement some logical theory expressed in terms of axioms and rules of inference. The reason for this is that efficient graph theoretic algorithms can be defined for inheritance. This is in sharp contrast with the case of first-order logic, which is not even decidable. (1990, 312)

Furthermore, since inheritance networks are theoretically equivalent to directed acyclic graphs, they benefit from all the already well-known algorithms and procedures that compute graphs. Moreover, this simplicity would render these representational structures, easier to implement, as Brewka explains:

Since the expressiveness of the language of and IHS [inheritance system] is usually very restricted it is hoped that quite efficient, yet theoretically well-founded, implementations can be found. Some existing implementations support this expectation. (1991, 108)

Second, inheritance networks are attractive to the extent that they offer greater *modularity* than other more standard logical formalisms. This modularity feature of inheritance network is a byproduct of these networks following the graph-oriented theme for knowledge representation. That is, unlike several standard logic formalisms, modifications and revision to a given point within a network ought to be more straightforward inasmuch as it would only require altering the specific section with the graph rather than a revision to the theory as a whole. Carpenter & Thomason discuss this aspect as follows:

A network localizes the information it encodes in a way that first-order theories do not. That is, a distance metric can be defined between nodes in a network by the number of links in the paths connecting them. The idea is then that a network will be straightforward to modify, since revisions "about" a concept will only require the modification of links close to it and not links which are arbitrarily distant. (1990, 313)

The above stands in opposition to what happens in much of the standard logic in the sense that information is not localised at some given point within a theory, but rather it follows from the theory together with the axioms (Carpenter & Thomason, 1990)²⁴.

Third, and associated with the previous feature, inheritance networks are attractive to the extent that they are more intuitive and easier to operate about and with due to their graph-oriented nature. This means that a given system based on such representational framework would entail as an easier entry point, which in turn renders it easier to build and maintain. In this sense, inheritance networks are said to be an attractive tool due to their *buildability* and *maintainability*²⁵. One key aspect

²⁴ "Networks not only store explicit knowledge as patterns of connected nodes, but store it in the expectation that reasoning procedures will be local. Logic, on the other hand, doesn't lend itself to such modularity. In general, it makes no sense to ask "where" a natural consequence of an axiomatized theory is stored in the axioms." (Carpenter & Thomason, 1990, 313).

²⁵ "For a knowledge representation scheme to be technologically useful, rather than, say, a research tool, it should be possible for more or less ordinary people without a large amount of specialized training to maintain it. Networks have a kind of intuitive simplicity that makes them perform well in this respect. And the iconic nature of nets allows many of the important properties to be easily gleaned by visual inspection, thus allowing relatively untrained people to work with them [...]" (Carpenter & Thomason, 1990, 313).

associated to the buildability and maintainability lies on the fact that a given system is always subject to addition or retraction of information by its designers and users, due to changes occurring in the domain of knowledge for which it is being applied. The modularity feature makes it easier to build and maintain larger systems without having to undergo major revisions to the system as a whole²⁶. The above can be regarded as a minor and secondary trait that ought not to play a role in the context of theoretical concerns. Nevertheless, it has been noted that these traits are important to the extent that systems based on such networks are to be designed, used and maintained by experts in a domain of knowledge different from the system itself (Carpenter & Thomason, 1990)²⁷.

Fourth, and despite the limitations already discussed, inheritance networks do offer a satisfactory expressive capability to the extent that they can portray in an intuitive and simple fashion what is otherwise less straightforward by means of a more robust logical formalism (Horty, 1994). This makes inheritance networks a tool that better serves the purpose of studying defeasible reasoning. Conversely, this straightforward and easy access entry to core problems in defeasible reasoning paves the way to a better understanding and insight into the field. This is explained by Horty as follows:

[...] the explicit representation of arguments in these direct theories allows for a very fine-grained analysis of the structure of defeasible reasoning, which it is often difficult to achieve using more general nonmonotonic logics. For this reason [the explicit representation of arguments] also, because of their sensitivity to the detailed structure of argument, the techniques of path-based inheritance allow a good deal of

²⁶ "Maintenance is as important as buildability. No knowledge based system is ever complete. As more information is gained about a particular domain, or errors are discovered in the original data, it is necessary to modify the original knowledge base. It is important to develop knowledge representation systems that can readily be updated; among other things, this means that changes that seem intuitively to be local can be made by local modifications of the knowledge system. Without this feature, the system is especially liable to lose its integrity as it becomes larger. (Carpenter & Thomason, 1990, 314)

²⁷ It would be desirable as well to have a system that could be maintained as well as built primarily by users or domain experts, rather than requiring a computer scientist or knowledge engineer. Networks do not provide perfect solutions to these problems, but do deliver techniques that work acceptably in many practical situations, as is witnessed by their widespread application in commercial expert systems." (Carpenter & Thomason, 1990, 314).

versatility in the definition of particular theories, enabling us to articulate a variety of different intuitions about defeasible reasoning. At times, the contrast among the different theories can suggest, in a very simple environment, both options and problems that may not have been apparent form a more general point of view. (1994, 116)

It is clear now that a representational framework should aim to depict and express as much as it can be said of a given domain of discourse through enhanced expressive capabilities. Nevertheless, this is often achieved only at the expense of some of the various features that render inheritance networks as an attractive formalism. In such vein, the tradeoff between expressibility and other features such as efficiency, tractability and ease of use is inescapable:

Of course, we want a representation scheme that will allow us to encode all of the relevant knowledge for our application domain. But as usual in computational implementations, there is a tradeoff between expressive power and computational efficiency. Adding expressive power will usually result in a less computationally tractable theory. That is, the power is inversely related to the speed of the procedures that may be designed to manipulate it. (Carpenter & Thomason, 1990, 315)

In this sense, inheritance networks face the same challenges as any other logic formalism face in terms of serving as a knowledge representation tool. This important aspect is remarked by Thomason & Horty as follows:

It is important to realize that inheritance theory, which deals with objects like graphs and paths rather than with formulas and proofs, is an independent theoretical pursuit with its own methods, intuitions, and results. In some ways it is like logic. But it is not the same as logic. The emergence of inheritance theory as a separate area of inquiry is an important step in relating theories to actual knowledge representation technology. (1989, 221).

In such a vein, inheritance networks trade some of the expressive capabilities located in other logic formalisms for a series of traits that seem to render these networks as an attractive and legitimate

alternative (Carpenter & Thomason, 1990)²⁸.

This subsection summarises some of the most important features associated with inheritance networks, which render them as a desirable tool for studying defeasible reasoning. As we have stated, despite the numerous limitations, there are ample motives for the adoption of these network-like structures as a serious alternative to other non-monotonic logics.

3.2.2 Attractive Shortcomings

In this section, we examine how some of the inherent shortcomings of inheritance networks turn into positives features *for the purposes of our research*. We discuss how several limitations can be used to strengthen the main point of contention underlying our overall work.

First, and as we have previously stated, inheritance networks can be regarded as an extremely simple and limited language to the point that some of the most basic logical principles cannot be expressed within the language in the absence of connectives:

The language with which we have been dealing is so weak that it is incapable of representing many of the principles that typically distinguish classical from nonclassical logics. In particular, *Excluded Middle* can't be expressed, since disjunction is not available. (Thomason et al. 13)

The above has been referred to as an *impoverished language*, but it has also been recognised as an important framework *despite* some of its salient limitations. This specific aspect is highlighted by Thomason et al. as follows:

To a logician, the [simple and restricted] context will seem unusually simple (there are, for instance, no genuine propositional connectives);

²⁸ "In our approach to inheritance theory, expressive power is given lower priority than efficiency. We ensure a computationally tractable knowledge representation system by relentlessly maintaining efficiency, while gradually enlarging the system to achieve more expressive power." (Carpenter & Thomason, 1990, 315).

but the texture of the resulting logic is surprisingly rich, given the impoverishment of the background language. (1987, 1)

In a similar vein, Carpenter and Thomason explain that:

When properly understood, networks are formalisms for understanding natural patterns of tractable reasoning that are as intelligible and theoretically respectable in their own right as familiar logical formalisms. The calculi of symbolic logic were designed to account for mathematical reasoning, and as a result are expressively powerful, intractable, and rather distant from commonsense language and thought patterns. Networks are expressively weak, tractable, and tend to be closer to commonsense language and thought patterns. (1990, 311)

In this sense, inheritance networks offer a sensible aid in the task of studying and understanding various forms of defeasible reasoning which render these structures as a legitimate framework "in their own right" (Carpenter & Thomason, 1990). Nevertheless, as we already mentioned, research in the field has explicitly pointed out that such "weak expressive power" is not to be equated to "theoretical triviality". That is, even in these very simple inheritance structures, the theoretical problems are conceptually intricate and complex enough to pose a challenge (Gabbay, & Schlechta, 2016)²⁹. Inheritance networks do offer an instance in which simplicity does not lead to any form of theoretical insubstantiality as stated by Carpenter and Thomason:

We have stressed all along the weak expressive power of networks, and that this power is sacrificed to obtain computational tractability. But computational tractability and expressive weakness should not be conflated with theoretical triviality. Even though networks can be thought of as impoverished logics, the definition of inheritance and the proofs of basic theorems are often complex. (1990, 340)

²⁹ "Inheritance diagrams are deceptively simple. Their conceptually complicated nature is seen by, e.g. the fundamental difference between direct links and valid paths, and the multitude of existing formalisms, upward versus downward chaining, intersection of extensions versus direct scepticism, on-path versus off-path preclusion (or preemption), split validity versus total validity preclusion etc., to name a few, see the discussion in Sect. 3.5." (Gabbay, & Schlechta, 2016, 75)

Having stated the above, the question revolves around whether such array of limitations might be of any positive use for the purposes of our research. In particular, we have two important claims towards a positive answer to the aforesaid. First, the limited expressive power of these formal structures allows us to focus on the core problem underlying our research without getting distracted by any form of technical intricacy. This, in turn, has a two-fold grounding. On the one hand, it serves us to the extent that our inquiry into the problem is more expedite and focused. On the other hand, and once having identified the problem in the most straightforward way, we can draw attention to the most relevant features of the problem as well as our approach to address the problem without getting distracted by extraneous technical intricacy. Second, and moving past beyond the clarity and simplicity previously remarked, there lies a deeper theoretical claim in favour of such a restricted class of representational structures. In particular, if the problem under scrutiny can be instantiated in such a limited and restricted framework, the problem itself is not a byproduct of an overly elaborate representational framework, but rather a deeply rooted difficulty which can be instantiated using some of the most limited tools at our disposal. Either aspect in support of such a restricted framework suffices as a methodological justification. Nevertheless, the conjunction of both features provides a strong case not only of the adequacy, but rather the adequacy of inheritance networks as a tool for the purposes of our research.

In this subsection, we acknowledged some of the various limitations associated with inheritance networks. Nevertheless, we have not fully developed how such limitations serve the purpose of enhancing our approach to the particular problem of study. The aforesaid is explained in thorough detail over the remainder of our thesis.

3.3 Delimitation

In this last section, we examine the type of inheritance networks we are using to conduct our research. In particular, we go over some of the aspects that can extend or limit inheritance networks as representational structures and define the type of traits and networks we are using. Moreover, we specify the type of problems we are leaving out of our discussion that allow us to focus and narrow our attention into the type of problems we are undertaking. The aim of this section is to provide a clear demarcation of two vital aspects of our work: the type of networks that we employ and the type of problems that although interesting, we are leaving behind. This section also aims to critically asses how the aforesaid shapes our research.

3.3.1 Type of Networks

As we have extensively pointed out, defeasible reasoning is an important field which has taken inheritance networks as a research tool. We also stressed the imprint placed by non-monotonic reasoning to these network-like formalisms, and how this drew a departure from Quillian's semantic networks research programme. Thus, the *raison d'etre* of inheritance networks as a serious tool for understanding and modelling defeasible reasoning lies upon the positive and negative links among the nodes within a network having a defeasible interpretation. This stands in direct opposition to statements that hold without exception, such as definitions (Carpenter & Thomason, 1990). Given that inheritance networks are primarily a tool for defeasible reasoning and nonmonotonic logics, the defeasible account of *isa* and *not-isa* links is usually the default one (Gabbay & Schlechta, 2010). This defeasible reading of links within networks is what usually renders these class of networks as defeasible networks.

Nevertheless, the defeasible interpretation of *isa* and *not-isa* links as defeasible associations is not the only form of reading such connections. Another approach relies on interpreting *isa* and *not-isa* links as strict and non-defeasible connections. Under this approach, the direct association between any two given nodes cannot include any form of retraction or exceptions. This view only allows connections that have an infallible nature (Carpenter & Thomason, 1990). Such reading of *isa* and *not-isa* links is less attractive for knowledge representation and artificial intelligence in general. Thus, this interpretation is less explored and discussed throughout the literature.

In line with the above, and although the strict approach is less interesting or useful, a different approach which accounts for this is more common. In particular, there has been research into what is known as networks of *mixed inheritance*. Such frameworks account for both types of readings of *isa* and *not-isa* links, employing two different relation symbols. First, we can express the positive *isa* defeasible association between any two given nodes x and y as follows: $x \rightarrow y$. Moreover, to express the positive *isa* strict association between any two given nodes x and y as and y as follows: $x \rightarrow y$. Likewise, we can express the negative *isa* defeasible association between any two given nodes x and y as $x \neq y$, and to express the negative *isa* strict association between any two given nodes x and y as $x \neq y$.

Following the above convention, an inheritance network can mix both types of positive and

negative associations between nodes. This type of network can be said to have a far more expressive power as it allows us to represent a larger class of statements. The above being said, and *for the purposes of the problem under scrutiny*, we need not the greater expressive power of mixed inheritance networks. That is, the problem that lies at the heart of our research can be instantiated using only defeasible links.

In line with the previous delimitation, there is also an important point to be made. As we noted earlier in this chapter, rather than a single and narrow formalism, inheritance networks are more like a family of network-like representational structures. In such a vein, we do not address all the various network-like formalisms, but rather focus on a very simple type of inheritance network.

On the one hand, assessing an ample class of networks would undermine the key objectives of our research. In particular, this research is not about inheritance networks, nor any specific non-monotonic formalism. This research is about an epistemological problem of defeasible reasoning, which can be instantiated in the most simple, elementary and straightforward formalisms. Thus, according to our stance, the network-like formalism of our choosing is rather a gateway to introduce an epistemological problem. As such, the formalism is simply a way in which we can inquire into the main concern of our research, not the other way around.

On the other hand, although the various formalisms are part of the same family of network-like structures, they do have subtle differences that would render the endeavour of talking uniformly about the same problems and their inner workings counterproductive. This is not to say that they exclude one another, but rather that some of their differences do not translate directly into one another. Again, like the previous point, our research is not about inheritance networks, but rather a specific problem within defeasible reasoning, which so happens that a specific type of inheritance network serves as the simplest way to instantiate such a problem. Thus, we deliberately determined to delimit the type of network-like formalisms that we employ throughout our research.

This concludes our justification for delimitation of the type of inheritance networks we are considering throughout our research. This is not to say that other approaches or types of networks-like structures are not viable options. Nevertheless, our delimitation has been deliberately chosen to strike a compromise between the simplest type of networks in which we can instantiate the problem that lies at the heart of our research without sacrificing any form of generality.

3.3.2 Type of Conflicts

In the first section of the previous chapter, we emphasised that one of the central concerns for defeasible reasoning is the problem presented by conflicts. In such vein, and following the picture portrayed by Strasser & Antonelli (2018), conflicts have a multifaceted nature. More precisely, these authors make a twofold division of conflicts within defeasible reasoning. On the one hand, resolvable conflicts, which, as their name suggests, have procedures to address the nature of the conflict within the bounds of a formal framework. On the other hand, non-resolvable conflicts, for which a solution needs to call upon to extra-systemic criteria. We briefly consider some of the *resolvable* types of problems within inheritance networks, as they will be short-lived in the context of our research.

To understand resolvable conflicts, we need to recall that inheritance networks can contain multiple inheritance. In such cases, any given node can be associated with more than one other node. This produces the problem of redundancy.



Figure 3.3.2-a Basic Redundancy

Figure 3.3.2-a depicts the network Γ in which we have the nodes {*Tweety*, *Penguin*, *Bird*} and the relations {*Tweety* \rightarrow *Penguin*, *Penguin* \rightarrow *Bird*, *Tweety* \rightarrow *Bird*}. Thus, we have the following paths:

 σ_1 : Tweety \rightarrow Penguin σ_2 : Tweety \rightarrow Penguin \rightarrow Bird σ_3 : Tweety \rightarrow Bird

In this example, the starting *Tweety* node has a direct connection with not one, but two different nodes, which renders the network as a multiple inheritance structure. On the one hand, the *Tweety* node inherits the attributes from the *Bird* property, and it does so through the implicit connection it has to such class through the *Penguin* node. That is, the *Tweety* node inherits the *Bird* attribute through σ_2 . On the other hand, we also have that the *Tweety* node inherits the attributes from the *Bird* property through the direct connection it has with such node, i.e. through σ_3 . Thus, the direct link of the *Tweety* node to the *Bird* node creates a redundant situation. That is, there is an implicit and explicit connection between both nodes at the same time. Now, as far as all the links are positive, there seems no reason for such phenomena to have unwanted side effects. Nevertheless, when negative information is taken into account, redundancy can have undesirable effects.



Figure 3.3.2-b Basic Redundancy with Negative Links

Figure 3.3.2-b depicts network Γ in which we have the nodes {*Tweety*, *Penguin*, *Bird*, *Flying*.*Thing*} and the relations {*Tweety* \rightarrow *Penguin*, *Penguin* \rightarrow *Bird*, *Tweety* \rightarrow *Bird*, *Penguin* \rightarrow *Flying*.*Thing*}. Thus, we have the following paths:

 σ_1 : Tweety \rightarrow Penguin σ_2 : Tweety \rightarrow Penguin \rightarrow Bird σ_3 : Tweety \rightarrow Penguin \rightarrow Bird \rightarrow Flying.Thing σ_4 : Tweety \rightarrow Penguin \rightarrow Flying.Thing

The *Tweety* node is implicitly linked to the *Flying*. *Thing* attribute via its positive connection through the *Penguin* and *Bird* nodes, i.e. σ_3 . Nevertheless, the *Tweety* node is at the same time, implicitly connected by a negative link to the same *Flying*. *Thing* attribute, in this case through the *Penguin* node and from there by the explicit negative link to the *Flying*. *Thing* attribute, i.e. σ_4 . In this case, the question revolves around as to whether the *Tweety* node should be given the *Flying*. *Thing* attribute or not, since it seems possible to draw a connection to this attribute as well as to its negation.

One possible solution to deal with the aforementioned lies in taking into account the distance between the various nodes involved, as a criterion to solve such conflicting situations arriving from redundant links. The criterion is referred to as the *inferential distance ordering*. The solution is based on the idea that in cases such as the one depicted by Γ , the paths involving shorter or more direct associations, i.e. paths with a smaller number of intermediary nodes should take primacy over paths with conflicting information, but involving a longer route, i.e. requiring a greater number of intermediary nodes. Given the information provided by the inheritance network Γ , one notices that the *Tweety* node is connected to the *Flying*. *Thing* attribute through the following paths:

> σ_3 : Tweety \rightarrow Penguin \rightarrow Bird \rightarrow Flying.Thing σ_4 : Tweety \rightarrow Penguin \rightarrow Flying.Thing

On the one hand, the *Tweety* node is negatively associated with the *Flying*. *Thing* attribute through a path involving one intermediary node, i.e. the intermediary *penguin* node. On the other hand, the *Tweety* node is positively associated with the *Flying*. *Thing* attribute, but this time it does so through a path involving two intermediary nodes, i.e. the *Penguin* node and *Bird* intermediary nodes. Thus, the shorter σ_4 path would trump σ_3 , and *based on the inferential distance ordering*; we would state that the *Tweety* node inherits the negative connection to the *Flying*. *Thing* attribute.

The previously stated structural arrangements produce various problems within inheritance networks. Nevertheless, these types of problems are not the focus of our research. We focus on a

specific type of non-resolvable conflict within inheritance networks. In the same vein, the specific type of network-like formalism, we have chosen allows us to inspect the problem of our interest without needing to undertake problems such as the ones hereby highlighted.

The aforesaid is a crucial methodological decision that enables us to focus on the problem of our choosing with the simplest formal framework. Such decision is grounded on the fact that our research is not about inheritance networks, but rather a specific problem of defeasible reasoning for which inheritance networks offer the most straightforward formalisms to instantiate, study and inquire about. Having stated the above, the connection between inheritance networks and defeasible reasoning, in general, is far deeper, and it would demand a thorough research by itself. Here we limit ourselves to having acknowledged the connection, but sidestep that line of inquiry to address a specific problem. The link between both endeavours is an underlying motivation throughout our work. The nature of this motivation will surface in the final chapters of the thesis when we examine the epistemological aspects of inheritance networks.

This concludes the demarcation of the class of problems we will be addressing throughout our research. Again, this leaves out many problematic aspects surrounding defeasible reasoning, non-monotonic logics and inheritance networks. Nevertheless, this demarcation also lets us be precise about what are we going to be looking for throughout the rest of our work.

3.3.3 Stein Networks

So far, our depiction of inheritance networks and the notions we have built around these networks has made a few tacit assumptions that are yet to be discussed. In this brief subsection, we examine the aforesaid assumptions as well as some their implications. To address the above, we consider a particular network.



Figure 3.3.3 Multiple Inheritance Network

Figure 3.3.3 depicts the network Γ in which we have the nodes {*Nixon*, *Republican*, *Pacifist*, *Conservative*} and the relations {*Nixon* \rightarrow *Republican*, *Republican* \Rightarrow *Pacifist*, *Republican* \rightarrow *Conservative*}. As we have previously stated, this is a structure involving multiple inheritance. The *Nixon* node inherits the *Republican* attribute, which in turn inherits, on the one hand, a positive connection to the *Conservative* attribute and, on the other hand, a negative connection to the *Pacifist* attribute. Thus, we have the following paths:

 $\sigma_1: Nixon \rightarrow Republican$ $\sigma_2: Nixon \rightarrow Republican \not\rightarrow Pacifist$ $\sigma_3: Nixon \rightarrow Republican \rightarrow Conservative$

Such network endorses the *Conservative* attribute to the *Nixon* node, and at the same time, it fails to endorse the *Pacifist* to the same node. Nevertheless, it is important to notice that, following the same network, we can focus on the *Republican* node and examine which properties could be attributed *with respect to* such node. On the one hand, the *Republican* node inherits the positive

link to the *Conservative* attribute. On the other hand, the *Republican* node inherits the negative link to the *Pacifist* attribute. In particular, we have the following:

 σ_4 : Republican \Rightarrow Pacifist σ_5 : Republican \Rightarrow Conservative

Thus, depending on the node we fix as the starting node, we have different paths. The same principle can apply to the terminal node. This feature is very specific to Stein's (1991) depiction of inheritance networks and was referred to by Simonet & Ducournau (1994) as a *subgraph-base* approach as opposed to a *path-based* approach. As Simonet & Ducournau (1994) point out, this very subtle variation is far from how inheritance networks are usually depicted throughout the literature. Nevertheless, the purpose for complying with Stein's approach to constructing inheritance networks relies on the fact that the key problem we examine in our research (*floating conclusions*) can be instantiated in a much more straightforward fashion. The simplicity of addressing the key problem of our research relies on two important reasons.

First, by complying with Stein's approach to inheritance networks, we no longer allow for either cycles within our networks or mixing strict and defeasible forms of *isa* and *not-isa* links. This makes our exposition of inheritance networks much simpler and easier to go through without taking a detour into various technical subtleties within these representational structures. Second, following Stein's approach, we can focus solely on the problem of floating conclusions without taking a detour into other problems for these networks, such as the problem of *preemption*.

The problem of preemption would inevitably lead us to address some of the main responses such as the *inferential distance ordering* and *specificity*, which are important and relevant in their own right, but are not necessary if our objective is to instantiate the problem of floating conclusions. Thus, in the process of adhering to Stein's approach, our depiction of inheritance networks is simpler as well as our understanding of the problem at hand.

The above, although abstract to some extent at this stage is addressed and made explicit in its due course over the next chapters of our investigation. Nevertheless, it suffices at this stage to briefly mention some of the most relevant restrictions and delimitations that shape our work.

3.4 Chapter Summary

In this chapter, we assessed the strengths and weaknesses associated with inheritance networks, such as their limited expressive capability and the overall lack of a well-defined semantics. We discussed how the same features that render inheritance networks as a limited tool, better serve the objectives of our research. Moreover, we delineated the kind of problems we look into, and the specific type of networks we use to assess such problems.

The above being said, we have not mentioned any problematic arrangements within inheritance networks that we analyse. In particular, we have put to the side one interesting and central problem of inheritance networks, namely the problem of *ambiguity*. In the next chapter, we examine situations known as ambiguous networks and assess the proposed solutions to the problem.

Chapter 4 Ambiguity

In the previous chapters, we introduced inheritance networks as a representational framework. We addressed various topics such as nodes, links, paths and the mechanisms employed to extract information. Throughout this chapter, we analyse a crucial problem within inheritance networks, namely the problem of ambiguity. We assess the extent and limitation of ambiguity, but more importantly, we use this problem as a gateway into the two main strategies used to compute information within inheritance networks. The aim of this chapter is twofold. On the one hand, we aim to assess the problem of ambiguity as a core problem within the context of our research. On the other hand, we aim to describe how the problem itself produces two information-processing policies within inheritance networks. These approaches or information-processing mechanisms, along with the problem of ambiguity, are the gateway to the central problem of our research.

4.1 The Problem

In this section, we analyse the problem of ambiguity. This problem serves as a gateway to various information-processing mechanisms within inheritance networks. Although ambiguous networks have a very simple structural arrangement, they present one of the central debates of information processing within inheritance networks. This section aims to address a foundational problem that produces the core problem of our research (addressed in the forthcoming chapter), i.e. *floating conclusions*. At this stage, it suffices to assess the conditions by which the problem emerges.

4.1.1 The Nixon Diamond

In this subsection, we look into a case study that instantiates the ambiguity problem. We examine this problem through the example known as the *Nixon Diamond*. We set the stage of the problem, yet the discussion of its extent and limitation is carried over throughout the rest of the thesis. Furthermore, the problem instantiated by the Nixon Diamond is only *a gateway to the more central problem that occupies our research*, namely, floating conclusions. As such, we do not devote much attention beyond what is strictly necessary to understand the technical context that produces

floating conclusions (chapters five and six).



Figure 4.1.1 Nixon Diamond

Figure 4.1.1 depicts network Γ in which we have the nodes {*Nixon*, *Republican*, *Quaker*, *Pacifist*} and the relations {*Nixon* \rightarrow *Quaker*, *Quaker* \rightarrow *Pacifist*, *Nixon* \rightarrow *Republican*, *Republican* \nleftrightarrow *Pacifist*}. Thus, we have path σ_1 : *Nixon* \rightarrow *Quaker* \rightarrow *Pacifist* such that $\sigma_1 \in \Gamma$, and the path σ_2 : *Nixon* \rightarrow *Republican* \Rightarrow *Pacifist* such that $\sigma_2 \in \Gamma$. We also have the path σ_3 : *Nixon* \rightarrow *Quaker* with $\sigma_3 \in \Gamma$, as well as the path σ_4 : *Nixon* \rightarrow *Republican* with $\sigma_4 \in \Gamma$. In a more schematic fashion, we have the following:

$$\sigma_1: Nixon \rightarrow Quaker \rightarrow Pacifist$$

 $\sigma_2: Nixon \rightarrow Republican \rightarrow Pacifist$
 $\sigma_3: Nixon \rightarrow Quaker$
 $\sigma_4: Nixon \rightarrow Republican$

The aforesaid network is known as the Nixon Diamond. First, it is important to notice that the paths $Nixon \rightarrow Quaker$ and $Nixon \rightarrow Republican$ (i.e. σ_3 and σ_4) do not conflict with one another. The above being said, we have two features that do provide a challenge. On the one hand, we have a path (namely σ_1) starting from the *Nixon* node that goes through the *Quaker* attribute such that it

is positively connected to the *Pacifist* attribute. On the other hand, have a path (namely σ_2) that starts with the *Nixon* node, that goes through the *Republican* attribute, which in turn, is negatively connected to the *Pacifist* property.

The above produces the problematic situation in which we have two equally legitimate paths pointing to conflicting information. The nature of the problem lies in the fact that the conflict occurs from completely unrelated attributes, namely the *Quaker* and *Republican* attributes. Hence, the *Pacifist* property with respect to the *Nixon* node is ambiguous since both the positive and the negative connection can be drawn, eliciting a conflicting situation. Such a state of affairs is what we refer to as an ambiguous network, or simply as an ambiguity. The key question is how to resolve this type of conflict.

Despite the simple structural arrangement that produces ambiguities, the problem presents an important difficulty that sets the stage for a wide and opposing views as of how information within the above type of representational frameworks should be processed. That is, the problem of ambiguity poses one of the simplest, yet fundamental problems within inheritance networks. Different stances toward this problem produce the various approaches to some of the key features of this family of knowledge representation frameworks. How we decide to process the information contained in the pair of conflicting paths is what produces an array of stances towards inheritance networks. As Horty puts it "What you say about inheritance depends crucially on your treatment of nets like the Nixon Diamond" (1990, 317).

One important aspect to take into account concerning the concept of ambiguity lies in the fact that this notion is not equivalent to the common sense of ambiguity. A common every-day use of ambiguity might appeal to some form of vagueness or imprecision. Nevertheless, in the context of inheritance networks, an ambiguity is to be understood as a very specific type of conflict. As we showed in section 3.3.2, there are various ways to produce conflicting scenarios within inheritance networks. More specifically, ambiguity reflects the type of conflict in which a starting node reaches a given terminal node by two different and completely unrelated but conflicting routes, such that, *prima facie*, there is no way to resolve the conflict. The absence of a salient mechanism to decide how to draw information from an inheritance network involving multiple extensions that support opposing conclusions is what typifies the scenario as ambiguous, since there is no clarity as to what is it that we are warranted to infer. In this sense, an ambiguity will be treated as a specific *type of conflict* rather than a conflict *simpliciter*.

Before proceeding to present the various strategies that have been put forward to address ambiguities, in the next subsection we look into a notion that is used in cases in which there are not one, but multiples streams of conflicting information associated to a given network, like in the Nixon Diamond.

4.1.2 Extensions

The existence of multiple paths within a given network raises a problem when some or all of the paths conflict with one another, i.e. in the case they support conflicting information, like the Nixon Diamond. Networks in which the various paths conflict with one another, like in the case of networks involving an ambiguity have no straightforward strategy to determine the information that we can extract. This type of conflicts compel us to revise our concept of a path in a way that such conflicting situations can be accommodated. In this brief subsection, we review the informal and intuitive idea behind the notion of an *extension*.

In the context of a network where the paths associated with the networks are at odds with each other, we need for a mechanism or strategy to decide what is to be drawn from a given network. An *extension* is the notion that determines the various subsets of coherent and consistent information within the network. That is, an extension represents a non-conflicted unit of information within a network.

In a more technical sense, an extension represents a maximally consistent unit of information. That is, an extension refers to a fragment of information associated with a given network, such that no other information of the network can be added to such subset without loss of consistency. The maximality of extensions in the context of inheritance networks and non-monotonic logic is relative to a given structural arrangement and not absolute, as the addition or subtraction of links within a given inheritance network reshapes the extensions associated to it.

One of the main uses of the concept of an extension is that it allows us to identify the specific subset of paths that express the greatest amount of information associated with a network (the relative maximality property) in such a way that dismisses all the paths from the same network that conflict with the information subset in question. In this sense, the idea of an extension carries informational utility and allow us to identify the subsets of paths that can support a given

conclusion to be inferred from a network. To be more precise, extensions stand for subsets of *non-conflicted sets of paths* within a given network.

This account of an extension as an elementary unit of non-conflicted information associated with a given network addresses two core tasks. On the one hand, an extension classifies and distinguishes the conflicting paths associated with the network. That is, an extension guarantees subsets of information that preserve coherence and consistency. On the other hand, an extension captures subsets with as many paths as possible associated with the network, without loss of coherence or consistency. That is, an extension seizes as much information as possible, while preserving coherence and consistency.

In networks where the paths do not clonflict with one another there is only one extension that contains all the paths of the network. In such case, the extension of the network would collapse in a set that contains every single path associated to the network. This is the trivial and less interesting case of an extension. Nevertheless, if there is a path that stands in conflict with at least one other path associated to the network, such conflicting path will be part of a separate "informational unit" of the network, i.e. it instantiates another extension. It follows that a given inheritance network may have multiple extensions, and this is the most common and interesting case for networks.

The above provides a general sketch regarding the notion of an extension. Moreover, we discussed how this concept provides an elementary and basic currency to process information associated with inheritance networks. In the following subsection, we study the role of this concept within networks such as the Nixon Diamond.

4.1.3 Extracting Information

Having presented the concept of ambiguity contained in structures like the Nixon Diamond, and the notion of an extension used to accommodate the problem, we now have the tools to examine the problem of extracting information from inheritance structures involving ambiguities in a more precise fashion.

As we pointed out, the Nixon Diamond is a network Γ composed of the nodes {*Nixon*, *Republican*, *Quaker*, *Pacifist*} and the relations {*Nixon* \rightarrow *Quaker*, *Quaker* \rightarrow *Pacifist*, *Nixon* \rightarrow *Republican*, *Republican* \rightarrow *Pacifist*}. Based on the above, we have the following paths:

 $\sigma_1: Nixon \rightarrow Quaker \rightarrow Pacifist$ $\sigma_2: Nixon \rightarrow Republican \rightarrow Pacifist$ $\sigma_3: Nixon \rightarrow Quaker$ $\sigma_4: Nixon \rightarrow Republican$

Like we have previously stated, paths σ_3 and σ_4 do not conflict with one another, and as such, they do not pose any intrinsic difficulty. The problem lies on the fact that the other two paths associated to Γ provide us reasons to support that *Nixon* inherits the attribute of being a *Pacifist* (through σ_1) and that *Nixon* inherits the attribute of not being a *Pacifist* (through σ_2). The coexistence of such conflicting information leads us to put forward the concept of an extension. We noted that an extension is a set that contains every single path associated with a given network, such that none of the paths in the set are conflicted. To put it another way, an extension is a maximally consistent unit of information associated with a given network. In the context of the Nixon Diamond, we have the following two extensions.



Figure 4.1.3-a Nixon Diamond Extension 1

On the one hand, and as portrayed in Figure 4.1.3-a, we have that one of the units of information, i.e. extensions of the Nixon Diamond, which we will refer to as Φ_1 can be characterised as follows Φ_1 : {{*Nixon* $\rightarrow Quaker$ }, {*Nixon* $\rightarrow Republican$ }, {*Nixon* $\rightarrow Quaker \rightarrow Pacifist$ }. Again, this is a maximally consistent unit of information associated with the network.



Figure 4.1.3-b Nixon Diamond Extension 2

On the other hand, and as portrayed in Figure 4.1.3-b, we have that the other unit of information, i.e. extension of the Nixon Diamond, which we will refer to as Φ_2 can be characterised as follows Φ_2 : {{*Nixon* \rightarrow *Republican*}, {*Nixon* \rightarrow *Quaker*}, {*Nixon* \rightarrow *Republican* \Rightarrow *Pacifist*}}. Like before, Φ_2 stands for a maximally consistent unit of information associated with the network. Based on the above, the Nixon Diamond has the following extensions:

$$\Phi_{1} = \{ \{Nixon \rightarrow Quaker\}, \\ \{Nixon \rightarrow Republican\}, \\ \{Nixon \rightarrow Quaker \rightarrow Pacifist\} \}$$

$$\Phi_2 = \{ \{ Nixon \rightarrow Republican \}, \\ \{ Nixon \rightarrow Quaker \}, \\ \{ Nixon \rightarrow Republican \not\rightarrow Pacifist \} \}$$

The above complies with the notion of an extension since both sets contain every single path associated with the network such that the paths in each extension are not conflicted. The aforesaid captures the intuitive notion that when facing conflicts, we branch the information into non-conflicted units. It is important to notice that each extension can be easily populated with as many non-conflicted paths as the network in question supports.

Based on the feature of networks allowing multiple extensions, there are particular arrangements which produce a poignant situation. In particular, there are networks with multiple extensions leading to conflicting conclusions, such as the Nixon Diamond. The above prompts the question of how to process a case in which alternative extensions lead to conflicting conclusions?

It turns out that there are two major strategies to address such a problem. One relies on the idea that more information is better, even if it is at the cost of conflicting information. The other strategy presents itself as a more conservative stance and only accepts information that is in the intersection of the extensions. In the following two sections, we examine both alternatives, and in doing so, we provide a progressive sketch to the most important approaches towards knowledge representation within inheritance networks.

4.2 Credulous Approach

In this section, we examine one of the two major strategies that addresses ambiguities. We analyse the credulous approach to ambiguity, but more importantly, the reasons and underlying intuitions for this information-processing mechanism. Here, we aim to demonstrate the impact that this approach to ambiguity has on the amount and type of information we are ultimately able to extract from inheritance networks.

4.2.1 Mechanism

The credulous approach to inheritance structures is based on the intuition that whenever possible, the information that is ultimately endorsed should be as much as it is available from a given network. This decision-making strategy assumes that more information is better, despite any conflicts that may come along. To put it another way, the credulous approach indulges in endorsing as much as it can extract from a given network, even if it is at the expense of having as a result opposing and conflicting information.

To address the aforementioned, the credulous approach operates on what it refers to as *credulous extensions*. A credulous extension is nothing more than an extension of a given network. That is, by definition *every extension of a given network is a credulous extension*. Such extensions are referred to as credulous in virtue of the fact that conflicting information is not seen as problematic given the information-hungry attitude of this stance. Thus, after having identified the credulous extensions of a given network, which are nothing more than the extensions of the network itself, the credulous approach states that the information that ultimately can be endorsed is the information contained in *any* of the credulous extensions (Horty, 2002)³⁰.

In light of the above, and recalling the structural arrangement of the *Nixon Diamond* we have a network Γ with the nodes {*Nixon*, *Republican*, *Quaker*, *Pacifist*} and the relations {*Nixon* \rightarrow *Quaker*, *Quaker*, *Quaker* \rightarrow *Pacifist*, *Nixon* \rightarrow *Republican*, *Republican* $\not\rightarrow$ *Pacifist*}. Based on the above, this network has the following extensions:

$$\Phi_1 = \{ \{Nixon \rightarrow Quaker\}, \\ \{Nixon \rightarrow Republican\}, \\ \{Nixon \rightarrow Quaker \rightarrow Pacifist\} \}$$

$$\Phi_2 = \{ \{ Nixon \rightarrow Republican \}, \\ \{ Nixon \rightarrow Quaker \}, \\ \{ Nixon \rightarrow Republican \rightarrow Pacifist \} \}$$

³⁰ [...] either endorsing the set of conclusions supported by an arbitrary one of the several argument extensions, or perhaps endorsing a conclusion as believable whenever it is supported by some extension or another. (Horty, 2002, 59)

The question revolves around what information are we to endorse from such network, i.e. should we take the stance that Nixon inherits the attribute of being a Pacifist or should we adopt the view that Nixon inherits the attribute of not being a Pacifist? Despite which position we take, we need to provide reasons as to why such an approach is better than the alternative.

As we have stated, the first task is to identify the credulous extensions of the network. Nevertheless, all of the extensions of a given network are, by definition, credulous extensions. In our particular case study, we have that both Φ_1 and Φ_2 are to be taken as credulous extensions of the network. Having identified the extensions of the network, we pointed out that the credulous approach states that any of the information contained in the credulous extensions is legitimate information. Based on the credulous approach, we have equal grounds to infer that the Nixon node inherits the *Pacifist* attribute, as well as stating that the Nixon node inherits the attribute of not being a *Pacifist*.

The above, summarises the credulous reasoning policy for inheritance networks, along with its grounding intuitions. In the following subsection, we assess the extent and limitation of this approach of information processing within inheritance networks.

4.2.2 Discussion

It is important to recall that one feature of ambiguity involves networks having multiple extensions (contrary to a unique extension). The credulous strategy presents itself as a mechanism to process information in the aforesaid circumstances. As we have previously stated, the credulous stance endorses an attribute whenever such attribute belongs to any of the available extensions. Nevertheless, it is commonly agreed that the credulous approach as a mechanism to compute information in the context of inheritance networks is conspicuously unsatisfactory. The credulous stance provides no useful decision mechanism when the inheritance structure produces a set of conflicting or mutually exclusive conclusions, and leaves as the only alternative a choice that seems to be an unsound foundation to ground a generally extensible inference mechanism.

In line with the above, the initial question that lies at the core of the mechanisms to compute information within inheritance structures remains unaddressed. In the case of the Nixon Diamond, *Nixon* inherits the attribute of being a *Pacifist*, and it also inherits the attribute of not being a

Pacifist. Thus, the credulous "solution" seems to entirely bypass the problem itself, since the precise nature of the problem is deciding which of the attributes *Nixon* ought to inherit (either being or not being a *Pacifist*). Hence, the proposed solution of allowing either of them to be endorsed seems to elude the problem rather than addressing it as noted by Horty:

[..] this variant of the second option also manages to sidestep our original question. We wanted to know what conclusions we should actually draw from the information provided by a default theory--- whether or not, given the information from the Nixon Diamond, we should conclude that Nixon is a Pacifist, for example. But according to this variant, we are told only what there is good reason to believe---that both B(Pn) and $B(\neg Pn)$ are consequences of the theory, so that there is good reason to believe that Nixon is a Pacifist, but also good reason to believe that he is not. This may be useful information, but it is still some distance from telling us whether or not to conclude that Nixon is a Pacifist. (2002, 57)

This common criticism is one of the bases why this reasoning policy is not taken as a sensible solution to the problem of ambiguity. Thus, when there are multiple extensions associated with a particular network, a more prudent form of handling information seems to be required if the inferences are going to have a sound justification.

Having stated the above, in the next section, we proceed to inquire in detail into one alternative mechanism that is placed in direct opposition to the credulous strategy for cases involving ambiguity, i.e. *skeptical* approaches. The skeptical approach is placed as an improvement over the credulous account of information processing within networks.

4.3 Skeptical Approaches

In this section, we analyse the skeptical account of information-processing within inheritance networks. Moreover, we emphasise the main ideas of the skeptical approach and how it stands in direct opposition to the credulous strategy. The main objective of this section is to provide a review of skeptical approaches to processing information within inheritance networks. More precisely, we examine how both approaches address the Nixon Diamond. Our aim at this stage is to stress the various skeptical strategies to approach ambiguities as information processing mechanisms. More importantly, we aim to present each account as an information extraction policy, which substantially varies the information we are ultimately able to extract from inheritance networks.

4.3.1 Mechanism

In this subsection, we provide a general depiction of the underlying idea grounding the skeptical reasoning policy within inheritance networks in the context of ambiguity. Nevertheless, given that this particular stance has at least two ways of being instantiated, we address each specific implementation in further subsections.

First, we recall the problem at hand. We have a network Γ with the nodes {*Nixon*, *Republican*, *Quaker*, *Pacifist*} and the relations {*Nixon* \rightarrow *Quaker*, *Quaker* \rightarrow *Pacifist*, *Nixon* \rightarrow *Republican*, *Republican* \rightarrow *Pacifist*}. Based on the above, this network has the following extensions:

$$\Phi_{1} = \{ \{Nixon \rightarrow Quaker\}, \\ \{Nixon \rightarrow Republican\}, \\ \{Nixon \rightarrow Quaker \rightarrow Pacifist\} \}$$

$$\Phi_2 = \{ \{ Nixon \rightarrow Republican \}, \\ \{ Nixon \rightarrow Quaker \}, \\ \{ Nixon \rightarrow Republican \not Pacifist \} \}$$

The question is whether we should take the stance that *Nixon* inherits the attribute of being a Pacifist or should we adopt the view that *Nixon* inherits the attribute of not being a Pacifist. Regardless of the stance we take, we need to provide reasons as to why either approach is better than the alternative.

The skeptical strategy, in the context of inheritance networks, is proposed as a more sensible solution to cases involving ambiguity. The underlying intuition of this approach is that, the information that should be endorsed from a particular structural arrangement with multiple extensions is not whatever is contained in any of the extensions, like in the credulous stance (which is why the credulous stance is often seen as a gullible approach to discern and endorse information

from ambiguous and otherwise conflicting networks). To the contrary, a skeptical stance states that an attribute should be endorsed whenever such attribute is contained in the *intersection of all extensions* associated with the network. This extension-computing policy aims to provide a much more prudent scrutiny of the information it is willing to endorse. The skeptical reasoning policy stands in direct opposition to the credulous approach, as the former is much more conservative than the latter.

Despite the straightforward and simple motivation behind the skeptical approach, as one that endorses attributes or conclusions only when such information is contained in the intersection of all the extensions, there are two different ways to accomplish such operation. On the one hand, one could collect the attributes endorsed by all the extensions associated with a given network, *and then* apply the intersection operation on such a set of endorsed attributes. This form of skepticism is referred to as an *indirect* approach. On the other hand, the operation of intersecting all extensions associated to the network can be applied not to the set of attributes endorsed by the extensions associated with the network, but rather to the paths contained in the extensions. This form of skepticism is referred to as a *direct* approach.

Although the distinction is subtle, it relies on the same skeptical intuition. Nevertheless, it provides two completely different mechanisms to compute information within networks. In some cases, they provide the same information, whereas, in other cases, they produce different results. This is the core problem of our research.

The above concludes our general sketch of the underlying motivations behind the skeptical stance towards ambiguous cases as a reasoning policy. In the following subsections, we present the specific details and implementation of each version of skeptical reasoning.

4.3.2 Indirect Skepticism

As we previously stated, the crucial intuition underlying the skeptical approach for networks involving multiple and potentially conflicting information lies in the idea of intersecting all the extensions. The indirect version of this mechanism is instantiated through the process of determining when an *attribute* is common to all extensions. This approach is also referred to as an *attribute intersection* approach (Horty, 2012), but we will refer to it as the *indirect* approach

(Makinson & Schlecta, 1991). This approach can be divided into a three-step process.

First, we assemble all the extensions associated with a given network. Like we already mentioned, every extension associated with a given network is, according to credulous standards, a credulous extension. For such reason, skepticism (both direct and indirect) is often described as operating on credulous extensions. Thus, the first step, according to direct skepticism, is as follows:

$$\{\Phi_n \text{ is an extension } \in \Gamma\}$$

Second, and after having identified all the credulous extensions, the indirect approach proceeds to construct a further set that is composed of the attributes endorsed by every credulous extension. The aforementioned set is referred to as the set of *credulously endorsed attributes*. To account for this, we use a consequence function cn() that takes as input the credulous extensions associated with the network and return a set of attributes. We apply the cn() function to the extensions associated with the network to get the endorsed attributes. Thus, this step operates as follows:

 $cn(\{\Phi_n \text{ is an extension } \in \Gamma\})$

Third, the indirect approach applies the intersection operation to the set containing the *credulously endorsed attributes*. Thus, the result of the intersection operation contains not paths (the elementary units of the extensions) but rather *attributes*. The reason why the resulting set contains attributes rather than paths lies in the fact that the set on which the intersection operation proceeds is composed of attributes rather than paths. The resulting set of attributes contains the values generated by *cn()* from a given network *following an indirectly skeptical approach*. The general mechanism of indirect skepticism can be summarised in the following procedure:

 \cap (cn({ Φ_n is an extension $\in \Gamma$ }))

The above states that the intersection operation \cap is applied to the output of cn(). Here, the cn() function takes paths (the base units of extensions) as inputs and returns attributes. Thus, the intersection operation computes attributes and not paths within extensions.

According to this form of skepticism the legitimacy to endorse an attribute is determined by such

attribute being common to *every* extension, i.e. such attribute is endorsed by *every* extension. Thus, whenever an attribute is contained in every extension of a given structure, the attribute is referred to as a *skeptically endorsed statement* (Makinson & Schlecta, 1991). The indirect approach ultimately seeks to identify *skeptically endorsed statements*, and only such statements are taken as legitimate conclusions to be inferred from a given network.

This concludes the general description of indirect skepticism. In the following subsection, we use this skeptical strategy to address the *Nixon Diamond*.

4.3.3 Indirect Approach to the Nixon Diamond

As previously remarked, the ambiguity revolves around the *Pacifist* attribute and the problem is whether such attribute or its negation should be inferred from the network. Following indirect skepticism, the problem can be reduced to determine if either the positive or the negative attribution to the Pacifist property is common to every credulous extension.

To address the Nixon Diamond according to an indirect approach, we proceed with the same threestep process. First, we identify all the credulous extensions. In the context of the *Nixon Diamond*, i.e. the network Γ , we have the following extensions:

$$\Phi_1 = \{ \{Nixon \rightarrow Quaker\}, \\ \{Nixon \rightarrow Republican\}, \\ \{Nixon \rightarrow Quaker \rightarrow Pacifist\} \}$$

$$\Phi_{2}=\{\{Nixon \rightarrow Republican\},\\ \{Nixon \rightarrow Quaker\},\\ \{Nixon \rightarrow Republican \not\Rightarrow Pacifist\}\}$$

Second, and after having identified the credulous extensions, we construct a set containing all the attributes endorsed by each of the credulous extensions. To account for this, we make use of the consequence operation cn() that takes extensions as input, and outputs the attributes endorsed by the extension. Thus, we apply the cn() operation to the extensions associated with the network. In our specific case, we apply the operation as follows:

 $cn(\Phi_1, \Phi_2)$

The above is just a shorter version of the following operation:

$$cn(\{\{Nixon \rightarrow Quaker\}, \\ \{Nixon \rightarrow Republican\}, \\ \{Nixon \rightarrow Quaker \rightarrow Pacifist\}\}, \\ \{\{Nixon \rightarrow Republican\}, \\ \{Nixon \rightarrow Quaker\}, \\ \{Nixon \rightarrow Republican \Rightarrow Pacifist\}\}) \\ = \{\{Quaker, Republican, Pacifist\}, \}$$

{*Quaker*, *Republican*}}

Third, and once we have constructed the credulously endorsed attributes set (in our particular example the resulting set of the *cn()* operation is {{*Quaker*, *Republican*, *Pacifist*},{*Quaker*, *Republican*}}), we apply the intersection operation to such set. Thus, the intersection operation is applied to the output of the *cn()* operation. As such, we have the following:

 $\cap(\{Quaker, Republican, Pacifist\},$ $\{Quaker, Republican\})$

= {Quaker, Republican}

This operation yields the {*Quaker*, *Republican*} set as a result. In particular, since neither the positive *isa* nor the negative *not-isa* connection to the *Pacifist* attribute is contained in every credulously endorse attribute set, the result of the intersection operation yields a set which contains nothing *with respect to such attribute*. The aforementioned amounts to the fact that, in the case of positive *isa* connection to the *Pacifist* attribute, although it is contained in the Φ_1 extension, such attribute is not contained in the Φ_2 extension. As such, the attribute cannot be said to be common to all extensions. The same occurs with the *not-isa* attribute of the Nixon node to the *Pacifist* node, which is contained in the Φ_2 extension, but it is not contained by the Φ_1 extension. Thus, the property attribution cannot be said to be common to all extensions. In such a vein, we have that

neither the positive nor the negative attribution to the *Pacifist* property can be inferred. Thus, judgement on whether Nixon is a *Pacifist* or not is suspended.

The above concludes our presentation of the indirect version of the skeptical reasoning policy within inheritance networks, and how it processes conflicting information in ambiguous networks. Furthermore, we inquired how it diverges from the credulous approach, and how such divergence can be seen as a further refinement of the credulous account of information processing. In the next subsection, we look into direct skepticism.

4.3.4 Direct Skepticism

The direct approach to skeptically process information is different from the indirect approach, in a subtle, yet significant way. Where the indirect approach applies the intersection operation to the set of attributes endorsed by the extensions associated with the network, direct skepticism applies the intersection operation to *paths*. Therefore, the goal from the directly skeptical point of view is to identify a *path that is common to all extensions*. Such a path is what determines the information that can be inferred from a given network. The direct approach can be divided into a three-step process.

First, the direct approach collects all the credulous extensions associated with a given network (which is the only commonality between these two procedures). As we have stated, due to the way this is conceived, every extension associated with a given network is, by definition, a credulous extension. Hence, this version of skepticism can also be said to operate on credulous extensions as a preliminary step as well. The first step according to direct skepticism is as follows:

$\{\Phi_n \text{ is an extension } \in \Gamma\}$

Second, once the direct account has constructed the set of all extensions associated to the network, instead of constructing a further set containing the attributes endorsed by all of the credulous extensions employing the cn() operation, it applies the intersection operation directly to the paths contained in the credulous extensions. In such a vein, and unlike in the case of indirect skepticism, given that the direct account operates on a set containing not attributes but paths (as the elementary units of extensions are paths rather than attributes), the intersection operation returns a set
containing a paths, not attributes. The resulting paths can be seen as the elements of what is referred to as a *skeptically endorsed extension*. This *skeptically endorsed extension* is the output of the following procedure:

$$\cap$$
 ({ Φ_n is an extension $\in \Gamma$ })

Third, direct skepticism proceeds to apply the cn() operation to the aforesaid *skeptically endorsed extension*. Thus, the skeptical approach allows inferring whatever is supported by the skeptically endorsed extension. In particular, we have the following final and general mechanism of direct skepticism:

$$cn(\cap(\{\Phi_n \text{ is an extension } \in \Gamma\}))$$

The above states that the consequence operation cn() is applied to the output of the intersection operation \cap . Like in the case of indirect version of skeptical the consequence operation cn() acts on paths (i.e. the elementary components of extensions) rather than attributes. Nevertheless, in the case of direct skepticism it operates on what is referred to as a *skeptically endorsed extension*. This subtle distinction is what sets apart both approaches to skeptical inference.

According to direct skepticism the legitimacy to endorse an attribute is determined by such attribute being contained in a path that is common to all extensions. Thus, the sole aim of this version of skeptical inference is to identify a *skeptically endorsed extension*, also referred to as *skeptically acceptable arguments*. Horty describes the aforesaid as follows:

In contrast to the multiple-extension accounts considered so far, that theory first defined a single argument extension that was thought of as containing the "skeptically acceptable" arguments based on a given inheritance network. The skeptical conclusions were then defined simply as the statements supported by those skeptically acceptable arguments. (2002, 61)

Then, we would be able to infer whatever is supported by such a skeptically endorsed extension (or skeptically acceptable argument).

The above concludes the sketch of direct skepticism. In the following subsection, we use this skeptical strategy to address the Nixon Diamond.

4.3.5 Direct Approach to the Nixon Diamond

Like in the previous subsection, the problem for direct skepticism remains fixed on the ambiguity revolving around the *Pacifist* attribute and whether such attribute or its negation can be inferred from the network. Nevertheless, following direct skepticism, the problem is reduced to determining if either the positive or the negative attribution to the Pacifist property is endorsed by a path common to all credulous extensions of the network.

To address the Nixon Diamond according to a direct approach, we follow the same three-step process. First, we identify all the credulous extensions. In the context of the Nixon Diamond, we have the following extensions:

$$\Phi_1 = \{ \{Nixon \rightarrow Quaker\}, \\ \{Nixon \rightarrow Republican\}, \\ \{Nixon \rightarrow Quaker \rightarrow Pacifist\} \}$$

$$\Phi_2 = \{ \{ Nixon \rightarrow Republican \}, \\ \{ Nixon \rightarrow Quaker \}, \\ \{ Nixon \rightarrow Republican \not\rightarrow Pacifist \} \}$$

Second, and unlike the indirect approach, which extracts the credulously endorsed statements from the credulous extensions, the direct approach intersects the set containing the credulous extensions, to determine the *skeptically endorsed extension*, i.e. the {{ $Nixon \rightarrow Quaker \rightarrow Pacifist$ }, { $Nixon \rightarrow Republican \Rightarrow Pacifist$ } set. Thus, we have the following:

$$\cap(\Phi_1, \Phi_2) = \{\{Nixon \rightarrow Republican\}, \{Nixon \rightarrow Quaker\}\}$$

The above is just a shorter version of the following operation:

$$\bigcap(\{\{Nixon \to Quaker\}, \\ \{Nixon \to Republican\}, \\ \{Nixon \to Quaker \to Pacifist\}\}, \\ \{\{Nixon \to Republican\}, \\ \{Nixon \to Quaker\}, \\ \{Nixon \to Republican \not Pacifist\}\}) \\ = \{\{Nixon \to Republican\}, \{Nixon \to Quaker\}\}$$

The process of intersecting the credulous extensions involved in the Nixon Diamond returns the $\{\{Nixon \rightarrow Republican\}, \{Nixon \rightarrow Quaker\}\}$ set. This result lies in the fact that only paths σ_3 and σ_4 are paths common to all *extensions*.

Finally, we apply the consequence operation cn() to the output of the $\cap()$ operation. Unlike the indirect approach, here we take in as input the result of the intersection of the credulous extensions rather than every credulous extension. To put it another way, the cn() now operates on the *skeptically endorsed extension* rather than on the collection of credulous extensions. Thus, we have the following:

 $cn(\{Nixon \rightarrow Republican\}, \\ \{Nixon \rightarrow Quaker\}) \\ = \{Republican, Quaker\}$

The previous operation yields the {*Republican*, *Quaker*} set as a result. Like the indirect approach, we first identify the credulous extensions of the network, which are the same given that we are operating on the same network, i.e. both Φ_1 and Φ_2 are credulous extensions of the *Nixon Diamond*. Second, and unlike indirect skepticism, which further extracts the credulously endorsed statements from the credulous extensions, the direct approach immediately intersects the set containing the credulous extensions, to determine the skeptically endorsed extension. In the context of the Nixon Diamond, only σ_3 and σ_4 are paths that are common to all extensions. Therefore, from a directly skeptical stance, the result of intersection of the extensions in the Nixon Diamond is the {*Nixon* \rightarrow *Republican*}, {*Nixon* \rightarrow *Quaker*} set, which in turn would be our skeptically endorsed extension. Consequently, the *cn()* operates on the aforesaid set and yields {*Republican*, *Quaker*} as the resulting set.

In conclusion, from a directly skeptical account, we have that, *with respect to the Pacifist attribute*, the *Nixon* node inherits neither the positive connection nor the negative connection to the *Pacifist* attribute. The reason is that, in the context of the Nixon Diamond, no path is common to all extensions, which enables a positive or negative attribute to such node.

To summarise, we have provided a depiction of the direct account of the skeptical reasoning, along with its underlying intuitions. Furthermore, we have seen how this version of skepticism processes conflicting information in cases involving ambiguities, and how it differs from the indirect approach. We have also pointed out how, despite the difference in the approach to ambiguous cases, direct skepticism ultimately endorses the same results as indirect skepticism for the Nixon Diamond. In the next subsection, we provide a general overview of the various reasoning policies we have reviewed, and assess such strategies.

4.3.6 Discussion

It is important to recall that a feature of ambiguous cases relies on a network having more than one extension. Thus, both mechanisms of information processing within inheritance networks address this specific condition. Skeptical approaches are regarded as a more conservative approach to such cases, and as a further refinement of the credulous strategy. That is, the skeptical strategy to computing information can be seen as an extension of the credulous approach (Schlechta, 1993)³¹. In particular, skeptical approaches to computing ambiguities also build on top of the notion of a credulous extension. As previously mentioned, this is one of the few similarities between both versions of skepticism. The other common thread between both approaches underlies in the methodological thesis that skeptically acceptable information is to be determined by the intersection operation applied to all of the involved extensions. Nevertheless, this is precisely what produces the divergence between these two mechanisms.

The difference between both versions of skepticism revolves as to how the idea of intersecting the credulous extensions is carried out. On the one hand, the indirect approach applies the intersection operation to the set of the *credulously endorsed statements*. Indirect skepticism looks for those attributes that are common to every extension. Skeptically acceptable conclusions are, according

³¹ The extensions approach can, in a second step, be turned into an (indirectly) sceptical one by considering only the information contained in all extensions. (Schlechta, 1993, 455)

to this view, the attributes that are contained in the resulting set of intersecting the credulously endorsed attributes. On the other hand, the direct approach applies the intersection operation not to sets of endorsed attributes, but rather to paths that belong to the *skeptically endorsed extension*. Direct skepticism does not look for the attributes that are common to every extension, but rather for a path that is common to every extension. This alternative view states that the skeptically acceptable information is that which is supported by a path within the skeptically acceptable extension.

Thus, the main difference between both versions of skepticism is that according to the indirect approach the input to which the intersection operation is applied to are sets containing attributes, whereas the direct approach applies the intersection operation not to attributes, but rather to paths. The result of the intersection operation is an attribute in the case of indirect skepticism, and a path in the case of direct skepticism.

Finally, and although some common terminology has been taken for granted it is important to take note of some relevant features of these concepts. When we refer to skepticism, we are very well aware that this term is entrenched in philosophical discourse. Skepticism in this "classical" sense is associated with the various challenges raised against the certainty of knowledge and is allocated within the field of epistemology. Nevertheless, as we have employed this term so far, and will continue to do so, it has nothing to do with such sense of skepticism.

Skepticism, in the context as employed here and throughout the literature on inheritance networks and non-monotonic logics, is regarded as a stance towards inferences involving multiple extensions. That is, skepticism should be regarded as strategy, mechanism or reasoning policy that tries to address adequately scenarios like the Nixon Diamond and other networks involving multiple extensions. In this sense, Skepticism is presented as a prudent or conservative strategy when facing these scenarios involving multiple extensions insofar as they tend to draw less information when multiple and potentially conflicting extensions are involved. This reasoning policy is better understood when compared to the more licentious or liberal reasoning policy offered by credulous approaches to networks with multiple extensions.

As we have pointed out, the credulous approaches take the stance that more information is better even if it comes at the expense of conflict or inconsistencies. In this sense, the stance that deviates from such liberal policy and endorses only the information that is in the intersection of the various extensions associated to a given network is to be regarded as skeptical, prudent or more conservative, but has no relation to the common understanding of the notion in current philosophical discourse. In the same vein, when we refer to an inference being drawn in a skeptical fashion or through a skeptical standard, we will appeal to this idiosyncratic notion of skepticism as a very conservative attitude towards information extraction within inheritance networks as opposed to a notion of skepticism more general or commonly used in the philosophical discourse.

Thus, the central problem of our research can be rephrased as the pursuit of assessing and determining the best form of skeptical standards for information processing within inheritance networks for non-monotonic reasoning. As such, we will be flexible with this notion of skepticism as it is going to be an object of debate and scrutiny throughout our research.

The above concludes the preliminary exposition of the alternative approach of the credulous strategy to ambiguous situations. We highlighted how this approach has not one but two different procedures through which it can be instantiated and how they coincide and diverge.

4.4 Chapter Summary

In this chapter, we reviewed the problem of ambiguity. We discussed the extent and limitation of such problem, and how it produces the notion of extensions of a network, as the first step to account for ambiguities. We reviewed two approaches to address ambiguous situations, namely credulous and skeptical approaches. Both approaches represent different accounts of information processing within inheritance networks. Lastly, we emphasised the existence of a twofold solution to the intuition behind skepticism, and its significance.

Despite the above, we have yet to address, in a comprehensive fashion, the possibility of both versions of skepticism prompting different results, how such situations might emerge and what problems could this entail.

In the next chapter, we address the aforesaid as well as a crucial corollary of this problem. More precisely, we consider cases in which the two versions of skeptical reasoning provide different results, which pose a central problem for the skeptical stance. This is the core problem of our resear

Chapter 5 Floating Conclusions

Throughout the preceding chapters, we have examined inheritance networks as representational structures. We discussed the ambiguity problem and the different strategies to approach such networks. In this chapter, we analyse a problem that emerges from the various strategies to address ambiguities, namely the problem of floating conclusions. The aim of this chapter is to address the key problem of our research, namely, the problem of floating conclusions. The aforesaid leads to the proposal of a framework of analysis (chapter 6), which we later apply to the problem of floating conclusions (chapter 7).

5.1 Producing Floating Conclusions

In this section, we examine the preliminary aspects of the central problem of our research. We analyse the problem of floating conclusions. In particular, we examine the structural conditions through which floating conclusions emerge, as well as the technical features surrounding such networks. The main objective of this section is to have a preliminary overview of the main aspects surrounding the phenomena of floating conclusions. In particular, we aim to provide a technically accurate but intuitive exposition of the underlying problem, and to present a case study through which we can analyse the precise dynamics surrounding the problem.

5.1.1 Emergence

One of the salient features of inheritance networks is the way in which information is inferred through the concept of an *extension* Horty (2002). In such a vein, we have considered networks involving multiple extensions that lead to conflicting information. To address such problem, we reviewed two of the most common solutions to extract information from such situations, namely: *credulous* and *skeptical* approaches. As we mentioned, the credulous approach is the most liberal attitude towards information processing, since all that it takes for an attribute to being endorsed from a given network is for that attribute to be allocated in *some extension*. We highlighted that this strategy, rather than addressing the problem seems to bypass it entirely. In the same vein, we

reviewed the skeptical stance, which operates on the assumption that only information that is contained in the *intersection of all extensions* is suitable to be information inferred from a given network.

Within the skeptical approach, we examined two ways of addressing the *dictum* of inferring information contained on the intersection of all credulous extensions, namely the *indirect* and *direct* approach. We noted the commonalities and the divergence between both approaches. Nevertheless, and what is important to point out is that, despite being different mechanisms to instantiate the intersection of extensions operation, in the case of the Nixon Diamond, they both produced the same outcome. Despite this, we have not yet addressed in detail networks in which both mechanisms provide different results.

There are particular structural arrangements which, when considered through a skeptical lens, will provide different results depending on which type of skeptical mechanism is applied. The reasons and conditions upon which the aforesaid discrepancy emerges is the *quid* of this entire research. In light of the above, we analyse how an attribute can be contained in every credulous extension associated with a network (i.e. such an attribute is considered a *skeptically endorsed statement*) but it is not contained in a path common to all of the extensions of the network (i.e. it is not supported by a *skeptically endorsed argument*). Such situations are known as *floating conclusions*.

Floating conclusions refer to cases in which a specific attribute can be located in every credulous extension associated to a given network, yet such attribute is not endorsed by a path common to every single extension. These conclusions are referred to as *floating conclusions*, precisely because they "float" above all of the credulous extensions, yet they cannot be located within a path common to all of the credulous extensions (Prakken, 2002). As defined by Makinson and Schlechta:

[...] there are quite simple examples of nets N and propositions a such that a is skeptically acceptable on even the most rigorous intuitive standards, even though there is no skeptically acceptable path to support it. This is because in the examples, under any reasonable account of extensions, every extension contains some path supporting a, but there is no path supporting a that is common to all extensions. We shall call such a proposition a floating conclusion. (1991, 202)

More importantly, the existence of floating conclusions produces the philosophical question regarding the underlying epistemological, structural or otherwise reasons that produce such scenarios, and how should we proceed.

5.1.2 Case Study: Paths

In this subsection, we consider a case study in which floating conclusions are instantiated. In particular, we sketch the particular structural arrangement, but more importantly, we depict in detail the various paths associated with this network. Identifying of the extensions of this network is the topic of the following subsection.



Figure 5.1.2-a Seedless-Grape-Vine network (sgv-net)

To address the phenomena of floating conclusions, we take the network depicted in Figure 5.1.2a as a working example. We refer to this network as the sgv-net. The sgv-net is composed of the nodes {Seedless.Grape.Vine, Grape.Vine, Vine, Arbor.Plant, Plant, Seedless.Thing, Fruit.Plant, Tree} and the links {Seedless.Grape.Vine \rightarrow Grape.Vine, Grape.Vine \rightarrow Vine, Vine \rightarrow Arbor.Plant, Arbor.Plant \rightarrow Plant, Seedless.Grape.Vine \rightarrow Seedless.Thing, Seedless.Thing \neq *Fruit.Plant, Grape.Vine* \rightarrow *Fruit.Plant, Fruit.Plant* \rightarrow *Tree, Tree* \rightarrow *Plant, Fruit.Plant* \rightarrow *Arbor.Plant*}. The sgv-net entails various paths that contain different information with respect to the Seedless.Grape.Vine node. We consider them one by one.



Figure 5.1.2-b SGV-Net First Path

First, we have the route that follows the *Grape.Vine* node. The simplest path would be the *Seedless.Grape.Vine* \rightarrow *Grape.Vine* \rightarrow *Grape.Vine* \rightarrow *Grape.Vine* \rightarrow *Vine*, *Seedless.Grape.Vine* \rightarrow *Grape.Vine* \rightarrow *Vine*, *Seedless.Grape.Vine* \rightarrow *Grape.Vine* \rightarrow *Vine* \rightarrow *Arbor.Plant*, and *Seedless.Grape.Vine* \rightarrow *Grape.Vine* \rightarrow *Vine* \rightarrow *Arbor.Plant*, this last path is shown in Figure 5.1.2-b).



Figure 5.1.2-c SGV-Net Second Path

Second, we have the route that follows the *Fruit.Plant* property through the *Grape.Vine* node. Here we have the *Seedless.Grape.Vine* \rightarrow *Grape.Vine* \rightarrow *Fruit.Plant* and the *Seedless.Grape.Vine* \rightarrow *Grape.Vine* \rightarrow *Fruit.Plant* and the *Seedless.Grape.Vine* \rightarrow *Grape.Vine* \rightarrow *Fruit.Plant* \rightarrow *Arbor.Plant* paths (the last one depicted in Figure 5.1.2-c)



Figure 5.1.2-d SGV-Net Third Path

Third, following the same thread as the above, i.e. going through the *Fruit.Plant* attribute through the *Grape.Vine* node, we have the *Seedless.Grape.Vine* \rightarrow *Grape.Vine* \rightarrow *Fruit.Plant* \rightarrow *Tree* and *Seedless.Grape.Vine* \rightarrow *Grape.Vine* \rightarrow *Grape.Vine* \rightarrow *Fruit.Plant* \rightarrow *Tree* \rightarrow *Plant* paths. Unlike the previous paths, this variation, as depicted by Figure 5.1.2-d follows the *Tree* node and ultimately connects to the *Plant* attribute.



Figure 5.1.2-e SGV-Net Fourth Path

Finally, we have the *Seedless.Grape.Vine* \rightarrow *Seedless.Thing* and the *Seedless.Grape.Vine* \rightarrow *Seedless.Thing* \rightarrow *Fruit.Plant* paths (the last one depicted in Figure 5.1.2-e). Unlike the previous paths, it takes the *Seedless.Thing* node as a route, and from there it is negatively linked to the *Fruit.Plant* property. This last *Seedless.Thing* \rightarrow *Fruit.Plant* not-isa connection terminates this short-lived path.

It is important to stress that the previously listed possibilities exhaust the paths associated with the sgv network. Thus, we have the following set:

 $\sigma_{1}: Seedless.Grape.Vine \rightarrow Grape.Vine$ $\sigma_{2}: Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine$ $\sigma_{3}: Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine \rightarrow Arbor.Plant$ $\sigma_{4}: Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine \rightarrow Arbor.Plant \rightarrow Plant$ $\sigma_{5}: Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant$ $\sigma_{6}: Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant \Rightarrow Arbor.Plant$ $\sigma_{7}: Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant \rightarrow Tree$ $\sigma_{8}: Seedless.Grape.Vine \rightarrow Seedless.Thing$ $\sigma_{10}: Seedless.Grape.Vine \rightarrow Seedless.Thing \Rightarrow Fruit.Plant$

So far, we have described the various paths associated to the *sgv-net*. Nevertheless, we need to further examine the extent to which they can consistently coexist as a unit of information. Unsurprisingly, there is no way in which we can allocate them to a single set without producing conflicting situations.

5.1.3 Case Study: Extensions

In the context of the *sgv-net*, the question revolves around whether the *Seedless.Grape.Vine* inherits the *Plant* attribute. To address this, we encounter a first problem emerging from such network. The problem comes from the ambiguous *Fruit.Plant* node, from which we have two choices. We can either endorse the *Fruit.Plant* attribute or decline to endorse such attribute.



Figure 5.1.3-a SGV-Net First Conflict

First, as Figure 5.1.3-a depicts, we can decide to endorse the *Fruit.Plant* attribute to the *Seedless.Grape.Vine* node. Under this account, the *Seedless.Grape.Vine* \rightarrow *Grape.Vine* \rightarrow *Fruit.Plant* path is upheld. Nevertheless, such state of affairs conflicts with the *Seedless.Grape.Vine* \rightarrow *Seedless.Thing* \Rightarrow *Fruit.Plant*, *Seedless.Grape.Vine* \rightarrow *Grape.Vine* \rightarrow *Grape.Vine* \rightarrow *Wine* \rightarrow *Arbor.Plant* and *Seedless.Grape.Vine* \rightarrow *Grape.Vine* \rightarrow *Vine* \rightarrow *Arbor.Plant* and *Seedless.Grape.Vine* \rightarrow *Wine* \rightarrow *Arbor.Plant* \rightarrow *Plant* paths.



Figure 5.1.3-b SGV-Net Second Conflict

Second, as Figure 5.1.3-b depicts, we can choose the negative connection of the *Seedless.Grape.Vine* node to the *Fruit.Plant* attribute, through the *Seedless.Grape.Vine* \rightarrow *Seedless.Thing* \Rightarrow *Fruit.Plant* path. Under this account, we do not have the *Seedless.Grape.Vine* \rightarrow *Grape.Vine* \rightarrow *Fruit.Plant*, *Seedless.Grape.Vine* \rightarrow *Grape.Vine* \rightarrow *Fruit.Plant*, *Seedless.Grape.Vine* \rightarrow *Grape.Vine* \rightarrow *Grape.Vine* \rightarrow *Grape.Vine* \rightarrow *Grape.Vine* \rightarrow *Grape.Vine* \rightarrow *Grape.Vine* \rightarrow *Fruit.Plant* \rightarrow *Tree* and the *Seedless.Grape.Vine* \rightarrow *Grape.Vine* \rightarrow *Fruit.Plant* \rightarrow *Tree* \rightarrow *Plant* paths. In particular, these last paths would stand in conflict with the *Seedless.Grape.Vine* \rightarrow *Seedless.Thing* \Rightarrow *Fruit.Plant* path.

Thus, the *sgv-net* instantiates a network in which we have conflicting paths as a byproduct of ambiguity. Like with the Nixon Diamond, this prompts the use of extensions to distinguish the units of non-conflicting information associated with the network. In this sense, we can arrange all the paths associated with this network into two separate units of information that can stand without conflict.



Figure 5.1.3-c SGV-Net First Extension

First, as Figure 5.1.3-c depicts, one way to resolve the ambiguity located at the *Fruit.Plant* node. This account relies on upholding the positive *isa* connection to such attribute. If this is the case, paths σ_1 , σ_2 , σ_5 , σ_6 , σ_7 , σ_8 , and σ_9 do not stand in conflict. As such, they stand as a consistent unit of information associated with the *sgv-net*. Under this state of affairs, the first extension is $\Phi_1: \{\sigma_1, \sigma_2, \sigma_5, \sigma_6, \sigma_7, \sigma_8, \sigma_9\}$. The extension Φ_1 can be expanded as follows:

$$\begin{split} \Phi_{1} &= \{ \{ Seedless.Grape.Vine \rightarrow Grape.Vine \}, \\ \{ Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine \}, \\ \{ Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant \}, \\ \{ Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant \Rightarrow Arbor.Plant \}, \\ \{ Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant \rightarrow Tree \}, \\ \{ Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant \rightarrow Tree \rightarrow Plant \}, \\ \{ Seedless.Grape.Vine \rightarrow Seedless.Thing \} \} \end{split}$$



Figure 5.1.3-d SGV-Net Second Extension

Second, as Figure 5.1.3-d depicts, the other way in which we can resolve the ambiguity contained in the *Fruit.Plant* node. This account relies on choosing the *Seedless.Grape.Vine* \rightarrow *Seedless.Thing* \rightarrow *Fruit.Plant* path, ie. path σ_{10} . Under this approach to the ambiguous node, the *not-isa* connection to the *Fruit.Plant* property is upheld. The side effect of this way to resolve the ambiguity is that paths σ_1 , σ_2 , σ_3 , σ_4 , σ_9 , and σ_{10} are not in conflict. Under this state of affairs, the second extension is Φ_2 : { σ_1 , σ_2 , σ_3 , σ_4 , σ_9 , σ_{10} }. The extension Φ_2 can also be expanded as follows:

$$\begin{split} \Phi_2 &= \{ \{ Seedless. Grape. Vine \rightarrow Grape. Vine \}, \\ \{ Seedless. Grape. Vine \rightarrow Grape. Vine \rightarrow Vine \}, \\ \{ Seedless. Grape. Vine \rightarrow Grape. Vine \rightarrow Vine \rightarrow Arbor. Plant \}, \\ \{ Seedless. Grape. Vine \rightarrow Grape. Vine \rightarrow Vine \rightarrow Arbor. Plant \rightarrow Plant \}, \\ \{ Seedless. Grape. Vine \rightarrow Seedless. Thing \}, \\ \{ Seedless. Grape. Vine \rightarrow Seedless. Thing \Rightarrow Fruit. Plant \} \} \end{split}$$

Thus, although the set of paths of the *sgv-net* { σ_1 , σ_2 , σ_3 , σ_4 , σ_5 , σ_6 , σ_7 , σ_8 , σ_9 , σ_{10} } cannot stand together as a consistent unit of information, we have two subsets associated with the network that can serve to such effect. In particular, and like we pointed out before, we have the following consistent units of information: $\Phi_1 = {\sigma_1, \sigma_2, \sigma_5, \sigma_6, \sigma_7, \sigma_8, \sigma_9}$ and $\Phi_2 = {\sigma_1, \sigma_2, \sigma_3, \sigma_4, \sigma_9, \sigma_{10}}$.

We now reach the point in which we have a network that involves multiple extensions. Thus, the extensions can be computed from either of our so far reviewed approaches. In the context of the *sgv-net*, the question is whether the various approaches to computing information within inheritance networks yield the expected results. In the following section, we review each approach separately.

5.2 Approaches

In this section, we go into further detail as to how the various information-processing strategies address the case study we presented in the previous section. We present the credulous approach, but more importantly and central to our research; we examine both versions of skeptical inference. The main goal is to address the technical aspects surrounding the phenomena of floating conclusions. We aim to discuss the problematic nature behind the different intuitions underlying each account of skeptical inference. Moreover, we aim to emphasise the subtle but significant differences between the two versions of skepticism.

5.2.1 The Credulous Solution

As we mentioned in previous sections, the credulous approach is an information-hungry strategy. This reasoning policy requires that we first identify all of the extensions associated with a given network, which are referred to as credulous extensions. Then, we are able to infer, according to the credulous approach, those attributes that are supported by any of the credulous extensions.

In the context of the sgv-net, and as we mentioned in the previous subsection, we have the following extensions:

$$\Phi_1 = \{ \{Seedless. Grape. Vine \rightarrow Grape. Vine \}, \\ \{Seedless. Grape. Vine \rightarrow Grape. Vine \rightarrow Vine \}, \\$$

 $\{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant\},\\ \{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant \Rightarrow Arbor.Plant\},\\ \{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant \rightarrow Tree\},\\ \{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant \rightarrow Tree \rightarrow Plant\},\\ \{Seedless.Grape.Vine \rightarrow Seedless.Thing\}\}$

$$\begin{split} \Phi_2 &= \{ \{ Seedless. Grape. Vine \rightarrow Grape. Vine \}, \\ \{ Seedless. Grape. Vine \rightarrow Grape. Vine \rightarrow Vine \}, \\ \{ Seedless. Grape. Vine \rightarrow Grape. Vine \rightarrow Vine \rightarrow Arbor. Plant \}, \\ \{ Seedless. Grape. Vine \rightarrow Grape. Vine \rightarrow Vine \rightarrow Arbor. Plant \rightarrow Plant \} \\ \{ Seedless. Grape. Vine \rightarrow Seedless. Thing \}, \\ \{ Seedless. Grape. Vine \rightarrow Seedless. Thing \Rightarrow Fruit. Plant \} \} \end{split}$$

Having already identified the proper extensions associated with the sgv-net, coupled with the fact that every extension of a given network is by definition a credulous extension, we end up with the credulous extensions of this network. The next step is to endorse any attribute supported by any of the credulous extensions.

Given this approach, we have that, following either Φ_1 or Φ_2 we have credulous grounds to endorse the *Plant* attribute. Thus, whichever extension we choose, we end up endorsing the *Plant* attribute to the *Seedless.Grape.Vine* node. Such endorsement is, as usually the case with credulous approaches, unproblematic. In such a vein, we can safely state that, according to the credulous reasoning policy, the sgv-net enables the *Seedless.Grape.Vine* node to inherit the *Plant* attribute.

As we have highlighted, the credulous account for processing information within inheritance networks is regarded, more often than not, as a highly unsatisfactory mechanism. In particular, when it comes to conflicting and inconsistent information, this information-hungry strategy seems to bypass the problem rather than address it. In such a vein, the credulous approach is not of interest for our purposes. Furthermore, and in line with the underlying thread within this chapter, the phenomenon of floating conclusions does not emerge from this form of processing information within inheritance networks. As such, throughout the remainder of our work we will not consider the credulous approach.

The above concludes the description of how the credulous approach handles cases like the sgv-

net. We devote the following subsections to examine how skeptical approaches address the same network. In the next subsection, we start with the indirect version of skeptical inference.

5.2.2 The Indirectly Skeptical Solution

From an indirectly skeptical approach, and unlike a credulous stance, there is a more rigorous and higher standard to satisfy before we can endorse a given attribute. In particular, indirect skepticism states that, as with the credulous approach, we first identify the credulous extensions. Once the credulous extensions are identified, we do not infer anything that is endorsed by some of the credulous extensions, but rather we endorse those attributes that are common to all credulous extensions, i.e. the *skeptically endorsed attributes*. This approach is the foundation of the indirectly skeptical mechanism that we will apply to the *sgv-net*.

As previously remarked, in the context of the sgv-net the problem revolves around whether the *Plant* attribute can be inferred. Following the indirect account of skeptical inference, the problem can be reduced to determine if such attribute is common to every extension of the network, i.e. whether the Plant attribute is a *skeptically endorsed statement*.

To address the *sgv-net* from an indirect approach, we comply with a three-step process. First, we identify all the credulous extensions. In the context of the *sgv-net* we have the following extensions:

$$\begin{split} \Phi_{1} &= \{\{Seedless.Grape.Vine \rightarrow Grape.Vine\}, \\ \{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine\}, \\ \{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant\}, \\ \{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant \rightarrow Arbor.Plant\}, \\ \{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant \rightarrow Tree\}, \\ \{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant \rightarrow Tree \rightarrow Plant\}, \\ \{Seedless.Grape.Vine \rightarrow Seedless.Thing\}\} \end{split}$$

$$\begin{split} \Phi_2 = \{ \{Seedless. Grape. Vine \rightarrow Grape. Vine \}, \\ \{Seedless. Grape. Vine \rightarrow Grape. Vine \rightarrow Vine \}, \\ \{Seedless. Grape. Vine \rightarrow Grape. Vine \rightarrow Vine \rightarrow Arbor. Plant \}, \end{split}$$

$$\{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine \rightarrow Arbor.Plant \rightarrow Plant\},\\ \{Seedless.Grape.Vine \rightarrow Seedless.Thing\},\\ \{Seedless.Grape.Vine \rightarrow Seedless.Thing \Rightarrow Fruit.Plant\}\}$$

Second, and after having identified the credulous extensions, we construct a set containing all of the attributes endorsed by each of the credulous extensions. To account for this, we make use of the consequence operation cn() that takes as input extensions and outputs the attributes endorsed by the extensions. Thus, we apply the cn() operation to the set of extensions associated with the network. In our case, we apply the operation as follows:

 $cn(\Phi_1, \Phi_2) = \{ \{Grape.Vine, Vine, Fruit.Plant, Tree, Plant, Seedless.Thing\}, \\ \{Grape.Vine, Vine, Arbor.Plant, Plant, Seedless.Thing\} \}$

The above is just an abbreviation of the following operation:

$$cn(\{ \{Seedless.Grape.Vine \rightarrow Grape.Vine\}, \\ \{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine\}, \\ \{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant\}, \\ \{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant \rightarrow Arbor.Plant\}, \\ \{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant \rightarrow Tree\}, \\ \{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant \rightarrow Tree \rightarrow Plant\}, \\ \{Seedless.Grape.Vine \rightarrow Seedless.Thing\}\}, \\ \{\{Seedless.Grape.Vine \rightarrow Grape.Vine\}, \\ \{Seedless.Grape.Vine \rightarrow Grape.Vine}, \\ \{Seedless.Grape.Vine}, \\ \{Seedle$$

Third, and once we have identified the credulously endorsed attributes (in our case the resulting set of the *cn()* operation is {{*Grape.Vine, Vine, Fruit.Plant, Tree, Plant, Seedless.Thing*}, {*Grape.Vine, Vine, Arbor.Plant, Plant, Seedless.Thing*,}), we apply the intersection operation to such set. Thus, the intersection operation is applied to the output of the *cn()* operation. As such, we have the following:

∩({{Grape.Vine, Vine, Fruit.Plant, Tree, Plant, Seedless.Thing}, {Grape.Vine, Vine, Arbor.Plant, Plant, Seedless.Thing}})

= {*Grape.Vine*, *Vine*, *Seedless.Thing*, *Plant*}

The previous operation yields the {*Grape.Vine*, *Vine*, *Seedless.Thing*, *Plant*} as a result. Thus, from an indirect approach, we have that the *Seedless.Grape.Vine* node inherits the *Plant* attribute.

The above can be summarised as follows. First, we take every credulous extension, which like we previously mentioned, are all of the extensions associated with the network, i.e. Φ_1 and Φ_2 . Then, and unlike the credulous approach, we do not infer an attribute just because it is supported by some credulous extension, but rather we endorse the attributes that are *in every one of the credulous extensions of the network*, which we refer to as the *skeptically endorsed attributes*. In our case, we have that *Plant* is an attribute that is contained in every single one of the network, we are justified in endorsing such conclusion. That is, we endorse the *Plant* attribute to the *Seedless.Grape.Vine* node. As with the credulous form of computing information, we support the *Plant* attribute to the same starting node. Nevertheless, and contrary to the credulous approach, there seems to be a more sensible justification for this endorsement. In particular, the conclusion is grounded not on the fact that such attribute is contained in *some* of the available extensions, but on the fact that such attribute is contained in *some* of the available extensions, but on the fact that such attribute is contained in *some* of the available extensions, but

The above concludes our description of how the indirect version of skeptical inference addresses the sgv-net. In the next subsection, we examine how the direct version of skeptical inference operates on the same network.

5.2.3 The Directly Skeptical Solution

Like the indirect counterpart, direct skepticism aims to provide a better standard to endorse information than the credulous approach. According to this account of skeptical inference, we first identify the credulous extensions associated with the network. Then we identify the paths that are common to all of the involved extensions. Finally, the attributes that can be inferred from the network are those attributes that are supported by the paths that are common to all of the involved extensions, i.e. the paths belonging to the *skeptically endorsed extension*. This is the strategy adopted by the direct version of skeptical inference that we immediately proceed to apply to the *sgv-net*.

Like in the previous example, the problem remains whether the *Plant* attribute can be inferred from the network. Nevertheless, following the direct version of skeptical inference, the problem is now reduced to determine if the attribute is contained in a path that is common to all credulous extensions of the network, i.e. if the *Plant* attribute belongs to a path contained in the *skeptically endorsed extension*.

To address the *sgv-net* according to a direct approach, we comply with a three-step process. First, we identify all the credulous extensions. In the context of the *sgv-net* we have the following extensions:

$$\begin{split} \Phi_{1} &= \{\{Seedless.Grape.Vine \rightarrow Grape.Vine\}, \\ &\{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine\}, \\ &\{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant\}, \\ &\{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant \rightarrow Arbor.Plant\}, \\ &\{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant \rightarrow Tree\}, \\ &\{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant \rightarrow Tree \rightarrow Plant\}, \\ &\{Seedless.Grape.Vine \rightarrow Seedless.Thing\}\} \end{split}$$

$$\begin{split} \Phi_{2} &= \{\{Seedless.Grape.Vine \rightarrow Grape.Vine\}, \\ &\{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine\}, \\ &\{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine \rightarrow Arbor.Plant\}, \\ &\{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine \rightarrow Arbor.Plant \rightarrow Plant\}, \\ &\{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine \rightarrow Arbor.Plant \rightarrow Plant\}, \\ &\{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine \rightarrow Arbor.Plant \rightarrow Plant\}, \\ &\{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine \rightarrow Arbor.Plant \rightarrow Plant\}, \\ &\{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine \rightarrow Arbor.Plant \rightarrow Plant\}, \\ &\{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine \rightarrow Vine \rightarrow Arbor.Plant\}, \\ &\{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine \rightarrow Vine \rightarrow Arbor.Plant\}, \\ &\{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine \rightarrow Vine \rightarrow Arbor.Plant\}, \\ &\{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine \rightarrow Vine \rightarrow Arbor.Plant\}, \\ &\{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine \rightarrow Vine \rightarrow Arbor.Plant\}, \\ &\{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine \rightarrow Vine \rightarrow Arbor.Plant\}, \\ &\{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine \rightarrow Vine \rightarrow Arbor.Plant\}, \\ &\{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine \rightarrow$$

{Seedless.Grape.Vine
$$\rightarrow$$
 Seedless.Thing},
{Seedless.Grape.Vine \rightarrow Seedless.Thing \rightarrow Fruit.Plant }}

Second, and unlike the indirect approach, which extracts the endorsed statements from the credulous extensions, here we intersect the set containing the extensions to determine the *skeptically endorsed extension* as follows:

$$\bigcap(\Phi_1, \Phi_2) = \{ Seedless. Grape. Vine \rightarrow Grape. Vine, \\ Seedless. Grape. Vine \rightarrow Grape. Vine \rightarrow Vine, \\ Seedless. Grape. Vine \rightarrow Seedless. Thing \}$$

The above is just a shorter version of the following operation:

 $\bigcap(\{\{Seedless.Grape.Vine \rightarrow Grape.Vine\}, \\ \{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine\}, \\ \{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant\}, \\ \{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant \rightarrow Arbor.Plant\}, \\ \{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant \rightarrow Tree\}, \\ \{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Fruit.Plant \rightarrow Tree \rightarrow Plant\}, \\ \{Seedless.Grape.Vine \rightarrow Seedless.Thing\}\},$

 $\{\{Seedless.Grape.Vine \rightarrow Grape.Vine\},\\ \{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine\},\\ \{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine \rightarrow Arbor.Plant\},\\ \{Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine \rightarrow Arbor.Plant \rightarrow Plant\},\\ \{Seedless.Grape.Vine \rightarrow Seedless.Thing\},\\ \{Seedless.Grape.Vine \rightarrow Seedless.Thing \Rightarrow Fruit.Plant\}\}$

= {Seedless.Grape.Vine → Grape.Vine, Seedless.Grape.Vine → Grape.Vine → Vine, Seedless.Grape.Vine → Seedless.Thing}

The process of intersecting the credulous extensions involved in the sgv-net returns the following

set: {Seedless.Grape.Vine \rightarrow Grape.Vine, Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine, Seedless.Grape.Vine \rightarrow Seedless.Thing}. This result comes from the fact that the aforesaid paths are the only ones common to all extensions.

Finally, we apply the consequence operation cn() to the output of the $\cap()$ operation. Thus, unlike the indirect approach, here we take as input the result of the intersection of the credulous extensions rather than every credulous extension. To put it another way, the cn() operates on the *skeptically endorsed extension* rather than on the collection of credulous extensions. Thus, we have the following:

 $cn(\{Seedless.Grape.Vine \rightarrow Grape.Vine, \\Seedless.Grape.Vine \rightarrow Grape.Vine \rightarrow Vine, \\Seedless.Grape.Vine \rightarrow Seedless.Thing\})$

This operation yields the set {*Grape.Vine*, *Vine*, *Seedless.Thing*} as a result. Hence, from a direct approach, the *Seedless.Grape.Vine* node does not inherit the *Plant* attribute.

The above can be summarised as follows. First, and following the thread of the indirect approach, the direct account of skeptical inference takes as the basic input credulous extensions. As we have pointed out, in the context of the *sgv-net*, Φ_1 and Φ_2 are the credulous extensions. Then, unlike the indirect approach, in this version of skeptical inference, we do not look for *attributes* that are common to every credulous extension, but rather, we look for *paths* that are common to every credulous extension. In our case, we have that, *with respect to the* Plant *attribute*, there is no path whatsoever that is common to all of the credulous extensions. Thus, in the context of the sgv-net, there is no skeptical inference and with respect to the *Plant* attribute, we can infer nothing from the sgv-net. Hence, the direct account of skeptical inference suspends judgment in cases like the sgv-net. This result is significantly different from the indirect approach.

Having reviewed how direct skepticism addresses the *sgv-net*, we continue in the next and final section, with a critical assessment of the extent and limitations associated with the divergence of both approaches to skeptical inference.

5.3 Discussion

In this section, we analyse some of the debates associated with the diverging views of skeptical inference. We examine the problem of the twofold approach to the same intuition behind skeptical inference and how it relates to floating conclusions. Moreover, we review some of the consequences of the distinction between direct and indirect approaches to skeptical inference. The main objective of this section is to consider some of the preliminary discussions of a twofold approach to the idea of skeptical inference. We also set the stage for the debate that we address in the forthcoming chapter concerning, floating conclusions.

5.3.1 An Irreconcilable Mismatch

Skeptical approaches are based on the idea that the information that is to be endorsed from a given network with multiples extensions is that which is contained in the intersection of all extensions. Despite this basic intuition underlying the skeptical approach towards inheritance networks, there are two different ways in which the intersection of multiple extensions can be carried out. The existence of a twofold solution to the same foundational intuition underlying skeptical approaches opens a series of questions concerning the mechanisms behind each approach, the extent of their divergence and what does it mean for inheritance networks as representational structures.

In the context of the sgv-net, and following an indirect approach, no matter how we choose to resolve the conflicts, we end up somehow endorsing the *Plant* attribute. On the one hand, if we dismiss the *Fruit.Plant* node, but follow the *Vine* and *Arbor.Plant* nodes, we end up with the result that the *Seedless.Grape.Vine* node inherits the *Plant* attribute. On the other hand, if we follow the *Fruit.Plant* and *Tree* nodes, we also end up with the result that the *Seedless.Grape.Vine* inherits the *Plant* attribute (Stein, 1989). Contrary to the credulous account, the indirect approach seems to provide a better justification for this endorsement, since it is grounded not on the fact that the *Plant* attribute is contained in some of the available extensions, but on the fact that such attribute is contained in *every* extension provided by the network. Nevertheless, the directly skeptical approach offers a different result, given that there is no path common to *all* extensions. Contrary to indirect skepticism, the *Plant* attribute is supported by *every* extension, but only from *different* paths. Thus, despite having the same working assumption of inferring a conclusion only when it can be located in the intersection of all extensions, the skeptical approach produces different results

depending on *how* the operation of the intersection of all extensions is computed. The difficulty concerning the nature and extent of the divergence between both approaches was addressed by Schlechta, who concluded that they cannot be made equivalent (1993).

One of the first criticisms that are raised by networks that contain floating conclusions is that the directly skeptical approach fails to infer information that seems reasonable to endorse, even though it does not satisfy the technical restrictions posed by the directly skeptical procedure. As such, direct skepticism has been seen as an explicitly faulty mechanism, as pointed out by Maier & Nute:

The majority of researchers accept floating conclusions as reasonable and view direct skepticism's inability to derive floating conclusions as a defect. Makinson and Schlechta (1991) call the directly skeptical approach "just wrong". (2010, 247)

One of the fundamental objections to Horty's notion of a directly skeptical approach is that it fails to account for information that lies beyond a path defined as the intersection of all paths contained in the credulous extensions. This line of criticism challenges the directly skeptical approach to information within inheritance networks. In such a vein, Makinson & Schlechta points out that one of the problems with direct skepticism is the requirement for attributes to be located in what is regarded as a *skeptically acceptable path* (1991). Even if such path is attainable, there are cases in which an attribute seems reasonable to endorse under a highly skeptical analysis, but fails to be contained in a skeptically acceptable path, not because it is not skeptically reasonable to endorse such information.

One crucial point made by Makinson & Schlechta (1991) rests on the distinction between information endorsed by a *skeptically acceptable path* and information endorsed by a *skeptical criterion*. The former is what lies at the heart of direct skepticism, whereas the latter is what the authors esteem as being portrayed by the indirectly skeptical approach, i.e. being endorsed by every single extension of the inheritance structure.

So far we have provided a preliminary discussion over the inherent mismatch between the two approaches to the same intuition behind skeptical stances in the context of inheritance networks. In the next subsection, we continue to discuss one remaining important problem associated with this mismatch between the indirect and direct accounts of skeptical inference.

5.3.2 A Unified Notion of Skeptical Inference

Regarding the extent and limitations of what should count as a genuinely skeptical attitude for handling information, it has been considered that such a skeptical attitude has a different motivation. In particular, Horty explains:

As a policy for reasoning with conflicting defaults, the notion of skepticism was originally introduced into the field of nonmonotonic logic to characterize the particular system presented in [6], which did not involve the assumption that one of a theory's multiple extensions must be correct, and did not support floating conclusions. (2002, 71)

Thus, skeptical approaches refer to stances that do not assume that, when multiple lines of reasoning (i.e. extensions) are involved, one particular line of reasoning should be held as correct. In this vein, a skeptical stance would inherently imply the omission of floating conclusions. Nevertheless, Horty states that:

By now, however, the term is used almost uniformly to describe approaches that do rely on this assumption, so that the "skeptical conclusions" of a theory are generally identified as the statements supported by each of its multiple extensions, including the floating conclusions. (2002, 71)

As it stands, the notion of skepticism and skeptically adequate information, according to indirect skepticism, is deflated to include situations in which a particular piece of information is contained in every line of reasoning (i.e. extension) of the involved formalism (in our case and for our particular interest inheritance networks). The direct notion of skeptical inference would deter floating conclusions, whereas the indirect account of skepticism would take them as legitimate information.

Horty (2002) suggests that any form of skeptical inference, contrary to what is done through an indirectly skeptical stance, should account for the intersection of extensions themselves contrary to simply taking the intersection of the statements endorsed by the extensions. Furthermore, these

kinds of mechanisms to compute the information of a network should have a place in more nuanced forms of skeptical reasoning. In particular, Horty claims that:

[...] if we think of skepticism as the general policy of withholding judgment in the face of conflicting defaults, rather than arbitrarily favoring one default or another, there is a complex space of reasoning policies that could legitimately be described as skeptical, many of which involve focusing on the arguments that support particular conclusions, not just the conclusions themselves. (2002, 71)

Thus, Horty's stance goes beyond the mere idiosyncratic use of the term of skepticism. What lies in Horty's critique is tied to the foundational intuition behind the concept of skeptically acceptable information, and how such deeply rooted intuitions can shape the idea of skeptically acceptable information. Concerning this, Horty states that:

Of course, there is nothing wrong with this usage of the term, as a technical description of the statements supported by each extension---- except that it might tend to cut off avenues for research, suggesting that we now know exactly how to characterize the skeptical conclusions of a theory, so that the only issues remaining are matters concerning the efficient derivation of these conclusions. (2002, 71)

The above concludes the preliminary discussion of the extent and limitation of the possibility for a common and unified notion of skeptical inference. In particular, we described how floating conclusions represents an inescapable problem that impedes moving towards finding a common ground.

5.3.3 Floating Conclusions in the Context of Classical Logic

The problem of floating conclusions can be framed in terms of classical logic. More precisely, it has been raised the point that the disjunction elimination rule employed in classical logic displays a similar dynamic instantiated by floating conclusions in non-monotonic formalisms (Stein, 1989).

The disjunction elimination rule, also referred to as proof by cases, is an argument form which one can infer *R* if we have the premises $P \rightarrow R$, $Q \rightarrow R$ and $P \lor Q$. This rule from classical logic is based on the assumption that if both *P* or *Q* implies *R* and we also have that *P* or *Q*, then it's safe to deduce *R*. That is, if we have a disjunction and both disjuncts imply the same proposition, then we can deduce such proposition based on the fact that at least one disjunct holds.

Horty (2002) suggests that on the assumption that a given agent that believes that B implies A, and that C implies A, and that either B or C holds, following the disjunction elimination rule from classical logic, the agent can conclude that A holds. That is if one knows that the following holds:

(1.a)
$$B \rightarrow A$$

(2.a) $C \rightarrow A$
(3.a) $B \lor C$

Then one naturally endorses A.

The dynamics of the disjunction elimination rule clearly has a substantial resemblance with floating conclusions, since floating conclusions occur when different arguments or lines of reasoning converge on the same conclusion, regardless of which argument or line of reasoning one is inclined to believe. In this sense, we could try to rephrase the disjunction elimination rule in terms of defeasible logic as follows: given a default defeasible theory with two extensions, E_1 and E_2 , assume that the extension E_1 contains the statement A and that the extension E_2 also contains the statement A. That is, no matter which extension one chooses, one will end up with the statement A (Horty, 2002). More precisely, if we accommodate the notion of extensions and attributes in the same fashion by presumably holding the following:

(1.b) $E_1 \rightarrow A$ (2.b) $E_2 \rightarrow A$ (3.b) $E_1 \lor E_2$

Then, given 1.b, 2.b and, 3.b one should be allowed to endorse A.

Horty points out that "The standard view is that a skeptical reasoner should then draw A as a conclusion, even if it is not supported by a common argument in the two extensions" (2002, 70).

Here, what Horty refers as the standard view is indirect skepticism, which supports floating conclusions. Nevertheless, he raises a substantial point as to why the scenario in classical logic is not entirely analogous to what happens in defeasible logic. More precisely, he points out that:

Notice the difference between these two cases, though. In the first case, the classical reasoning agent believes both that *B* and *C* individually imply *A*, and also that either *B* or *C* holds. In the second case, we might as well suppose that the skeptical reasoner knows that *A* belongs to both the extensions E_1 and E_2 , so that both E_1 and E_2 individually imply *A*. (Horty, 2002, 71)

Nevertheless, Horty (2002) points out that there is a distinction to be made from such analogy.

The reasoner is therefore justified in drawing A as a conclusion by something like the principle of constructive dilemma---as long as it is reasonable to suppose, in addition, that either E_1 or E_2 is correct. This is the crucial assumption, which underlies the standard view of skeptical reasoning and the acceptance of floating conclusions. But is this assumption required? (71, italics added)

Thus, Horty suggests that the whole case for accepting floating conclusions rests on the presumption of 3.a being the same made in 3.b, which clearly needs not to be the case, and there is no argument as to why it should be the case. This highlights one of the fundamental problems concerning floating conclusions, which is arguing in favour of the (indirectly skeptical) view that the same conclusion follows *no matter which extension one chooses*, and this is what should be relevant when accepting such conclusions. Nevertheless, this overlooks the very credible arguments raised against floating conclusions. Therefore, the problem of floating conclusions framed in terms of the disjunction elimination rule of classical logic only displaces the problem, rather than solving the dispute. In particular, Horty argues that:

Suppose that each of the theory's multiple extensions is endorsed by some credulous reasoner. Then the assumption that one of the theory's extensions must be correct is equivalent to the assumption that one of these credulous reasoners is right. But why should a skeptical reasoner assume that some credulous reasoner, following an entirely different reasoning policy, must be right? Of course, there may be situations in which it is appropriate for a skeptical reasoner to adopt this standard view [...]. But there also seem to be situations in which a deeper form of skepticism is appropriate—where each of the multiple extensions is undermined by another to such an extent that it seems like a real possibility that all of the credulous reasoners could be wrong (2002, 71).

Therefore, and following the licentious way in which we have been talking about classical and defeasible logic, we could state that the disjunction elimination rule of classical logic resembles an indirectly skeptical approach. Nevertheless, this does little to adequately explain *why* we ought to operate as such. That is, the move of reframing the problem of floating conclusions dismisses the problem entirely by assuming that the indirectly skeptical dynamic is the best approach, rather than making clear why is it that such strategy is better equipped in the context of defeasible reasoning.

5.4 Chapter Summary

In this chapter, we examined the counterintuitive results that emerge from the existence of a twofold approach to skeptical inheritance, namely, the problem of floating conclusions. We discussed the nature of the opposing stances towards floating conclusions. Lastly, we examined how floating conclusions represent an inescapable problem to the existence of a unified notion of skeptical inference. There are various challenging and intriguing theoretical questions associated with both versions of skeptical inference and its corresponding problem of floating conclusions. Such problems open further questions of methodological and practical nature, which compels us to examine various commitments towards inheritance networks as representational frameworks for defeasible reasoning. Nevertheless, a thorough and comprehensive discussion of the aforesaid was not addressed in this chapter. In the next chapter we inquire into the debate associated to the existence of different versions of skeptical inference and their relation to the problem of floating conclusions, and consequently each version of skeptical inference.

Chapter 6 The Debate over Floating Conclusions

In this chapter, we look into the problem of floating conclusions. In particular, we move from the technical approach we took in the previous chapter and focus entirely on the epistemological dimension of the problem. Furthermore, we examine the arguments that have been put for and against the acceptance of floating conclusions. We analyse how these arguments not only provide support for each particular stance towards floating conclusions but also for their corresponding version of skeptical inference.

The purpose of this chapter, unlike the previous, is to provide a detailed and critical assessment of the intuitions and reasons provided for each version of skeptical inference. Through the examination of the main cases in support of each version of skepticism, we can assess the grounding reasons for floating conclusions, and consequently, the corresponding version of skeptical inference.

6.1 Defending Indirect Skepticism

In this section, we analyse indirect skepticism as an account of skeptical inference. Furthermore, we examine the motivation for this approach to inheritance networks. We consider three cases that are used in support of this notion of skeptical inference. The aim of this section is to provide a thorough description and analysis of the cases used in support of indirect skepticism. The aforesaid will serve as an important input for the last chapters, in which we construct our stance towards floating conclusions.

6.1.1 Preliminaries

In this subsection, we discuss the overall strategy and cases that we examine with respect to indirect skepticism and the problem of floating conclusions. As we mentioned earlier, we have three specific cases to assess in support of indirect skepticism. First, we consider a case suggested by Ginsberg (1987) and further discussed by Horty (2012), which is the first case that suggested

and problematised the phenomenon of floating conclusions and is used by Horty to support indirect skepticism. Second, we inquire into a case proposed by Prakken (2002), which uses the same reasoning pattern, and involves attributing a leisure preference to a given individual with a particular background and social attributes. Third, we assess the example suggested by Stein (1992), which we have addressed to some extent in the previous chapter, but at this stage is undertaken with more of an epistemological concern, as opposed to the technical attention it received in the previous chapter.

The overall strategy to address each case used in support of indirect skepticism roughly follows a four-point structure of analysis. First, we provide a detailed exposition of the case study as it was presented in the literature. In such a vein, we discuss an extended description of the overall assumptions that were made or further discussed, as noted in the literature. At this stage, the work strictly recapitulates rather than pursue a critical assessment. Second, and after having provided all the details surrounding the case, we present the conclusions drawn. That is, we follow the information that was extracted as it is put forward in the literature. In this stage, we follow the original author's thread and no alternative lines of inquiry are considered. Third, we provide a critical assessment of both, the case as well as the conclusions extracted by the author. Here, we undertake a slightly more critical assessment than in the previous phases. Nevertheless, we do not go into deep detail with such critical assessment, as that is later addressed. Finally, and based on the previous aspects of the analysis, we examine how each specific case and the authors' approach to the case relate to the problem of floating conclusions and the indirect stance twoards skeptical reasoning.

It is important to mention that we make a significant departure from the technical analysis we have made in the previous chapter. The reason behind the aforementioned is that having presented the technical details surrounding certain cases such as the one discussed in the previous chapter, we now want to make a concerted effort to focus on the epistemological aspects of the examples, rather than the technical details surrounding each case. In the following subsections, we proceed to address each of the aforementioned cases having in mind the strategy set out in this subsection. Following the above, we aim to provide a thorough discussion for all of the cases addressed.

6.1.2 Ginsberg Diamond

The case proposed by Matthew Ginsberg in 1987, which by all accounts is the first case study that noted what would be later on known as floating conclusions. The scenario is a simple variation of the Nixon Diamond relies on the elementary tenets that, if a given individual is, in the context of international relations what is known as a *dove* then, all things considered, that same individual cannot be a *hawk*. Conversely, if a given individual is known to be, in the context of international relations as a *hawk*, he cannot be, all things considered, a *dove*. Based on the previous information, Ginsberg presents a case that involves two lines of reasoning. On the one hand, we assume that a given individual *Nixon* is a *Quaker*. Furthermore, *Quakers* are known to be *doves*. In the same vein, we also come to know that by a set of political standards, *doves* are held as *extremists*. On the other hand, we have that *Nixon* is a *Republican*. Furthermore, *Republicans* are known to be *hawks*. Likewise, by the same set of political standards, *hawks* are also regarded *extremists*. In addition to the previous information, we assume that one cannot be both a *hawk* and a *dove* at the same time.

In light of the previous case, *Nixon* is said to be a *Republican*, which leads him to be, all things considered, a *hawk* and consequently an *extremist*. Nevertheless, based on a different stream of information, we also have that *Nixon*, being a *Quaker*, which are known to be *doves*, leads to the classification of *Nixon* as an *extremist*. Thus, regardless of the line of reasoning we decide to take we end up with the conclusion that *Nixon* can be classified as an *extremist*.

According to this case, *Nixon* is both a *dove* and a *hawk*, which renders the case as troublesome. Nevertheless, we can also note that regardless of how we decide to resolve *Nixon*'s status as a *hawk* or a *dove*, it seems that we will end up endorsing that *Nixon* is an *extremist*. As we previously stated, no matter what route we end up taking, it ultimately leads to ascribing *Nixon* the attribute of being an *extremist*, either in virtue of him being a *hawk* or in virtue of him being a *dove*.

The aforementioned state of affairs seems to support that the floating conclusion *extremist* is a justified piece of information we could endorse, even by skeptical standards. More precisely, no matter what trait we decide to ascribe to *Nixon* (either *hawk* or *dove*) any of the involved alternatives will lead us to the conclusion that he is an *extremist*, all of which in turn supports the indirect stance towards skeptical inferences.

According to Horty (2012) this example illustrates that the sole absence of a common route towards a given conclusion seems to be an insufficient reason not to adopt information that can be derived from all the lines of reasoning.

6.1.3 Dutch Norwegian Iceskating

Here we consider one of Prakken's (2002) case studies used in support of indirect skepticism. Is important to point out that we adhere to the original version of Prakken's example and restrain from modifying it in pursuit of accuracy or clarity, as we emphasise the argument as it was originally conceived.

The case revolves around a given individual named Brygt Rykkje, to whom Prakken (2002)³² ascribes the following four basic claims:

- A- : Brygt Rykkje is Dutch since he was born in Holland.
- B- : Brygt Rykkje is Norwegian since he has a Norwegian name.
- A : Brygt Rykkje likes ice skating since he is Dutch.
- B : Brygt Rykkje likes ice skating since he is Norwegian.

According to this example, we first have that Brygt Rykkje is *Dutch* since he was *born in Holland*. Second, we have that Brygt Rykkje is *Norwegian* since he has a *Norwegian name*. In line with the previous claims, we have two further stances concerning Brygt Rykkje's preferences. On the one hand, given that Brygt Rykkje is *Dutch*, we have that he *likes ice-skating*. On the other hand, given that he is *Norwegian*, we also have that he *likes ice-skating*. Furthermore, we know that a given individual cannot be *Dutch* and *Norwegian* at the same time (following a sense of what it means to be *Norwegian* and *Dutch* such that it precludes the possibility that both can coexist).

Based on the previous example, Prakken (2002) notes that, like in the Ginsberg case, we have two streams of information. If we follow the thread of Brygt Rykkje being *Dutch* due to him being

³² Whichever way [the conflict] is decided, we always end up with an argument for the conclusion that Brygt Rykkje likes ice skating, so it seems that this conclusion is justified, even though it is not supported by a justified argument. In other words, the status of this conclusion floats on the status of the arguments [...]. (Prakken, 2002)
born in Holland, then we end up with the line of thought that Brygt Rykkje *likes ice-skating* because he is *Dutch*. Nevertheless, if we follow the line of reasoning that Brygt Rykkje is *Norwegian* because he *has a Norwegian name*, then we end up endorsing that Brygt Rykkje *likes ice-skating* due to him being *Norwegian*.

In the aforementioned state of affairs, the same less than comfortable situation emerges as it did with the Ginsberg Diamond, since both lines of reasoning conflict with one another. That is, the line of thought that ascribes to Brygt Rykkje the property of being *Dutch* conflicts with the line of thought that classifies Brygt Rykkje as *Norwegian*. Thus, both sources of information cannot coexist in a non-conflicting fashion. Beyond this troublesome fact, and however we resolve the conflict, we end up classifying Brygt Rykkje as a person who likes ice-skating. As Prakken observes:

Whichever way [the conflict] is decided, we always end up with an argument for the conclusion that Brygt Rykkje likes ice skating, so it seems that this conclusion is justified, even though it is not supported by a justified argument. In other words, the status of this conclusion floats on the status of the arguments [...]. (2002)

In such context, we can confidently state that regardless how we resolve the conflict, we can rest assured that it provides enough reasons to endorse the preference of ice skating. Thus, the problem seems to lie not in the confidence we could have for attributing the property of liking ice-skating to Brygt Rykkje, but rather in his origin. Nevertheless, as we have stressed, all the involved alternatives provide us with grounds to endorse his preference for ice-skating.

This example is used to support the indirect approach to skeptical inference since it points to a situation in which there seem to be better reasons for admitting floating conclusions rather than rejecting them. In particular, one would be hard-pressed to come up with grounds as to why endorsing the floating conclusion amounts to not having a skeptical enough attitude, which is what is at stake in the discussion. In such vein, it is important to note that either form of skepticism can be very easily mistaken with respect to the floating conclusion, but in this situation, the claim that the indirect approach is not skeptical enough seems not to stand. On the one hand, the involved sources conflict with one another, which is a problematic situation in and of itself. On the other hand, we do not have a commonly shared argument that supports the conclusion. The former

concern, as we have previously assessed, is not a good enough reason as to why endorsing the conclusion is not sufficiently skeptical. The latter, claiming the absence of a commonly shared path in support for the floating conclusion is a circular critique, as this is what needs to be proven. More precisely, a proper assessment would find no reason to claim that accepting the floating conclusion is not sufficiently skeptical and using the absence of a commonly shared path seems to set in place a requirement that appears to hold no grounds in the context of this example. In such vein, this case seems to provide grounds in favour of the indirect account of skeptical inference.

6.1.4 Seedless Grape Vine

Here, we analyse the third and final case used to support indirect skepticism and the acceptance of floating conclusions. This case was proposed by Stein (1992), and it was our base case in the previous chapter. Nevertheless, in this subsection, it receives a different treatment than when it was previously addressed.

As we mentioned in the previous chapter, the case proposed by Stein (1992) revolves around the *Seedless.Grape.Vine* entity, for which we have two threads of information. First, we have that such an entity is a *Grape.Vine*, and a *Vine*. From the fact that it can be classified as a *Vine*, we come to know that can be classified as an *Arbor.Plant*. Since *Seedless.Grape.Vine* can be classified as an *Arbor.Plant*. Second, we have that a *Seedless.Grape.Vine* is a *Grape.Vine*. Furthermore, we come to know that a *Grape.Vine* falls in the category of a *Fruit.Plant*, and consequently falls in the category of being a *Plant*. Third, and of vital importance for this case, we also come to know that a *Fruits.Plant* is not to be classified as an *Arbor.Plant*. That is, being a *Fruit.Plant* conflicts with being an *Arbor.Plant*, rendering the two categories as incompatible. Thus, we come to the problem of deciding whether *Seedless.Grape.Vine* is a *Plant* or not.

In such context, Stein (1992) suggests that, if we follow the thread that classifies *Seedless.Grape.Vine* as an *Arbor.Plant*, we end up rendering *Seedless.Grape.Vine* as a *Plant*. On the contrary, if we follow the thread that classifies *Seedless.Grape.Vine* as a *Fruit.Plant*, then *Seedless.Grape.Vine* falls into the *Tree* category and, consequently into the *Plant* category. In the same vein, we have the undesirable situation in which *Seedless.Grape.Vine* cannot be an *Arbor.Plant* and a *Fruit.Plant* at the same time, because such categories are incompatible.

Nevertheless, it is suggested that a *Seedless.Grape.Vine* should be classified as a *Plant*. As it happened with the Ginsberg Diamond and Prakken's discussion of Brygt Rykkje preferences, regardless of how we decide to resolve the conflict of *Seedless.Grape.Vine* being and *Arbor.Plant* or a *Fruit.Plant*, what seems clear is that we end up categorising *Seedless.Grape.Vine* as a *Plant*. If we opt to classify *Seedless.Grape.Vine* as an *Arbor.Plant* we certainly classify it as a *Plant*, and if we determine that *Seedless.Grape.Vine* is better classified as a *Fruit.Plant*, we equally end up claiming that *Seedless.Grape.Vine* is a Plant. As Stein (1991) observes:

The difficulty lies in the fact that some conclusions may be true in every credulous extension, but supported by different [argument] paths in each. Any path-based theory must either accept one of these paths---and be unsound, since such a path is not in every extension---or reject all such paths---and with them the ideally skeptical conclusion---and be incomplete [13, p. 284]. (284)

Thus, as Stein suggests, the main reason for direct skepticism failing to endorse such information for networks involving floating conclusions, is that it has to indulge in one of two different undesirable decisions: (a) take any of the available extensions as skeptically acceptable or (b) take no extension as skeptically acceptable. If one takes the route given by (a) then the whole skeptical approach collapses into a credulous approach, which is to be avoided, or if on the other hand, one takes the solution contained in (b) then one can be accused of not accounting for information that seems to be justifiably inferred. In such a vein, this case is used to support the indirect account of skeptical inference inasmuch as it instantiates a situation in which there seems to be a better case for accepting the floating conclusion than rejecting it. Thus, the detailed scrutiny of the case fails to provide a conclusive reason as to why the floating conclusion should be dismissed, other than conflict among the sources being present and the absence of an argument common to all lines of reasoning. That is, the sole existence of a conflict among the involved sources of information or the lack of an argument common to all extensions, seem not to provide reasons as to why endorsing the conclusion would not be sufficiently skeptical.

6.2 Defending Direct Skepticism

In this section, we analyse the arguments in support of the direct approach to skeptical information processing within inheritance networks. In particular, we discuss the motivation for this version of skeptical inference and analyse three cases supporting this approach. The aim of this section is to review cases and arguments used to support the direct version of skeptical inference. As we mentioned in the previous section, the material addressed in this section and the overall chapter is crucial for the last chapters (chapter 7 and 8) in which we develop our stance towards the debate of floating conclusions.

6.2.1 Preliminaries

In this subsection, we provide a general overview of cases considered in support of direct skepticism. Furthermore, we sketch the general strategy we follow to analyse each case in relation to the problem of floating conclusions.

Like in the previous section, here we have three specific cases in support of direct skepticism. First, we consider a case proposed by Horty (2002), in which we are to decide if the commander of a given military unit ought to launch an attack on a seemingly defensive enemy because different spies reach the same conclusion but through alternative and mutually undermining reasons. Second, we examine another example by Horty (2002). This case presents the problem of deciding if a specific individual has good reasons to make a significant financial investment based on the possibility of receiving a substantial inheritance. The caveat is that the sources that inform him of the possibility of receiving the inheritance money have alternative and opposing reasons. Third, we look into a final example, also presented by Horty (2002), in which a public policy needs to be enacted to address an imminent economic downturn, but the reasons for such prognostic are opposing and mutually undermining reasons.

Like in the previous section, the general strategy that we employ to address each of the cases used in support of direct skepticism follows a four-component plan. First, we offer an exhaustive description of the case under scrutiny as it was put forward in the literature. To address this point, we present not only the details concerning the case but also any relevant and auxiliary assumptions surrounding the example. This component adheres to a description of the case under scrutiny, and thus, only an expository style is followed. Second, and once all of the aforementioned aspects have been made sufficiently explicit, we provide a detailed description of the inferences and conclusions that are extracted from the case. In such a vein, it is important to point out that we present not only the author's example but also the reasons each case provides for accepting the involved floating conclusion. Third, we provide an assessment of the previous two components. That is, we contrast the acceptance of the conclusion from each case study with the author's original view towards floating conclusions. Again, we adhere as much as we can to the author's initial intuition to analyse the case, and as such, we do not seek to contest the original assessment. Having stated the above and following the same line of scrutiny as in the previous section, we do not exhaust the assessment of each case; this is the object of forthcoming chapters. Finally, after the case has been fully addressed, we assess how the example and the lines of reasoning relate not only to the problem of floating conclusions but more generally to the account of direct skepticism.

As we did with previous cases, in this section, we do not consider the technical details surrounding each example. To the contrary, what occupies our attention at this stage is the epistemological nuances regarding each case. This does not mean that the cases hereby analysed cannot comply with a formal representation through inheritance networks. Nevertheless, for the purposes of our research, such formal representation does not have a central role.

6.2.2 Military Attack

In this subsection, we assess one of the various examples proposed by Horty (2002) in support of the direct skepticism. This case involves making a strategic military decision of attacking an enemy's position based on the independent reports of two military scouts.

The example assumes that a given individual, which we call Ben, is the commander of a military unit. Ben, according to this example, is following the trail of the enemy, which, as it stands, holds a strong defensive position. Launching an attack would be lethally counterproductive due to the defensive position the enemy holds. Nevertheless, the overall goal of the unit, which is being constantly pressed by their superiors, is to proceed with an attacking strategy as soon as possible. Thus, Ben sends two of his trusted spies, Henry and Mark, who are held as reliable sources, to investigate the enemy's defensive positions. After some time passes, Ben independently hears from both of his spies. On the one hand, Henry reports to Ben that the enemy's defensive position is very weak as the main body of the enemy forces can be seen retreating to the mountains and that a diversionary group was sent out as well to make it seem like they were retreating along the river. On the other hand, Mark reports to Ben that indeed, the enemy's defensive position is very weak, but unlike Henry, Mark reports that the main body of the enemy forces can be seen retreating along the river and that a diversionary group was sent out as well to make it seem like they were retreating to the mountains.

The question is whether Ben should assume that, despite whichever spy he ends up trusting, the enemy's position is weak no matter what or if he should withhold his attack. If we are to believe Henry, then the enemy is weak as it is retreating (via the mountain). If we are to believe Mark, the enemy is equally retreating (via the river). No matter whom we believe, we reach the conclusion that the enemy is retreating. Like in the cases provided in support for indirect skepticism, here we have an instance of a conflict that no matter how it is resolved, we can rest assured it render us the same conclusion.

Nevertheless, Horty (2002) points out that, despite both sources reaching the same conclusion concerning the enemy's defensive position, they do so from different and possibly conflicting or at least mutually faulty-prone lines of reasoning. In particular, Horty observes that:

Based on this information, should you assume at least that the main enemy force has retreated from the defensive position---a floating conclusion that is supported by both spies---and therefore commit your troops to an attack? Not necessarily. Although they support a common conclusion, each spy undermines the support provided by the other. Perhaps the enemy sent out two diversionary groups, one through the mountains and one along the river, and managed to fool both your spies into believing that a retreat was in progress. Perhaps the main force still occupies the strong defensive position, awaiting your attack. (2002, 69)

Thus, Henry's account renders Mark's claim that the enemy is retreating via the river as mistaken, since this is only a diversionary group sent out to give the impression that they are retreating. Nevertheless, Mark's account renders Henry's claim that the enemy is retreating via the mountain as mistaken, since this is only a diversionary group sent to give the impression that they are conceding their position. Thus, it is equally plausible that even though both routes lead to the same

result, they should be thought as mutually undermining. That is, one cannot, or at least should not accept the claim that the enemy is retreating simply because, no matter how the conflict is resolved, we end up arriving at the conclusion that the enemy is retreating. Horty (2002) questions if one can skeptically infer from the available information that the enemy has retreated from the defensive position, because regardless which of the sources of information we choose, they both lead us to the same result, i.e. that the enemy has retreated from the defensive position.

This case is used in support of a direct stance towards skeptical reasoning. In particular, it highlights that floating conclusions should not be accepted but rather rejected. This case, unlike the previous cases supporting indirect skepticism, does seem to provide us with very good reasons as to why, accepting the floating conclusion might not be a prudent, let alone skeptical attitude. More precisely, we might conclude that no matter how we resolve the conflict, we end up endorsing that the enemy is retreating. Nevertheless, in this particular situation, such a claim seems to be more liberal than a skeptical one. The risk involved along with the error-prone sources of information given a highly sophisticated enemy can render the conclusion, a very likely false one. This first case study in support of direct skepticism highlight the circumstances in which the floating conclusion seems better off being dismissed than accepted, *from a skeptical point of view*. The aforementioned is not to say that the floating conclusion is more fallible, but rather that the surrounding conditions should make us adopt a more conservative attitude.

6.2.3 Inheritance Investment

Here, we consider another example used by Horty (2002) to argue in favour of the direct account of skeptical inference. The case involves making an investment decision based on the independent reports of two siblings.

Horty presents the case in which a given individual Adam, is faced with the tragic situation that both of his parents have a fatal disease, which leaves them with only one month to live. Both parents have a total net worth of a million dollars (each one having half a million dollars). Furthermore, Adam wants to purchase an expensive item, a yacht, which is set to be sold within a month. The price, although expensive by any objective measure is an attractive one for a short period of time. Nevertheless, the lowered price lasts until the end of the month. Thus, Adam finds himself in a financial predicament. On the one hand, Adam can seek to secure the deal on the expensive item by putting a large deposit. On the other hand, such a large deposit can be lost unless he gets at least one of the parents' inheritance to address the remaining of the balance.

To make up his mind, Adam seeks information from his brother and a sister, both already have talked to their parents regarding the matter of the inheritance. Adam reaches out to them to know if they have any information as to whether he will inherit any money whatsoever. On the one hand, Adam's brother tells him that the father will leave his part of the money to him (Adam's brother) while the mother will leave her part of the inheritance money to Adam, rendering the investment a good choice. On the other hand, Adam's sister tells him that the father will leave his part of the inheritance money to Adam, rendering the investment a good choice. On the other hand, Adam's sister tells him that the mother will leave its part of the money to her (Adam's sister) while the father will leave his part of the inheritance money to Adam, making the investment a good choice as well. The question is whether Adam should conclude that he will inherit half a million dollars and proceed to make the deposit on the yacht.

One could potentially be inclined to consider, as we did in the previous cases when we were analysing either the Ginsberg Diamond or Stein's Seedless Grape Vine object, that regardless of how we resolve the conflicting lines of reasoning, we end up with the claim that some parent will leave half a million dollars to Adam. For example, if Adam's brother is right then although the father will not leave his part of the money to Adam, his mother will do. If Adam's sister is correct, then although the mother will not leave her part of the money to Adam, his father will. Thus, regardless of what happens, one can accept that Adam will get someone's half a million dollars, despite any existing conflict between his brother's and sister's account or how one should decide to resolve such conflict.

Nevertheless, Horty (2002) points out that Adam should not infer that he will end up inheriting half a million dollars. Horty puts forward the idea that Adam should not proceed with the deposit because there are very strong reasons to suspect the *floating* claim that he will inherit a large sum of money. In particular, Horty (2002) points out that, despite having considered Adam's brother and sister as reliable sources, given the defeasible nature of the example, it could be very well the case that both accounts undermine each other in such a way that renders the conclusion of Adam inheriting half a million dollars as false. For example, it could happen that, the father will leave the money to Adam's brother, and the mother will do the same with Adam's sister. Appealing to the fallible nature of the information, such state of affairs is not implausible, and to the contrary seems like a possible arrangement of the inheritance money division based on account of both sources.

This case is presented as a specific example in favour of the direct approach to skeptical inference, because it illustrates a situation in which a floating conclusion seems a risky basis on which to act as long as one is trying to take a skeptical stance towards decision-making. More precisely, this example has important attenuating circumstances that should restrain Adam from accepting the conclusion that someone will give him his or her share of the inheritance money. In particular, Adam's decision involves a financial decision that if wrong could have severe and undermining consequences for Adam's financial well-being. These surrounding circumstances for the core reasoning process does point out to a good reason as to why there seem to be grounds for rejecting the floating conclusion.

6.2.4 Macroeconomics

This subsection examines the last of Horty's (2002) arguments in favour of direct skepticism. In particular, we take the hypothetical example of shaping a public policy based on financial information provided by experts.

For this case, we think of a head of state who is attending a macroeconomics forum in a time of economic health which is characterised by low inflation and economic growth. Two groups of experts are making the forecast with the same model. On the one hand, a group of experts, which we call Group-A, claims that the bold economic growth rhythm will trigger a high inflation which will lead to an economic downturn, according to their economic model. On the other hand, another group of experts, which we call Group-B, claims that given the same strong economic growth, but with minor changes in the parameters used in the same model, we should expect that the low inflation rate will not only persist, but it will continue to decline. This will produce a situation of deflation that will lead to an economic downturn. The question is whether the head of state should set in motion a series of public policies to address the imminent economic downturn or should the forecast of an economic downturn be disregarded.

In this case, like with the military example, there is a situation in which two distinct lines of reasoning lead to the same conclusion. Here, if we concede reason to Group-A, then the agreed-upon strong economic scenario will trigger higher inflation which will undermine the economy later on. Nevertheless, if we are to grant reason to Group-B, the strong economic situation will keep a low inflation that eventually will lead to a situation of deflation, which will undermine the

economy. Thus, one could be tempted to take the stance that regardless of which group is right, we ultimately reach the conclusion that there is an economic downturn at the horizon. That is, the situation might give us sufficient reasons to accept the floating conclusion. Thus, public economic policies should be adjusted to this inescapable forecast.

Nevertheless, Horty (2002) questions whether one should accept the inescapability of an economic downturn based on the fact that either line of reasoning reaches such conclusion. In particular, Horty points out that, all things consider, the extreme predictions endorsed by each group could potentially have an attenuating effect towards the opposing forecast. For example, the high inflation rate, which will lead to the economic downturn, predicted by Group-A might be attenuated by the deflationary process predicted by Group-B, thus undermining the claim of an economic downturn, might be attenuated by the inflation rate predicted by Group-A, all of which would potentially undermine the prediction of an economic downturn. Thus, there is room for a middle ground position, which would erode the plausibility of an economic downturn at all.

In line with the above, these alternatives do not suggest that the floating conclusion is wrong because, in the context of defeasible reasoning, all conclusions are open to correction or retraction. Nevertheless, the case does highlight that, if we are said to hold a skeptical attitude towards defeasible reasoning in general, then the floating conclusion regarding the inescapable economic downturn should not be accepted. Thus, like in the previous two cases proposed by Horty, and unlike the cases we reviewed in defence of indirect skepticism, here we seem to find strong rather than weak reasons to reject the floating conclusion, and hold an attitude that aligns with the directly skeptical stance.

6.3 Overview of the Debate

In this final section, we provide a general overview of the debate over direct and indirect approaches to skeptical inferences. Moreover, we examine how such accounts stand in relation to the problem of floating conclusions. The aim of this section is to construct a uniform systematisation of each stance of the debate. This section provides the elementary material to develop our particular approach to the problem of floating conclusions, in the forthcoming chapters.

6.3.1 Indirect Skepticism

One of the underlying and tacitly agreed-upon assumptions of the critiques against direct skepticism and, in general, any skeptical approach that does not make room or is unable to endorse what we have identified as floating conclusions is that its inability to account for such inferences is intrinsically mistaken. In this vein Horty observes:

What lies behind these various criticisms, of course, is the widely-held assumption that the second, rather than the first, of our two skeptical alternatives is correct---that floating conclusions should be accepted, and that a system that fails to classify them among the consequences of a defeasible knowledge base is therefore in error. (2002, 62)

From the standpoint of the indirect approach to skeptical inference floating conclusions are skeptically legitimate information, and as such, they should be accounted. Nevertheless, concerning this underlying assumption, Maier & Nute states "Acceptance of floating conclusions is not universal, however. Defeasible logic rejects them, as do other proponents of direct skepticism [...]" (2010, 247).

The feature instantiated by networks containing floating conclusions is that, although such conclusions are not accounted from a skeptically acceptable path, they seem to satisfy the skeptical scrutiny. Such defence of indirect skepticism is based on the idea that the technical constraints imposed by the directly skeptical approach do not match our intuition of what seems to be skeptically acceptable information. Based on the above, it has been argued that any form of skepticism that proposes to endorse information appealing to a unique skeptically acceptable extension fails to be a sensible mechanism to compute information within inheritance networks (Horty, 2002)³³.

³³ "From this they likewise argue, not only that the particular theory developed in [6] is incorrect, but more generally, that any theory attempting to define the skeptically acceptable conclusions by reference to a single set of acceptable arguments will be mistaken." (Horty, 2002, 62)



Figure 6.3.1 The Floating Conclusions Arguments for Indirect Skepticism

Figure 6.3.1 summarises the cases that we have previously reviewed in support of indirect skepticism. More precisely, such cases supported the idea that endorsing floating conclusions seems to be skeptically legitimate in the sense that no matter the route one decides to take, one ends up reaching the conclusion in question. All cases presented in support for the indirect version of skeptical inference share the same pattern, which appeals to our common understanding that no matter how the involved conflict is resolved, the floating conclusions seems to hold. For example, the "skating case" appeals to our intuition that regardless of whether the person involved is Dutch or Norwegian, we can safely assume that such a person likes ice skating. The aforesaid conclusion, although it follows from different and conflicting lines of reasoning seems to hold no matter which route we decide to choose.

The above provides a general overview of the indirect account of skeptical inference and how the various cases stand in relation to this approach. In the next subsection, we undertake the same task but for the direct version of skeptical inference.

6.3.2 Direct Skepticism

As we have pointed out, the indirectly skeptical approach assumes as skeptically acceptable information conclusions that are contained in *every* extension. This stance operates on the assumption that, given multiple and divergent routes of reasoning, at least one of such involved lines of inquiry ought to be right, rendering floating conclusions as information sufficiently

skeptical. This is the point that is contested, as Horty explains:

Is it necessary for a skeptical reasoner to assume, when a theory leads to multiple extensions, that one of those extensions must be correct? Suppose that each of the theory's multiple extensions is endorsed by some credulous reasoner. Then the assumption that one of the theory's extensions must be correct is equivalent to the assumption that one of these credulous reasoners is right. (2002, 71)

Horty (2002) is one of the few critics of the acceptability of floating conclusions as sufficiently skeptical information and argues that such a view ought to be revised in light of other conditions. In particular, Horty brings into question such assumption as follows:

But why should a skeptical reasoner assume that some credulous reasoner, following an entirely different reasoning policy, must be right? Of course, there may be situations in which it is appropriate for a skeptical reasoner to adopt this standard view---that one of the various credulous reasoners must be right, but that it is simply unclear which one. That might be the extent of the skepticism involved. (2002, 71)

Thus, from a skeptical point of view, when a network has multiple extensions, one should not sustain that either of the extensions is skeptically justified. Horty's view is grounded on the fact that when multiple lines of reasoning are involved, assuming by default that at least one of them must be correct is not a genuinely skeptical attitude.

The problem when different lines of reasoning are involved resides on the possibility that such lines of reasoning could conflict with each other. Although this might not be problematic in certain situations, it fails to consider what Horty calls a "*deeper form of skepticism*". Horty explains the aforesaid as follows:

[...] there also seem to be situations in which a deeper form of skepticism is appropriate---where each of the multiple extensions is undermined by another to such an extent that it seems like a real possibility that all of the credulous reasoners could be wrong. The yacht,

spy, and economist examples illustrate situations that might call for this deeper form of skepticism. (2002, 71)

This deeper form of skepticism would adopt a less "licentious" policy, that may restrain itself from taking as skeptically acceptable information that is an end product of every line reasoning in a given context. This oversees the plausibility that such different lines of reasoning might conflict or undermine each other in a meaningful way. Despite the above, it is important to stress the fact that Horty's stance is not a simple case against floating conclusions as skeptically adequate information. Horty's take on the debate rests on the common intuition behind the general framework of defeasible reasoning that conclusions by definition have some degree of fallibility (the contrary view would collapse into a monotonic framework). To the contrary, what he points out is that, when a conclusion is only endorsed or inferred as a floating conclusion, it has a less secure justification than conclusions endorsed by a skeptically acceptable path. In particular, Horty explains this as follows:

The point is not that floating conclusions might be wrong; any conclusion drawn through defeasible reasoning might be wrong. The point is that a statement supported only as a floating conclusion seems to be less secure than the same statement when it is uniformly supported by a common argument. As long as there is this difference in principle, it is coherent to imagine a skeptical reasoner whose standards are calibrated so as to accept statements that receive uniform support, but to reject floating conclusions. (2002, 65)

Thus, Horty defends the idea that there should be room for situations in which floating conclusions should not be taken as skeptically acceptable information.



Figure 6.3.2 The Floating Conclusions Arguments for Direct Skepticism

In light with the above, Figure 6.3.2 summarises the previous discussion. In particular, it points to cases that have bifurcating lines of reasoning in which no matter what route one chooses to opt for; it ends up endorsing the same result. The aforementioned are cases in which the same conclusions is supported by different and conflicting lines of reasoning, with the caveat that such conflicting lines of reasoning do not share a common argument. The aforesaid evokes a very reasonable doubt as to whether the floating conclusion can be endorsed while still claiming to have a skeptical attitude towards information processing within this type of structures. For this, we could "[...] imagine a situation in which two sources of information, or reasons, support a common conclusion, but also undermine each other, and therefore undermine the support that each provides for the common conclusion." (Horty, 2002, 68). Such situations raise the question as to whether a conclusion being supported by every line of reasoning should be taken to be skeptically acceptable (at least from the indirectly skeptical approach).

In the same vein, Horty states that "[...] it is easy to construct other examples along similar lines: just imagine a situation in which two sources of information, or reasons, support a common conclusion, but also undermine each other, and therefore undermine the support that each provides for the common conclusion." (2002, 68). Adhering to this pattern one can construct a myriad of cases that show the risk implied by endorsing such information. In this sense, the underlying assumption that floating conclusions should be accepted seems unaccounted for, and sometimes even unquestioned. As Horty observes:

Why is it so widely thought that floating conclusions should be accepted

by a skeptical reasoner, so that a system that fails to generate these conclusions is therefore incorrect? It is hard to be sure, since this point of view is generally taken as an assumption, rather than argued for [...]. (2002, 70)

In such a vein, the argument that a skeptical approach to defeasible reasoning that fails to endorse this type of conclusion seems to be wrong, is brought into question.

6.3.3 Summary

The direct approach to skeptical inference completely rejects the admissibility of floating conclusions as skeptically legitimate information. In particular, this account of skepticism operates on the assumption that, in the absence of a common path with respect to all credulous extensions, a conclusion cannot be skeptically endorsed from a given network. Second, the indirect account of skeptical inference, accepts floating conclusions as skeptically legitimate information. This version of skepticism takes the stance that it is not a common and unique path which makes an attribute or conclusion skeptically acceptable, but rather our ability to locate that attribute or conclusion in every credulous extension of the network. The aforementioned can be synthesised in a very simple schema that portrays each of the cases reviewed and their relation to the direct and indirect version of skeptical inference.

	Examples
Direct Skepticism	Macroeconomics (Horty, 2002)
	Inheritance (Horty, 2002)
	Military (Horty, 2002)
Indirect Skepticism	Extremist (Ginsberg, 1987)
	Skating (Prakken, 2002)
	Plant (Stein, 1992)

Figure 6.3.3 The Floating Conclusions Debate

Figure 6.3.3 depicts what we have discussed so far. In particular, we have both accounts of skeptical inference along with the cases used in support of each one. Such matrix highlights two important aspects. First, it portrays the stance towards floating conclusions for each version of skeptical inference. Second, it provides a schematic depiction of the fact that there is no middle ground between either stance concerning the acceptance or rejection of floating conclusions. These are the building blocks that we are going to continuously extend and further distinguish with a nuanced approach we develop in the forthcoming chapters. This simple schema will later serve as input in the context of our inquiry.

6.4 Chapter Summary

In this chapter, we reviewed the cases used in favour of the direct and indirect account of skeptical inference. We assessed the main arguments that have been put in favour of both versions of skepticism. Moreover, we analysed how each case study and their treatment of floating conclusions stands in relation to support each version of skeptical inference. This chapter has followed the thread of inquiry initiated in the previous chapters. Nevertheless, we have only presented the problem, but have not proposed a specific solution. The question remains as to under what circumstances we ought to reject or accept floating conclusions, and under which circumstances, direct or indirect accounts of skeptical inference are to be followed.

In the next two chapters, we construct our particular response to the problem of floating conclusions and the direct/indirect skepticism debate. In the next chapter (chapter 7), we develop our framework of analysis. Later (chapter 8), we apply such a framework to the problem of floating conclusions to ground our stance towards the problem.

Chapter 7 The Source Conflict Cost Criterion

In this chapter, we construct a framework of analysis to account for the diverging views that underlie the problem of floating conclusions. The framework consists of two key components; one addresses the extent of the conflict in a situation with a floating conclusion and the other addresses the cost associated with the context involving such floating conclusion. The main objective of this chapter is to provide a detailed account of the theoretical framework used in the forthcoming chapter to analyse and address the problem of floating conclusions. In such a vein, most of the chapter is mainly a descriptive construction of this theoretical tool that we later put to use.

7.1 The Notion of Source

In this section, we consider the notion of a *source*. We discuss the extent and limitation of this concept for the purposes of our research and how it is shaped to address our specific needs. The following subsection aims to provide the first component of our twofold framework to address floating conclusions, that later on is merged with the second component to ground our framework of analysis.

7.1.1 Examples of Sources

The notion of a source is one of the key aspects that lies at the centre of our approach to understanding the problem of floating conclusions. In this subsection, we put forward an example that is going to help us consider the different caveats we wish to highlight while having a concrete case that illustrates how these caveats play a role in such situations. Although the notion of a source can be extended to cover other areas of epistemology, for the purposes of our research we restrain the scope to inheritance networks and its associated problems.

To illustrate the notion of a source, let us take a medical example. Let us say a given patient, which we will refer to as P is registered at a medical facility. In the same vein, let us say that P is seen

by three different practitioners, which we will denote as D_1 , D_2 and D_3 regarding a heart condition. Furthermore, let us say that the practitioners have the following prognoses.

- D_1 : *P* needs to take medication *Z* that contains low quantities of substance X_1 which are sufficient to improve his condition.
- D_2 : *P* needs to take medication *Z* that contains low quantities of substance X_1 which he should avoid, but does not present a risk in low quantities and large quantities of substance X_2 which he desperately needs to improve his condition.
- D_3 : *P* needs to take medication *Z* that contains large quantities of substance X_2 which are needed given his condition.

The above example highlights several aspects concerning the extent and limitation of the notion of sources. First, different sources can contain different amounts of information with respect to a given claim. Second, despite the different amounts of information provided by any given number of sources, we are interested in considering *sources with respect to a given claim*. Third, different sources can contain incompatible information. Such features do not exhaust by any means all of what can be discussed in our example. Nevertheless, it is also important that in our ongoing process of constructing our framework of analysis, we restrict and limit the extent of the aspects that shape such framework. Thus, this concept, like the forthcoming ones that are going to be discussed, is significantly delimited as a trade-off to pursue the more pragmatic goal of providing a concise tool for our analysis in further chapters.

The above summarises an example that allows us to assess the extent and limitation of the notion of a source. In the following subsections we address each feature separately. The first two aspects are expanded in the immediately following subsection (7.1.2), whereas the third feature is the object of subsection 7.1.3. We devote separate subsections to each feature to better understand the aspects we undertake and the aspects we are going to leave behind.

7.1.2 Amount of Information

Having in mind the previous example, we tackle the first aspect of sources. That is, we discuss the problem of different sources providing varying amounts of information. This seems a natural feature of what sources of information ought to convey.

As shown by our example, different sources might contain varying amounts of information. For example, D_1 and D_2 offer information concerning substance X_1 , which D_3 lacks. In particular, the first two sources provide information as to whether P should take medication Z based on the amount of X_1 it contains, while D_3 is silent as the role of X_1 contained in Z. Conversely, D_2 and D_3 offer information concerning X_2 , which D_1 does not offer. In particular, both D_2 and D_3 provide information as to whether P should take medication Z based on the amount of X_2 it contains, while D_1 is silent as to the role of X_2 contained in Z. Furthermore, D_2 offers more information than D_1 or D_3 by themselves. In particular, while D_1 only provides information as to whether P should take medication Z based on X_1 and D_3 only provide information as to whether P should take medication Z based on X_2 ; D_2 provides information as to whether P should take medication Z based on X_2 ; D_2 provides information as to whether P should take medication Z based on X_2 ; D_2 provides information as to whether P should take medication Z based on X_2 ; D_2 provides information as to whether P should take medication Z based on X_2 ; D_2 provides information as to whether P should take medication Z based on X_2 ; D_2 provides information as to whether P should take medication Z based on X_2 .

In light of the above, it is clear that different sources provide not only different amounts of information but also varying reasons. The specific notion of a source we employ allows such features to coexist. That is, our account of what constitutes a source does not restrain all sources to conveying only the same type of information nor the same amount of information. To have such restriction would seem to place a counterintuitive limitation.

In line with the discussion over varying amounts of information, one can point out that such feature can be made more precise from a mathematical approach to information, which would be a legitimate endeavour. For example, the *Mathematical Theory of Communication* (MTC) as it is referred to by Floridi (2016) relies heavily on such mathematical approaches to information. Such approaches have their origin in the work of Claude Shannon (Floridi, 2004), and following Floridi (2011) it is classified as Shannon-type information. Nevertheless, *for the purposes of our research*, such measures of information are not necessary. This is not to say that a mathematical account of the information contained or endorsed by different sources in a given context is not compatible with our overall aim, but it is not indispensable. That being said, our primary and foremost concern is to assess how the available information interacts in various circumstances. In this sense, we are far more interested in what Floridi (2004) refers to as *semantic information*.

The compliance of our approach to information with the MTC paradigm in the context of the framework we are constructing is something that is not addressed here, as it escapes the central objective of providing a framework of analysis to understand the problem of floating conclusions. Such endeavour, although interesting and legitimate, would be a corollary inquiry of our main

goal.

Finally, we want to highlight that our proposed notion of source does not make a strong distinction with respect to the information or argument being offered by a given source and the channel through which the information flows. For example, and taking into consideration the previously depicted situation of three medical practicioners advising on a given medication to address some health condition, we say that each medical practicioner is a source insofar as it provides an argument or information in support of taking the specific medication. In this scenario, the arguments supporting the conclusion can come (i.e. be *sourced*) from other human cognitive agents, or we might also stipulate that the information is coming from different journals found in an online medical database. Thus, in our use of the notion of source the origin ought not to be conflated with the content of what is being provided nor with the channel through which the arguments or information flows.

Although such distinctions concerning the flow of information are a relevant and interesting matter on its own, it is something we will not further problematize for simplicity purposes. In this sense, our notion of source can be said to be licentious with respect to the aforesaid nuances, although such licentiousness is not one that undermines our framework or its further purpose.

The above summarises our stance towards the mathematical measure of information in the context of our inquiry. As we argue later on, such approach is not to be taken into consideration in this chapter nor further stages of our research. In the next subsection, we look into the second property of the notion of source.

7.1.3 Directedness

In this subsection, we look into a less intuitive caveat that we want to imprint into the concept of a source, and that is a sense of *directionality*. This is used not as a generally relevant feature of the concept of a source, but rather an operational property that simplifies our use of such concept.

All subtleties aside regarding the varying amount of information, in the context of our previous example, one of the constant features is that all sources ultimately lead to a claim regarding P_1 needing to take medication Z. For example, D_1 supports the idea of P_1 needing the medication Z

based on the low quantities of X_1 . In the same vein, D_2 converges on the same conclusion, but because Z contains large quantities of X_2 (and makes a side note about X_1 being harmless in low quantities). Lastly, D_3 agrees with D_1 and D_2 on the need for P_1 to take Z, but his stance unlike D_1 is solely based on the large quantities of X_2 and mentions nothing about X_1 unlike D_2 . The above highlights the idea that regardless of the specific route each stream of information takes, they aim towards P_1 's need to take medication Z. This feature of convergence towards a specific end is what we label as the *directedness* feature of a cluster of sources.

The idea behind the *directedness* of a set of sources is that at any given time, as cognitive agents, a multitude of sources is readily available on demand. That is, a human cognitive agent has, at any given time, a variety of perceptual modalities (visual, sensory, auditory), not all of which need to converge in a specific target or object of attentional focus. For example, an individual in a stadium watching a rugby match is subject to a plethora of sensory inputs, some of which are not directly related to the developing game (e.g. people talking about sport gossip completely irrelevant to the game). In such case, and despite all perceptual modalities that can be considered as sources, they are sources not with respect to a specific object of attentional focus. Thus, they do not possess the specific nuance of *directedness* that seems of interest to us. In the context of the person watching the rugby match unfold, we would be interested in the sources that can contribute in a meaningful fashion towards the agent's belief about the outcome of the game or any of the phases of the game.

Conversely, we can think of a variety of claims or information sources such as accounts or testimonies that might have a *directedness*, but not in a useful fashion. For example, a cognitive agent has, at a given point in time, a collection of sources that provide information for several and unrelated claims. In such a situation, we can clearly state that the set of available sources has the property of *directedness*. Again, considering a person watching a rugby match unfold, all the directed streams of information that such person might have still leave room for some level of imprecision. That is, not taking into account the irrelevant and non-directed sources (e.g noise and other forms of distracting sensory input) there are still a great deal of sources that although related to the game itself, might not be related to a specific aspect of the game (e.g. the outcome of the game, the development of the game, the best player of the game, etc.). In such a vein, a special case of *directedness* is what we refer to as *unique directedness* (*u-directedness* for short). The *u-directedness* feature highlights the specific interaction between sources when all sources converge in a given conclusion or claim.

For example, a given cognitive agent can have a huge input of sources that point to a variety of claims, events or situations (like the person watching the rugby match unfold). In such case, it is of our interest to evaluate the subset of sources that provide information not only with a sense of target or aim (*directedness*) but also that the subset under scrutiny should have a specific and unique target or object of scrutiny (*u-directedness*). In this sense, we might be interested in sources or streams of information that not only target or converge on the rugby game itself, but we might want to consider the subset of sources meaningful to a particular aspect of the game like the outcome of the game, or the overall performance of a specific player (among many others).

It is important to state that although *directedness* appears to be a relatively desirable property for a given set of sources, *u-directedness* seems to be of a slightly higher demand. In particular, *u-directedness* restrain those directed sources which converge on a unique target. Such a requirement serves the purpose of capturing the interactions of sources with respect to a given object of attention. Thus, it may very well be the case that there are ample cases in which *u-directedness* might not be a strictly necessary property. That is, for a set of purposes the uniqueness of the target for a given number of sources might not be of interest. Nevertheless, we will be interested in sources which have this additional property.

To summarise the previous discussion, we can consider these features in the context of our initial example, i.e. the case involving the decision of three practitioners as to whether P_1 needs to take medication Z. First, as we have already stated, the fact that the various sources involved in the case provide varying amounts of information need not worry us since this, more than being a rarity seems to be an ever-present feature of sources of information. This does not preclude all of those items being sources of information. Furthermore, there needs not to be a mathematical account of the measure of information available in each of the sources for them to count as sources (at least not for our purposes). Second, all of the sources in the example possess the *directedness* property since all of them are directed sources in the sense that they are *sources about something in particular*. That is, they have a directionality towards each of the involved source's targets. In the same vein, all of the sources not only have a clear directionality (*directedness*), but such directionality converges on the same object (*u-directedness*). For example, D_1 , D_2 and D_3 are all sources providing information that converge in the event as to whether P ought to take the medication Z, despite the sources providing distinct routes or reasons for the same converging claim. Thus D_1 , D_2 and D_3 are said to have *u-directedness*.

The above concludes our assessment of the notion of source and the properties that are of relevance for the purposes of our research. In the following subsection, we address one of the salient features associated with sources that was left unaddressed, which is the conflict among sources.

7.2 The Notion of Conflict

In this section, we discuss the notion of *conflict* and assess the basic elements composing our precise understanding and use of conflict. Furthermore, we use both, the notion of a source (depicted in the previous section) as well as the notion of conflict to construct the first component of our twofold framework of analysis, i.e. the *source conflict criterion*. This section aims to provide the first component of our twofold framework to address floating conclusions. Like the previous section, here we follow the same descriptive nature.

7.2.1 Conceptual Precision

One of the most problematic, yet widespread features of cases that involve multiple sources is the occurrence of conflict. In this subsection, we specify the extent and limitation of such a feature. Again, it is without question that this can be heavily stretched to cover multiple areas of contemporary epistemology or other fields but, in the context of our research, we limit such theoretical constructs to address the problems at hand.

To illustrate the notion of conflict, we present the following example. Consider a trial which involves determining whether the defendant, which we refer to as *D* is guilty of a given crime, say the murder of *V*. Furthermore, let us assume that the case involves five witnesses (W_1 , W_2 , W_3 , W_4 and W_5), who have a claim regarding the involvement of the defendant in the case. In particular, we are said to have the following:

 W_1 : D was at the time and place of the murder of V,

and D shot V on eight different occasions.

 W_2 : D was at the time and place of the murder of V,

and D shot V on five different occasions.

 W_3 : D was at the time and place of the murder of V,

but D was only indirectly and accidentally involved in the murder of V.

 W_4 : D was at the time and place of the murder of V,

but *D* was just an innocent bystander in the events that lead to the murder of *V*. W_5 : *D* was not at the time or place of the murder of *V*.

The above example portrays a situation in which five different threads serve as sources of information as to whether D was involved in the murder of V, and the extent of his involvement. In this example, the notion of conflict is illustrated by the fact that any given thread seems to be at odds with the remaining four, i.e. the sources are not entirely coherent among each other. This is the basic trait we examine. At this point, we emphasise two notable features. First, the level of disagreement seems to vary between testimonies, i.e. some sources disagree more or less with the rest of the testimonies. Second, it is not always *prima facie* clear nor explicit if two different sources conflict, as this notion comes in degrees. These two features surrounding the notion of conflict is what we immediately address.

Concerning the basic feature, i.e. sources can conflict with each other, it seems clear that if we take any of the available testimonies, and contrast them with the remaining four, such account is at odds with the rest of them. For example, if we take W_1 and contrast this source with W_5 it is patent that they disagree as much as any two sources in this particular example can. On the one hand, one source not only provides information that D murdered V, but it does so on the basis that D shot V on eight occasions. On the other hand, the alternative source of information supports the fact that there was no way D could be the perpetrator of the crime against V, since D was not at the place nor time of the events leading to V's assassination. This same process can be instantiated with whatever pair of sources we decide to compare, i.e. no pair of sources are in absolute agreement.

In light of the above, it is safe to state that a given number of sources can conflict with each other, despite providing different venues and amounts of information with respect to a specific claim. The possibility of conflict is one of the features of sources that we take as fundamental rather than subsuming it as a subsidiary property of sources. This is why the conflict feature in the context of sources is of instrumental value in the development of our framework to understand floating conclusions. Having stated the above, we proceed to analyse the two features of conflict we have emphasised.

First, concerning the level of disagreement, we note that depending on which subset of the

available sources we wish to focus, there are various degrees of disagreement. That is, conflict comes in degrees rather than being a discrete matter. As an example, let us take again sources W_1 and W_5 . According to the former source, not only there is information that D murdered V, but we also have information on how such crime happened. According to the latter source, not only do we have information as to D not being the murderer of V, but also, we have information that would make it physically impossible for D to have murder V given that he was not at the place or time of the assassination of V. This sole aspect points out to the existence of a disagreement. Nevertheless, we can also compare sources W_3 and W_4 . According to W_3 , D was involved in the murder of V, but only indirectly and accidentally. Whereas W_4 states that D was only an innocent bystander in the events that lead to the assassination of V. In light of the above, we can intuitively say that there is a level of disagreement between sources W_3 and W_4 , but such disagreement is less than when we contrasted W_1 and W_5 . In the process of comparing the two pair of sources, we notice not only that there is disagreement between the sources, but also that the disagreement, but also there are major disagreements, as well as minor disagreements.

Second, we notice that it is not always *prima facie* clear if two sources of information are in conflict. That is, the point at which we determine that a conflict emerges is not self-evident, but more often than not a matter of intersubjective agreement. To illustrate this second feature of conflict, we recall that of all five sources, W_1 and W_5 are the pair of sources that provide the most profound level of disagreement. In such a vein, we also note that the other end of the conflict spectrum happens between sources W_1 and W_2 . To support the aforementioned, we notice that both provide information towards the claim that D murdered V, and not only that but also D's way to carry such crime. To that extent, we can state that both sources are compatible. Nevertheless, W_1 and W_2 disagree in the number of times D shot V, the former supporting the claim that D shot V eight times, while the latter supporting the claim that D shot V on five different occasions. In this context, an intersubjective agreement amounts to a conflict between sources. That is, the sole disagreement between sources might not be sufficient to qualify such disagreement as an explicit conflict to the extent that the disagreement places the sources in a mutually undermining relation.

7.2.2 Source Conflict Criterion

Having reviewed the main features of the notions of source and conflict, in this subsection we construct a precise account of these two aspects as a criterion, which is of further use in this and later chapters. To approach this task, we first present the *source conflict criterion*. Later on, we explain several features that are assumed within our discussion of both notions. This analysis serves to further ground the criterion itself.

Source Conflict Criterion (SCC): A given number of sources are said to be conflicting if the information endorsed, contained or supported in each source is incompatible with information endorsed, contained or supported by one or more of the remaining sources or, in the worst case they mutually undermine each other. Otherwise, the sources are said to be non-conflicting.

First, the source conflict criterion is thought of as a criterion that operates on sources that are *uniquely directed*. That is, the criterion is relevant and non-trivial in those cases in which the sources under scrutiny refer to or are directed towards the same claim or information. For instance, let us recall the case in which we had five different testimonies concerning the murder of a person. Here, inquiring whether the different testimonies stand in conflict with one another, is meaningful to the extent that we have a common target motivating our inquiry.

For example, let us briefly consider the case in which we have three different sources of information. Let us assume that the first source provides us with information concerning Adam's involvement in the murder of a given individual; the second source provides us with information concerning Ben's involvement in a traffic infraction, and the third and last source provide us with information concerning Cody's involvement in an insider trading accusation in the stock market. Furthermore, let us imagine that none of these events are related with respect to a given target. Given these three sources of information, here we can inquire the extent to which those sources are in conflict or not (setting aside the potentially enormous complexity to determine if this is the case). Nevertheless, there seems to be no useful purpose of determining the existence of conflict, beyond the remote, yet not impossible case in which such events are related to an unmentioned piece of information, claim or situation (earlier precluded). The aforesaid exemplifies a case in which although theoretically not impossible it would certainly render instrumentally or practically useless to engage in determining the existence of conflict among sources which are not uniquely

directed.

Second, and having considered the notions of source and conflict, we discuss the problem of thresholds. One of the first things we wish to point out is that, in our depiction of the notion of conflict, we stressed that it operates on a continuum since it comes as a matter of degree. Thus, one can argue that determining conflict is a matter of degree and not a discrete state. In such vein, one could push the line by which determining whether two or more sources are in conflict would be impractical or vague to establish. Nevertheless, in the context of the source conflict criterion, and to escape the above, we can appeal to the notion of thresholds.

According to the notion of a threshold, we do not require the object of analysis to be discrete. To the contrary, it is not relevant to the nature of the object of analysis, because it can accommodate continuous or discrete phenomena. The reason is that a threshold does not need to be an intrinsic feature, but rather a commonly agreed intersubjective property.

For example, we can say that the point at which the physiological process of lactic acidosis occurs during an athletic performance for a given individual is an objective feature. Nevertheless, the point at which a given testimony in a trial is considered self-defeating, or the point at which two different testimonies undermine each other, contrary to the process of lactic acidosis in athletes, is not an objective feature of them, but rather an intersubjectively agreed threshold. Thus, we can safely state that based on the notion of a threshold, the phenomenon under scrutiny does not need to be discrete. On the one hand, if the phenomenon under scrutiny is discrete, then we can specify the point in which a conflict occurs (because we would have an objective measure against we can assess that which is under scrutiny). On the other hand, if the object of inquiry is not discrete but rather of a continuous nature, we can still specify an intersubjectively agreed-upon point at which the object of analysis is said to fail or pass the threshold.

Third, if the sources under consideration do not qualify as conflicting, then we describe such sources as non-conflicting sources. The above state of affairs can entail various posibilities. On the one hand, it could very well be the case that none of the sources interfere with each other due to them involving different types of information. On the other hand, it might very well be the case that the sources under scrutiny reinforce and support each other. Both of the previously described cases can range from one end to the other in a continuum. Nevertheless, for practical purposes, we do not inquire further which end of such spectrum does non-conflicting sources fall into. In such

vein, we operate on the assumption that non-conflicting sources, in the worst-case scenario, have nothing to do with each other. Such state of affairs does not preclude the case in which they stand in a relation of mutual support. This is a deliberate decision for the purposes of simplification within our research. Whether opening up this aspect of our analytic framework has further repercussions *beyond the scope of the object of our research* is something that is left to be further explored.

The above summarises our assessment of the notions of *source* and *conflict*; their extent and limitation. Furthermore, we have highlighted the properties that are of central importance for the purposes of our research. Finally, we have placed both concepts in terms of a criterion. We proceed in the following section to address the notion of cost.

7.3 The Notion of Cost

Here, we follow the thread from the previous section and consider the second of our two-fold framework of analysis, namely the *cost criterion*. We examine the notion of cost and discuss how it shapes the criterion we construct around such notion. This section aims to present the notion of *cost* and the *cost criterion* and assess its role as the second component of our twofold framework of analysis to address floating conclusions.

7.3.1 Conceptual Precision

Having reviewed the source and conflict aspect of our framework of analysis, we sketch the third component, namely, the notion of cost. A similar concept is used across a wide variety of fields (such as economics, politics, computer science, etc.). As such, the extent and limitations of this notion fluctuates depending on the knowledge domain. Nevertheless, in this subsection, we highlight a particular understanding that points to a very general account to be employed in a variety of domains. To address this, we use an example to illustrate the particular approach to the notion of cost that we use in the context of our research.

Let us consider the case in which a judge needs to determine the culpability or innocence of a given individual Adam. This case could potentially result in Adam receiving a life sentence or

death penalty. Regardless of the specific circumstances in which the case is being discussed, which can complicate or facilitate the decision-making process, we can state that the cost of the decision and information processing is high, due to the possible consequences that such decision or information processing entails. Conversely, we can think of a case in which a data entry agent in a government census institute is trying to determine whether Adam should be added to the number of murderers for the year under study. In this second case, the legal or directly practical implications are far less severe. First, we can state that besides the need for accuarate government statistics, the sole addition or subtraction of one entry in a large database is not going to solely determine government policies nor change the political arena. Second, if we take aside the legal aspects of being listed in such a database (we can assume that this is being done *a posteriori* from legal action), or specify that such a database has no legal consequences for the individuals (as most of the census-driven-databases lack), then we can also state that the cost ranges from low to non-existent.

Both of the previous examples highlight one salient property to further emphasise the extent and limitation of the concept of cost we have in mind in the context of our research. In particular, we want to frame the notion of cost as a practical or otherwise tangible and weighty consequence for some or all of the involved elements surrounding a case, rather than cost as the intrinsic burden of information processing. To address this, we first consider cost as a complexity measure of information processing.

In the example we previously stated, we could apply a notion of cost related to the measure of the intricacy of processing all the information related to the case. In the first example, we can measure the cost of compiling the testimony of all of the witnesses, or we can think of the cost in terms of fact-checking the information provided by all of the witnesses. In the second example, we can measure the cost to the government statistics agent's accessing the proper sources and determine if Adam is a murderer, like the time involved in such operation in contrast to the benefit of an accurate database. Both uses of the notion of cost are completely legitimate, but not quite appropriate for our purposes. Instead of proceeding with a notion of cost as a complexity measure of information processing, we emphasise the notion of cost as that of *consequence*.

To illustrate the underlying intuition of cost as consequence, we recall the legal example previously considered. In such example, and despite the potential complexity of processing all the information associated to the case, given that Adam is facing life in prison or the death penalty, we can state

that the cost of the decision made, regarding the innocence or culpability of the defendant, is high. Furthermore, we can state that, following the sense of cost as consequence, it is high regardless of two very legitimate variations of cost as complexity, which we immediately illustrate.

First, let us consider the case in which we have numerous witnesses, of whom a large percentage have left the country, and that it would be very intricate materially, legally and financially to track down. Furthermore, let us assume that the ties of some of the witnesses to the role Adam plays in the murder are somewhat unclear and that such witnesses could have a potential conflict of interest in Adam's culpability. Here, we can state that the cost of processing all of the information related to the case is high due to the aforesaid intricacies. Nevertheless, we claim that the cost of this case is high regardless of the previous reasons. The decision-making process is high because a person's life is on the line.

Second, let us assume we have a relatively small number of witnesses that are, all things considered relatively transparent and reliable sources of information, which can be very easily tracked down and whose testimonies can be very easily compiled and checked. Here, contrary to the above, we can have that the cost of assembling and processing all of the information involved in the case is immensely lower. As such, the case can be qualified as a low-cost situation *in the sense of information processing*. Despite the aforementioned, we still want to state that the cost of taking a stance in such a situation is high because a person's life is still on the line.

Thus, what both of the previous attenuating circumstances highlight is that cost, in the sense of cost as the consequences of the information processing or decision-making process, is not related in any form to cost as a complexity measure of the information processing or decision-making process itself. Given the specific needs of our research, we make use of the former notion of cost rather than the latter. Nevertheless, we do not claim that one account or approach is misleading or suboptimal in comparison to the other. Our choice is a strictly instrumental one given our particular purposes.

The above can be placed in the context of Wilson and Sperber's (2004) observations on the notion of *relevance* for the purposes of deriving meaning. In the process of presenting their theory of meaning, the aforesaid authors observed, among other features associated with such a process, the problem of cognitive effects and processing effort relative to a given agent. In particular, Sperber and Wilson state that:

a. Other things being equal, the greater the positive cognitive effects achieved by processing an input, the greater the relevance of the input to the individual at that time. b. Other things being equal, the greater the processing effort expended, the lower the relevance of the input to the individual at that time. (2004, 609)

We follow the approach of Wilson and Sperber (2004) only to state that regardless of the processing effort involved in a specific situation, there are instances in which processing effort is not inversely correlated to the cost of a given situation. In this sense, it is the practical consequences of a given situation that shape the high or low cost associated to it, as opposed to relying solely on the complexity of the information processing associated with the case.

Lastly, a salient feature associated with the process of determining the cost lies in the fact that we refer to the *net cost* of the overall process. That is, whenever we talk about a given decision-making process as costly, we assume that such property has been determined to take into consideration all of the intervening factors, and what we have is the net result. To illustrate this aspect of our notion of cost, let us recall the case in which we needed to determine whether Adam is a murderer for a specific case, which involves a fixed number of sources. To address the above, we look into the net cost feature at two different levels: (1) on the level of individual sources and (2) on the level of the overall cost as a reflection of what happens when all of the sources are considered, i.e. the net cost.

First, and to address (1), we can state that each independent source has taken a deliberation process by which it has arrived at the conclusion whether, all things considered, Adam is the murderer for case under consideration. Furthermore, in the ongoing deliberation of each source, one can expect that such decision-making process has taken into account the overall burden and importance of their testimonies towards the general process of determining whether Adam is the murderer for case under consideration.

Second, to address (2) and directly linked to the notion of cost at the level of sources, i.e. (1), whenever there is a claim that the consequence of a given information processing or decision-making deliberation is costly, such claim has the presupposition that it has taken into account the analysis of all other possibly conflicting factors. In the case of the decision-making process to determine whether Adam is a murderer, we operate on the assumption that if the case is of a high-

cost, then such high-cost would be the net cost of all things involved in the case.

For example, if (a) we move the decision-making process of determining whether Adam is a murderer from a legal trial to an entry in a database of a government statistics office, then the cost might be low. Nevertheless, if (b) such government statistics office database will be a key input into a forthcoming controversial legislative debate, then the cost can be described as high. Furthermore, it can also be the case that (c) even if such database will be a key aspect in the forthcoming legislative debate, Adam's entry will certainly not affect an established trend, rendering the whole situation as a low cost one. However (d) for political reasons, it is quintessential for the database to uphold the utmost accuracy, thus rendering the decision-making process as one of high cost.

The key aspect to note from the above variations is that there are many *attenuating circumstances* to determine if a given decision-making process or claim is of a high or low cost. Nevertheless, and *for simplification purposes*, such consideration is assumed to have taken into account all of the potential attenuating factors and reflects the overall net cost. The reason behind this working assumption lies in the fact that we want to simplify as much as possible all of the potentially complexity-increasing factors. The above being said, we do not hold the stance that such process is trivial or straightforward, but for simplicity sake, we operate on the assumption that it is a given.

In line with the above, and as we have pointed out, whenever we talk about the cost, such notion is used in the context of determining a claim, endorsing an attribute or simply put the burden of the consequences of a decision-making process. For simplicity purposes, we resume such a mouthful as the cost of *situations*. Furthermore, it is important to highlight that when we talk about a given situation having a given cost associated to it, we work on the assumption that the assessment of the cost is made by a cognitive agent facing such decision-making process that renders the assessment as a subjective evaluation. That is, the determination of a given situation as being of high or low cost resides on a given cognitive agent, presumably a human cognitive agent, which unquestionably has limited resources and information when determing such cost, which makes room for imprecision or innacuarracy from what one could refer to as the true and objective cost of the situation. Thus, the notion whe employ of cost is one that can de characterized as being subjective as opposed to a notion of cost that attempts any form of objective assessment. This latter notion of cost would demand stringent epistemic capabilities which are simply unavailable for human cognitive agents.

Finally, and following the previous thread, the cost associated to these situations does not take into account a time progression. That is, whenever a given situation is determined as being high or low cost, such assessment is assumed to be made at a fixed point in time, although our assessment of the cost of the situation might fluctuate over time. This of course can be framed either as a limitation, but seems nothing more than and intrisnsic feature of human cognition insofar as inferences are drawn in fixed points in time, regardless that such inferences could be corrected later on. The recognition of the defeasibility of inferences made by human cognitive agents based on time progression is an assumption underlying non-monotonic and defeasible reasoning. As such, we are not using a notion of cost that tries to capture some form of of atemporal accuaracy. This of course, is a corollary of the subjective nature of our notion cost.

7.3.2 Cost Criterion

In this subsection, we take our previously constructed notion of cost and place it in terms of a criterion. After stating the *cost criterion*, we explain four key properties that we attach to it, in the context of our previous discussion of the notions of sources and conflict respectively.

Cost Criterion: The cost of endorsing an attribute or statement through a specific informationprocessing mechanism or decision-making strategy is said to be high if the material and/or practical consequences of endorsing that attribute or statement in the given context is burdensome, by means of financial, legal, medical or other domain-specific criteria of burden, despite the complexity involved in processing the information or the complexity involved in the decision-making process that ultimately endorses the attribute or statement. Otherwise, the cost of the endorsed attribute or statement is said to be low.

First, one of the salient features involved in the cost criterion lies in that there are multiple venues through which the cost can be determined. That is, there is no sole factor that, *a priori*, supersedes any of the remaining factors involved in the process of determining the cost, but rather such precedence is to be decided *a posteriori*. For instance, the example of determining a patient's need for a particular medication can be said to be costly for a plethora of reasons. In particular, we can think of at least three potential reasons that can shape this as a high cost case: (a) it can be said that the intake of that specific medication will take an enormous toll on the patients overall health or, (b) it can be said that it will be an enormous financial burden to the healthcare system to provide

the patient such medication, and finally (c) it can also very well be the case that the prescription of such medication will have immense legal risks for the medical facility. The point is that, of all the possible determining factors, no single factor is said to have an *a priori* precedence over the others. As such, this precedence is to be intersubjectively determined given a sufficient amount of reasons and conditions which are knowledge-domain specific.

Second, and in relation to the immediately preceding aspect of the cost criterion, whenever we classify a case as high cost, what the cost criterion must reflect is the net cost of all the factors associated with the case. That is, if we determine that a situation is of high or low cost, such property reflects all of the potential attenuating circumstances. This is just an operational aspect of the cost criterion that comes as a consequence of the feature that we have stated in the previous subsection for the notion of cost.

Third, after having presented the cost criterion, it is important to address the problem of thresholds and its role in such criterion. In such a vein, we draw a similar argument regarding the nature of thresholds in the context of the cost criterion as it was put forward when we reviewed the source conflict criterion. Despite being a similar argument, we proceed with a step by step description of our stance. Before embarking on our very succinct sketch, we have to note that, as in the case of conflict, the cost is an external and intersubjectively determined benchmark that is used to determine if a given situation satisfies or suffices to be typified as high or low cost.

Cost, as we conceive it, lies not within the sources of information themselves but rather is an externally determined measure. To begin, the attenuating factors that shape a given decision-making process as a high-cost or low-cost might be measured or quantified in one of two ways. On the one hand, it can be the case that such attenuating factors are discrete. If this is the case, and since we operate on the assumption that despite the level of difficulty in finding a threshold, when one is given, such a threshold can be used to determine if the discrete value of the net cost is high or low. This represents no inconvenience whatsoever. On the other hand, it can be the case that the attenuating factors that shape a particular decision-making process as a high or low-cost are not discrete but rather continuous. If this is the case, we can contrast such a measure with the attenuating factors regardless of their continuous nature based on the existence of a threshold that guarantees a benchmark to assess whether such factors exceed or fall short of the threshold.

Fourth, and after having reviewed all the caveats associated with the cost criterion, we highlight

what is implied when, based on this criterion, we state that the burden or consequences of a decision-making process associated to a given situation are low. Here, we take the stance that the overall decision-making process and more importantly, its burden and practical consequences do not qualify as being costly. Nevertheless, the aforementioned does not imply that the decision-making process is trivial or irrelevant. To the contrary and given that the cost is determined by an external benchmark, a low-cost situation entails that it does not surpass *the established threshold*. To put it in another way, a low-cost situation refers to an external benchmark, but says nothing concerning its intrinsic or otherwise associated cost.

For example, determining Adam's culpability for the purposes of a government statistics database, we might end up establishing that the cost is low. Nevertheless, what this means is that when every attenuating circumstance is taken into account, then the cost (whether it be high or low) refers to a general assessment since we are dealing with the net cost given a particular threshold. This does not preclude that the cost of any of the individual attenuating factors may very well be contrary to the net cost. For example, in the previous example, we can assume that: (a) the database will play a strategic role in the discussion of a forthcoming controversial legislative debate, but that (b) Adam's entry will certainly not affect an established trend, then we might end up determining that the situation is of low cost. Despite all of the aforementioned, we can still take (a) and objectively state that this attenuating circumstance *by itself* renders the situation as a high-cost one. Thus, if a given situation is said to be of low cost, all that we can establish is that it does not exceed the overall net cost benchmarked against the determined threshold, but it does not entail that any of the particular constituents are of low-cost.

The above concludes our analysis of the extent and limitation of the notion of cost along with its role in terms of a criterion. In the following and last subsection, we merge the source conflict criterion and the cost criterion into a single criterion, which is the main tool of analysis to account for floating conclusions.

7.4 The Criterion

In this section, we merge the components that we have developed so far. We present the *source conflict cost criterion* (SCCC). We provide a methodological simplification of the SCCC and explain how some of the subtleties associated with this framework of analysis can be simplified in
terms of a binary model. This section aims to examine how both criteria can be merged for the analysis of a given situation. This has the objective of setting the stage to use this two-fold component in the following chapters to account for the problem of floating conclusions. Nevertheless, before embarking on the aforesaid, we need to consider the extent and limitation surrounding the interplay of the components in our framework, and what constraints we are going to set in place for such purposes.

7.4.1 A Binary Model of the SCCC Criterion

In the previous sections, we have reviewed the *source conflict criterion* and the *cost criterion*. In this section, we propose a simpler but equally legitimate use of this tool. In particular, we put forward a binary model of the source conflict and cost criteria. Moreover, we provide the reasons behind this simpler approach. Before embarking in the description of this simplified model, we discuss the foundations and reasons for such an approach.

As we have extensively stressed, the nature of a myriad of factors involved in determining conflict and cost are more often than not of a continuous or non-discrete kind. Nevertheless, we can often point out to external intersubjectively determined thresholds as a mechanism to, complexity aside, decide whether a given number of sources conflict or if the cost of a situation is high or low. That is, given the notion of a threshold, the problem of deciding whether a given number of sources conflict can be addressed in a simpler fashion. Having stated the above, we extend such notion to further simplify both, the source conflict criterion and the cost criterion, and we do so by considering a binary model of such notions. Our binary model of both the source conflict and the cost criterion relies on a two-value state for both at any given time for a specific situation. The basic underlying idea is that any of the two criterions can be either high (+) or low (-), and such states can be understood as binary values.

First, we have the source conflict criterion that establishes if a given number of sources conflict with each other. According to a binary model, we have the following:



Figure 7.4.1-a Binary Source Conflict Criterion Model

Following Figure 7.4.1-a, a given number of sources can either be conflicting (+) or nonconflicting (-). If a given number of sources are said to be +SC, then all things considered, we could state that they surpassed a given threshold, rendering them as conflicting sources. If such sources are determined to be -SC, then they might or might not have some level of disagreement, but all we can state is that they have not surpassed the agreed threshold that could describe them as conflicting sources.

Second, we have the cost criterion which determines whether the cost of a given situation is high or low. Thus, we have the following binary model:



Figure 7.4.1-b Binary Cost Criterion Model

According to Figure 7.4.1-b, the situation can be of high cost (+) or low cost (-). If a given claim, information or implication of a decision-making process is determined to be +C, the burden or implication of the decision-making process concerning the information or claim under scrutiny can be said to be of high cost (taking into account that here, cost refers to net cost). If a given claim or implication of a decision-making process is determined to be -C, we can safely say that the burden or implication of the decision-making process is low. The latter does not entail that there are no high-cost factors, only that when all factors are taken into account the net cost renders the situation as a low-cost situation.

The above provides a simplified binary model of both the source conflict and cost criterion. Such approach simplifies the evaluation process of a given situation that involves sources of information and the associated cost. In the next subsections, we further illustrate the binary model of the *SCCC*, by returning to a previously discussed example. Although we intend to use the *SCCC* as a single criterion, we first proceed to approach the example in terms of the source conflict criterion (section 7.4.2), then in terms of the cost criterion (section 7.4.3). Finally, we explain how both evaluations can be merged into a single framework of analysis (section 7.4.4).

7.4.2 An Example of the Binary Model of the SCCC: Source Conflict

In this subsection, we use the binary model of the source conflict criterion to assess the example given in section 7.3.1. The aforesaid case involved determining D's possible role in the murder of V, given the account of five different witnesses. The testimonies ranged from placing D at the time and place where the murder of V occurred and linking D with shooting V several times, to testimonies which do not place D at the time nor place where the crime against V happened. We explain how this example is accounted for in terms of a binary approach to the source conflict criterion.

First, we can have that, after taking into consideration all of the sources and level of disagreement, we conclude that the involved sources are in enough disagreement with respect to a given threshold to qualify as conflicting sources. In such a case, we state that the involved sources are +SC.

Source Conf	flict Criterion
+ SC	- SC

Figure 7.4.2-a High Source Conflict Criterion

For example, in the case of the testimonies surrounding D's role in the assassination of V, and after taking into consideration all five witnesses, we conclude that the involved sources substantially disagree. On the one hand, some testimonies locate D at the time and place of V's murder and link him as the author of the crime. On the other hand, we have a testimony that does not even locate

him at the time or place where the events ocurred. Given the above, and assuming that such a state of affairs surpasses an already given or agreed-upon threshold, we can state that the testimonies surrounding D's role in the assassination of V are in enough disagreement to qualify as conflicting sources, rendering such sources as +SC, as shown in Figure 7.4.2-a.

Second, we could also have that after taking into consideration all the involved sources, we conclude one of two things. On the one hand, we may find out that there is no disagreement whatsoever among the sources, such that, regardless of the working threshold, they will ultimately be classified as non-conflicting. On the other hand, we may find that there is a level of disagreement among the sources, but we can think of a more demanding and higher threshold, by which the disagreement of such sources does not amount to or qualify them as conflicting sources. Either way, we could have that the sources are non-conflicting. In such a case, we have that the involved sources are -SC.

ict Criterion
- SC
•

Figure 7.4.2-b Low Source Conflict Criterion

For example, after taking into consideration all of the same five different witness's testimonies, we realise that there is, in fact, a level of disagreement among the sources. Nevertheless, due to some legislative nuance we might have a higher threshold, by which the disagreement of the involved testimonies although existent, does not qualify the testimonies as conflicting sources. In such a vein, we can state that the testimonies surrounding D's role in the assassination of V are non-conflicting sources, i.e., -SC, as shown in Figure 7.4.2-b.

Despite the particular intricacies involving the working example, what we want to stress at this stage is that, for simplicity purposes, we are going to work on the assumption that it can all be reduced to a two-valued criterion which renders a set of sources as +SC or -SC. This, of course, is a significant simplification, but one for which we have taken care to extensively detailing not only the instrumental benefit but also the theoretical basis.

7.4.3 An Example of the Binary Model of the SCCC: Cost

Having undertaken the analysis of our working example in light of the binary model of source conflict criterion, in this subsection we address the same task, but now in the context of the binary model of the *cost criterion*.

To address the aforementioned, we consider a potential listing of all the relevant factors to be taken into account to determine whether such a situation is of high or low cost. In our example we assume the relevant factors are as follows:

- (a) The decision-making process is used for an official government database.
- (b) The database is a key input for a controversial legislative debate.
- (c) Adam's entry does not affect an established statistical trend.
- (d) The database's utmost accuracy is needed for political reasons.

First, after evaluating and balancing all of the attenuating circumstances for a given situation and in light of a specified threshold, we can conclude that the cost of a given situation is high. In such a case, we state that the situation is +C. Here, we are working with the assumption that this reflects the net cost, which might not reflect the cost of the individual factors.



Figure 7.4.3-a High-Cost Criterion

For example, in the context of our case study, after evaluating all of the initially described attenuating circumstances, we can decide that (a) provides reasons as to why the situation is of low cost. That being said, (b) inclines such cost towards the high end. The high cost shifts back to the lower end of the cost spectrum in light of (c). Nevertheless, when taking into consideration (d) we resolve that the case is best treated as a high-cost one. Therefore, we state that the situation is +C,

as shown in Figure 7.4.3-a.

Second, when all the attenuating factors for a given situation are considered and the overall net cost fails to exceed a previously established or agreed-upon threshold, we can state that the situation is a low-cost situation. Thus, we classify the situation as -C.

Cost	Criterion
+ C	- C

Figure 7.4.3-b Low-Cost Criterion

For example, in line with our current case study, and despite the linear progression of high/low attenuating circumstances starting with (a) and ending in (d), we can have that, for a given particular reason, factors (b) and (d) are to be of lower weight than factors (a) and (c). Thus, if we hold the same threshold, the case can be classified as a low-cost situation (as shown in Figure 7.4.3-b). Furthermore, we can also have that there is no variation in the weights of the attenuating factors, but we can be working with a different and stricter threshold. With a stricter and more demanding threshold, the attenuating factors, although providing insight as to how it can render the situation as one of high cost, may fail to surpass the new and stricter threshold.

At this stage, what we want to highlight is the simplification of all intricacies involved in determining the cost through a binary model of the cost criterion. Again, as with the case for the source conflict criterion, this is a significant simplification. Nevertheless, such simplification has been extensively grounded in a way that it does not represent a transgression of what we have developed so far.

7.4.4 SCCC: The General Schema

So far, we have presented two fundamental criteria, namely, the source conflict and the cost criterion. We have extensively detailed the caveats behind the elementary notions of each criterion. Finally, we have simplified all of the potential complexities to determine whether a given situation

is of high or low cost, or whether the involved sources are in conflict, through the use of a binary model of such criteria. Thus, all of what we have described so far can be summarised through the following schema:

	Source Conflic	t Cost Criterion	
+ SC		- SC	
+ C	- C	+ C	- C
1	2	3	4

Figure 7.4.4 Complete Source Conflict Cost Criterion Binary Model

The schema in Figure 7.4.4 is organised as follows: First, there is the source conflict criterion, which, as we have previously stated, can be high or low. Second, we have the cost criterion, which can also be high or low. The arrangement gives us all the possible high/low value combinations for the source conflict and cost criterion. Thus, given the specific properties of any given situation, we can classify it as belonging to one of the four categories. To illustrate this, we use our working example of D's involvement in the assassination of V in light of the five different testimonies.

First, reading the schema from left to right, in quadrant (1) we have the situation in which there is conflict among the sources, and the situation is a costly one. This situation can be better depicted, by the case in which the existing level of disagreement is sufficient to surpass or exceed a given threshold to qualify the witnesss' testimonies as conflicting, i.e. +SC. In this context, we also have that all attenuating factors reach the threshold rendering the situation as one of high cost, i.e. +C.

Second, quadrant (2) shows the case in which there is conflict among the sources, but the situation is of low cost. In this context, like the previous one, the testimonies are said to be in enough disagreement concerning a given threshold to qualify such sources as conflicting, i.e. +SC. Nevertheless, here the attenuating factors are regarded as being insufficient for the given threshold to qualify the case as a high cost situation. Notice this can be due to the weight given to the different

factors or to the fact that the threshold is a stricter, all of which renders the case as a low or no cost situation, i.e. -C.

Third, quadrant (3) depicts the case in which there is no conflict among the sources, and the situation is of a high cost. Here, we have that for whatever agreed upon working threshold; the testimonies are not in disagreement or that the existing disagreement does not qualify, for any particular set of reasons, as a set of conflicting testimonies. Thus, the sources are low or non-conflicting, i.e., -SC. Furthermore, despite having non-conflicting sources, we have that the attenuating factors do render D's involvement in the assassination of V as a situation of high cost, i.e. the situation is +C.

Fourth, quadrant (4) shows the situation in which there is no conflict among the sources and the situation is of low cost. In this case, like in quadrant (3), the testimonies are said to be in not enough of a disagreement to qualify them as conflicting sources, rendering the sources as -SC. Nevertheless, unlike the previous quadrant, the attenuating circumstances do not exceed the agreed-upon threshold of cost, rendering D's involvement in the assassination of V as a situation of low cost, i.e., -C.

The above concludes our overview not only of the binary model of source conflict and cost criterion (which we have previously undertaken) but also the interplay of both criteria merged as a single framework for the assessment of a particular situation.

7.5 Chapter Summary

In this chapter, we constructed two criteria, namely the *source conflict criterion* and *cost criterion*. We have examined the elements and building blocks underlying both, such as source, conflict, and cost. Furthermore, we discussed the extent and limitations that such notions have in the context of our inquiry. Lastly, we stressed the idea that the *source conflict cost criterion* is to be regarded as a two-fold framework to assess decision making and information processing situations.

Although we explained with detail all the aspects associated with both criteria, as well as using them in some preliminary examples, we restrained ourselves from applying both criteria to analyse in full extent the problem of floating conclusions. In the next chapter, we apply our framework to the problem of floating conclusions. Furthermore, we analyse our stance towards this problem, based on the framework of analysis that we have constructed throughout this chapter.

Chapter 8 Analysing Floating Conclusions through the SCCC

In this chapter, we analyse the problem of floating conclusions in the context of direct and indirect modes of skeptical inferences. Nevertheless, unlike chapters five and six, here we not only undertake a descriptive look into such problems, but we apply the tools proposed in chapter seven, namely the *source conflict and cost criterion* to understand floating conclusions. The main purpose of this chapter is to assess cases of floating conclusions which involve conflicting sources (+SC) and non-conflicting sources (-SC), along with the high cost (+C) and low cost (-C) attenuating conditions. The aim is to accommodate all of the above within the SCCC framework. We aim to examine how such modulating factors shape the problem of floating conclusions in the context of the direct/indirect skepticism debate. Furthermore, a last but important objective is to provide a normative stance towards floating conclusions and the direct/indirect skepticism debate.

8.1 Conflicting Sources with High Cost

In this section, we examine both direct and indirect approaches to skeptical inferences and floating conclusions. We undertake this analysis in the context of conflicting sources cases and high-cost situations. This section has three crucial objectives. The first objective is to provide a scrutiny of cases involving floating conclusions in which we have conflicting sources with a high cost. The second objective is to target the aforementioned with the aid of our framework of analysis. The third objective is to construct and justify our stance towards floating conclusions with these SCCC specific features.

8.1.1 Cases

In this section, we look into cases discussed throughout the fourth chapter that comply with the properties of having conflicting sources and a high cost. The main purpose is to assess the +SC +C value combination for our assessment of the acceptance/rejection of floating conclusions.

In the process of examining the arguments in favour of the direct account of skeptical inference

throughout the sixth chapter, we used three key case studies, all considered by Horty (2002). First, we reviewed a case involving an investment decision based on the potential inheritance of a large sum of money, but with a pair of conflicting reasons as to whether the money would be inherited ("inheritance case"). Then, we analyse a case, within a military context, that involved deciding whether to attack an enemy based on the account of two different spies which agreed on the enemy's retreat, but disagreed on the reasons as to why the enemy could be held as weak and retreating ("military case"). Finally, we examined a case, involving a public policy decision to face an inevitable economic downturn, with the caveat that such prediction was made by two different groups of economic experts who agreed on the result, but disagreed on the reasons that would lead such bleak economic outlook ("macroeconomics case"). All of the cases have various important features in common, which we immediately review.

First, and the most salient feature is that all cases have a significant degree of conflict. That is, in all three cases, the sources conflicted with one another to some extent. This conflicting feature, although by itself does not serve as a categorical reason to dismiss the involved floating conclusion, presents the first and initial aspect that evoked the need for further scrutiny to inquire for the grounds and justification of such floating conclusion. For example, in the inheritance case, the accounts of both siblings are in conflict. Likewise, the military case presents a situation in which the sources of information conflict with one another, as the reasons one provides undermines the alternate account.

Second, one aspect that was systematically pointed out in the fifth chapter consisted of the fact that, in such cases, there is no common, conflict-free line of reasoning shared among all of the sources. Under such circumstances, this feature provides a strong reason supporting the idea that the floating conclusion although plausible, certainly does not constitute the most conservative stance to take in the face of the situation itself. For example, in the macroeconomics case, both sources of information provided conflicting and directly opposing reasons to support their conclusions. Nevertheless, there is no other common argument through which both groups of financial experts coincide, beyond such conflicting lines of reasoning. Thus, the situation is one in which the only available lines of reasoning conflict among themselves, and there is no middle ground in which they converge. This provided strong grounds to support abstaining from endorsing the floating conclusion.

Third, and a recurrent feature of the cases in support for the direct approach to skeptical reasoning

is that they all share an explicit or implicit notion of cost associated with the situation. This is not the information-processing sense of cost, but rather, the sense of consequences entailed by the decision itself. For example, the inheritance case is a situation in which the cost associated with the decision or following from the decision is high. In such case, making a large investment relying on a dubious and less than secure promise of an inheritance, seems to be a decision that could potentially do more damage than good.



Figure 8.1.1 Floating Conclusions Debate and the +SC +C condition

Based on our previous assessment, and as sketched in Figure 8.1.1, we have it that the examples used in defence of the direct stance to skeptical inference have as an underlying theme the fact that they not only involve conflicting sources, but they also have a high cost associated to them.

In light of the above, we have a convenient twofold situation. First, and like we previously stated, all of the cases used in support of the direct account of skeptical inference fall in line with +SC and +C value combination hereby sought. Second, and in connection with the aforesaid condition, we can use the cases in support of direct skepticism as a gateway to analyse the acceptance/rejection conditions of floating conclusions under the +SC and +C values. That is, both features provide us with the conditions upon which we can set in motion a single targeted scrutiny

not only to inquire into the acceptance/rejection conditions of floating conclusions under the +SC and +C circumstances but also the acceptance/rejection conditions of the direct approach towards skeptical inference. This provides us with the methodological pivot to address two objectives through a single line of inquiry.

Having stated the above, we consider a further methodological nuance. Throughout the fourth chapter, we reviewed three different cases in support of direct skepticism. Nevertheless, to construct a normative stance towards both floating conclusions and direct skepticism in general, under +SC and +C values, we do not consider case by case to draw general conditions. To the contrary, we take one particular case study as an exemplary situation, which shares important features with the other cases. Thus, through the analysis of this case study, we forward general remarks concerning the acceptance/rejection of floating conclusions under the specific SCCC circumstances and direct skepticism as a whole.

The above settles the preliminary aspects that shape our analysis at this stage. We have presented the type of cases we review, the specific circumstances surrounding such cases and the aim of the scrutiny concerning the direct/indirect skepticism debate and their stance towards floating conclusions.

8.1.2 Analysis

In this section, we consider a particular case, to construct our stance towards the acceptance/rejection of floating conclusions and consequently, our position with respect to the direct/indirect skepticism debate under the +SC and +C conditions.

For this analysis, we review a case proposed by Horty (2002). The problem consists of deciding whether to make a significant investment on a luxury item based on two different, but conflicting sources that agree that a large sum of money from an inheritance will be later available. The aforesaid leaves us room to make the investment at the current time in light of such expectation. Nevertheless, although the involved sources of information agree on the result (one's access to a large sum of money in the near future), they provide conflicting and mutually undermining reasons in support of the availability of the inheritance money.

As we pointed out, there are two crucial reasons why the aforementioned case was used against indirect skepticism. In the context of the direct/indirect skepticism debate, such reasons have a distinctive interplay among themselves.

First, we have that the lines of reasoning are at odds among themselves. More precisely, the brother's claim conflicts with the sister's claim with respect to the possibility of getting the inheritance money. Furthermore, and what we want to stress here, that was not discussed thoroughly in the fifth chapter, is how this conflict is placed in the context of a very sensible investment decision. Thus, it is not just the conflict what seems to push Horty's intuition towards withholding the investment, but rather that such conflict has potentially severe financial consequences. This aspect of the conflict is strongly related to the +C feature of the case study. The aforesaid is one of the core reasons why the conflict, in this case seems to be a strong rather than a weak reason to restrain from taking an indirect approach.

Second, and in line with the above, the case explicitly lacks a commonly shared line of reasoning that sidesteps the conflicting alternative sources. That is, beyond the conflicting routes that converge on the floating conclusion, there is no common and conflict-free route shared among the different sources. This aspect builds on top of the first feature to further move the scale in favour of a directly skeptical approach.

The 'inheritance' case supports the view that the floating conclusion of taking the financial investment should not be given the status of skeptically endorsed information. This was already reviewed in the fifth chapter. Nevertheless, at this stage, we follow the aforementioned thread, which is where the discussion in the literature stops. More precisely, we link the previous discussion with the tools and resources provided by the SCCC, beyond its classification as a +SC and +C case.

On the one hand, it needs to be stressed that we agree on the final stance noted in the literature as to how this case study is indeed a strong case in support for the direct stance towards skeptical inference. That is, we agree that this example provides sufficient conditions that render the floating conclusion as skeptically unacceptable information. On the other hand, and contrary to the assessment provided in the literature, we have a specific framework that grounds our stance as to why the floating conclusion in the context of the aforementioned case is adequately classified as skeptically unacceptable information. In particular, we claim that the specific +SC and +C traits

associated with the case study is what renders the floating conclusion as skeptically unacceptable information.

First, we have a case with alternative but conflicting sources of information (+SC). This sole feature renders the situation as problematic, in the sense that, aside from the fact that we are dealing with floating conclusions, we are dealing first and foremost with conflicting information. Thus, the +SC trait is what makes the situation as one that needs to be treated with a skeptical attitude. For example, the fact that a conclusion needs to be reached in the face of conflicting sources of information signals that such a conclusion needs to be inferred with a skeptical attitude. In such a vein, a careful scrutiny of the sources is paramount since they involve opposing reasons, despite arriving at the same conclusion.

Second, and in line with the above, the +SC feature of the case is not the sole reason that renders the floating conclusion as skeptically unacceptable, but rather the combination of this feature with the +C trait. More precisely, the fact that the cost of the inference process is high, entails that actions based on such information have significant consequences. The aforementioned renders the situation as one demanding an even more skeptical attitude than if it would solely involve conflicting sources. For example, the fact that a decision or endorsement of an inference involves a big financial investment precludes the appeal of the view that, no matter how the conflict is resolved, we end up endorsing the same outcome as an unacceptable view for making the financial investment itself.

One could contest the aforesaid by pointing that this combination of features are peripheral and does not provide grounds to solve the decision problem of accepting or rejecting the floating conclusion. Nevertheless, and contrary to such reaction, we can easily notice how this whole situation differs if the venture to be pursued is sketched as a financially trivial one. For example, let us assume that following the same case, we introduce the sole change that the financial investment Adam is going to make is not a yacht but rather a \$5 lottery ticket. In such a vein, let us assume that, for whatever reasons, Adam wants to use not his own money, but he would rather spend someone else's money on such matter. In this situation, we have the same dilemma inasmuch as the involved sources conflict as to whether Adam is getting half a million dollars of the inheritance money. Nevertheless, the cost associated to the "investment" in this nuanced situation certainly is not a strong enough of a reason as to take the position that here the conclusion which endorses the purchase of a \$5 lottery ticket is not skeptical enough. That is, in this new case, Adam

can still withhold from proceeding with the "investment", but certainly this situation, unlike the previous, would not require such a strong and potentially counterproductive standard.

The aforementioned emphasises that the example proposed by Horty (2002) presents good and compelling reasons to regard direct skepticism as an appropriate approach to floating conclusions. Nevertheless, our position is that it does so not in virtue of an accidental aspect of the argument that somehow renders the floating conclusion as unappealing from a skeptical point of view. Rather, our stance is that the case does so because it combines a salient number of properties that render the floating conclusion as skeptically unacceptable. More precisely, it is that the floating conclusion emerges in a +SC and +C situation that renders it as skeptically unacceptable.

8.1.3 Prescriptive Stance (+SC + C)

In this final subsection, we pick up from what we previously proposed and draw a more general stance detached from the specific nuances associated with the case study we analysed (but keeping in mind the +SC and +C conditions). In particular, we revisit the findings of the previous subsections and provide a concise sketch of our stance with respect to the acceptance/rejection of floating conclusions and the direct/indirect skepticism debate in situations with conflicting sources and a high cost.

First, and as we have pointed out, the direct approach to skeptical inference is a more sensible attitude towards floating conclusions in the circumstances involving conflicting sources and a high cost. The reason is not that floating conclusions are skeptically unacceptable *simpliciter*, but rather that the combination of the +SC and +C attenuating factors favours the direct stance. More precisely, conflicting sources does trigger the need for a conservative attitude towards information extraction. Nevertheless, it is the conjunction of the above with a high cost what pushes the need for an even more conservative form skepticism to properly address floating conclusions. This is precisely what renders the direct stance as the safer option.

Second, and in contrast with the above, the indirect stance towards skeptical inferences under the +SC and +C circumstances fail to have a conservative enough attitude towards floating conclusions. In such a vein, when we have a situation involving conflicting sources and a high cost, floating conclusions cease to be skeptically acceptable information. Under such attenuating

factors, floating conclusions seem not to satisfy the restrained stance towards information extraction and information processing. Given the above, conflicting sources and high-cost conditions render floating conclusions as skeptically unacceptable information, which should lead us to decline the adoption of an indirect stance towards skeptical inference.



Figure 8.1.3 Floating Conclusions Stance and the +SC +C condition

Figure 8.1.3 depicts the full extent of our stance. It points out that we regard floating conclusions as skeptically unacceptable information in situations including conflicting sources and a high cost. This is the basis of our endorsement of the direct approach to skeptical inference, *under these specific circumstances of high cost*.

8.2 Conflicting Sources with Low Cost

In this section, we analyse the problem of floating conclusions in the context of both direct and indirect approaches to skeptical inferences, including the conflicting sources (+SC) and low cost (-C) conditions. This section aims to provide a detailed assessment of the cases discussed in the literature that satisfy the specific conditions of conflicting sources and low-cost. Furthermore, we inquire into such low-cost cases involving conflicting sources using the SCCC. Moreover, we aim

to provide a prescriptive stance for the acceptance/rejection of floating conclusions under the aforesaid conditions.

8.2.1 Cases

In this section, we look into the cases discussed in the fifth chapter involving conflicting sources but with low or no cost. The above helps us to determine a selection of cases that suits the +SC and -C conditions of our current inquiry. Based on such cases, we address the question of the acceptance/rejection of floating conclusions under such circumstances.

In our analysis of the arguments supporting the indirect version of skeptical inference, we used three case studies. First, we reviewed the case proposed by Ginsberg (1987) in which the problem revolved around whether to ascribe the attribute of being an extremist to Nixon based on two conflicting, yet converging sources (*extremist* case). Second, we inquired into the situation proposed by Prakken (2002) in which we needed to determine if we could ascribe someone the hobby of ice skating based on him being Norwegian or him being Dutch, despite that both streams of information conflicted (*skating* case). Third, we discussed the case proposed by Stein (1992), in which the question was whether we could infer if a seedless grape vine is a plant despite two conflicting sources ultimately agreeing in such matter (*seedless* case). These cases share various important similarities, which we proceed to enumerate.

First, and like in the case of direct skepticism, all of the examples used in support of the indirect account of skeptical inference showed an explicit level of conflict among the lines of reasoning. Nevertheless, this trait alone, like in the context of direct skepticism, did not serve as a sole reason for accepting or rejecting the conclusion. The existence of conflicting sources paved the way towards the stance that further scrutiny needed to be undertaken to determine whether the floating conclusion was skeptically acceptable. For example, the extremist and skating cases both involve sources of information that, beyond agreeing on the conclusion, they do instantiate situations in which the arguments for the conclusions conflict among themselves. In the extremist case, the fact that Nixon is said to be a hawk stands in direct opposition of Nixon also being a dove. In the skating case, Brygt Rykkje cannot be both Norwegian and Dutch.

Second, and in contrast with direct skepticism, throughout these cases, the absence of a non-

conflicting line of reasoning common to all sources of information did not serve as a strong enough reason to state that the conclusions are not skeptically acceptable information. That is, contrary to the cases supporting the directly skeptical stance; these cases did not show that the absence of such commonly shared argument was a conclusive reason to hold that the floating conclusion was not skeptically acceptable. For example, following the seedless case, although the sources conflict, it is clear that the absence of a common and non-conflicting path in which both sources coincide, does not provide a good enough reason to restrain from skeptically accepting and endorsing that a seedless grape vine is a plant.

Third, and also in contrast with the cases reviewed in support of direct skepticism, the situations discussed throughout the literature and used in support of the indirect stance, showed a significant and conspicuous absence of cost surrounding the situation (in the sense of the consequences associated to such decision). For example, in the skating case, there is a low cost associated with the acceptance of the floating conclusion. That is, all things considered, ascribing the hobby of ice-skating to Brygt Rykkje, seems *prima facie*, a low-cost decision.



Figure 8.2.1 Floating Conclusions Debate and the +SC -C condition

In light of the previous discussion, and as systematised in Figure 8.2.1 all of the cases used in support of indirect skepticism have as a common thread being situations of conflicting sources,

but with a low or no cost associated to them.

Thus, like in the previous section, we have a convenient twofold situation. First, and as previously noted, all of the cases used in favour of indirect skepticism comply with the specific values of the SC and C addressed in this section. Second, and in connection with the previous feature, here we can use the cases favouring indirect skepticism as a proxy to assess the conditions upon which we can accept or reject floating conclusions under the +SC and -C circumstances.

Mirroring the strategy of the previous section, both features not only provide the necessary conditions upon which we can take one line of inquiry but more importantly, it facilitates two important components. First, it provides us with the elements to determine when and why floating conclusions are skeptically justifiable under the +SC and -C circumstances. Second, it offers us the conditions under which the indirect stance towards skeptical inference is a skeptically acceptable strategy. Again, like in the previous section, one single line of inquiry facilitates the conditions for the analysis of the acceptance/rejection conditions of floating conclusions under the +SC and -C value combination as well as the acceptance/rejection conditions of floating conclusions of floating conclusions for the indirect account of skeptical inference.

In line with the above, we follow the same methodological approach as in the previous section. In particular, we do not go through each of the cases reviewed in the fifth chapter, but rather we analyse a single case to construct a general stance towards floating conclusions involving conflicting sources but with low cost. That is, we construct our position by taking one case study, which shares the most important features that all of the cases used in support of indirect skepticism and floating conclusions have. Through the analysis of a single but exemplary case, we construct a general stance towards the acceptance/rejection of floating conclusions in cases of +SC and -C and consequently of indirect skepticism in general.

The above concludes the specification of all the surrounding circumstances that shape our inquiry for the current section. More precisely, we discussed the cases we are considering as well as the underlying reasons that make such cases the target of this section. Furthermore, we have noted how these cases stand concerning the direct/indirect skepticism debate and the acceptance/rejection of floating conclusions. In the following subsection, we go through our analysis of a specific example to sketch our position towards floating conclusions and the direct/indirect skepticism debate.

8.2.2 Analysis

In this section, we analyse a case from the various examples reviewed throughout chapter six to set out our position concerning the acceptance/rejection of floating conclusions and the direct/indirect skepticism debate. To conduct our analysis, we use the case presented by Prakken (2002) in which we face the decision of whether to ascribe the hobby of ice-skating to a particular individual, Brygt Rykkje, based on some known background information of him. In particular, the predicament is produced because Brygt Rykkje is said to have the hobby of ice-skating based on him being Norwegian due to his name, but at the same time he is said to have the interest in ice skating based on him being Dutch due to his birthplace, namely Holland. In this example, we have two lines of reasoning converging in Brygt Rykkje liking ice skating, but providing incompatible reasons to endorse such a conclusion.

As we have previously pointed out, there are two crucial reasons why this situation was used as a case against direct skepticism. Furthermore, and like in the previous section, the two reasons taken into consideration interact among themselves in such a way that favours the indirect version of skeptical inference.

First, and like in the example of the previous section, in this situation, the different lines of reasoning are at odds with each other. In particular, the information in each line of reasoning conflicts with the information contained in the alternative route, with respect to Brygt Rykkje's origin and background. Nevertheless, and very much unlike the case assessed in the previous section, here the cost of the decision associated with the conflict seems to be of low (-C). That is, in the context of the debate over Brygt Rykkje like or dislike for ice-skating there seems to be no reason at all why such subject matter could be of a high cost. In such a vein, the conflict in this situation seems to be a rather *weak* reason to dismiss the floating conclusion. Furthermore, the lack of cost attached to the situation seems to play a crucial role in the acceptance of the floating conclusion.

Second, and in direct connection with the above, this case, like the one discussed in the previous section, lacks a commonly shared non-conflicting route among the different lines of reasoning. That is, the account provided by the two different sources of information lacks, beyond the conflicting and already discussed opposing lines of reasoning, a commonly shared route in support

for the floating conclusion. Nevertheless, given that the absence of a commonly shared route happens in the context of a situation of low cost, it provides not a strong enough reason, but rather a *weak* reason to decline the floating conclusion. This specific aspect placed in connection with the fact that the floating conclusion is already shared in every line of reasoning provided by the various sources of information sets a strong case in favour of accepting rather than rejecting such conclusion.

Thus, in this case, the floating conclusion that ascribes to Brygt Rykkje a fondness for ice-skating seems to qualify as skeptically acceptable information. That is, no matter how the discrepancy among sources is resolved, it certainly and unproblematically renders the same conclusion. Such was the force and reason behind this case study. This was already addressed in the fifth chapter, and at such point, the analysis stopped. Here we want to continue with the aid of the SCCC, which we employ beyond merely classifying the case as having a source conflict, yet a low cost.

On the one hand, we coincide with the stance taken in the literature concerning the case. In particular, we agree that the situation does provide good enough reasons to concede that the floating conclusion under these circumstances is skeptically acceptable information. On the other hand, and unlike the review such case received in the literature, we employ the SCCC framework to justify why the floating conclusion should be classified as skeptically acceptable information. More precisely, we claim that the particular +SC and -C features associated with the case study is what renders the floating conclusion as skeptically acceptable.

First, and like in the previous section, here we have a case with independently converging, but mutually conflicting sources of information (+SC). This feature produces the need to adopt a conservative attitude towards information processing. For example, Brygt Rykkje being Norwegian and born in Holland cannot, *prima facie* and non-trivially coexist. The fact that the lines of reasoning point to Brygt Rykkje interest for ice-skating should produce a careful, thorough and skeptical attitude towards the sources themselves. This coincides with the underlying intuition shared among the two forms of skeptical inference.

Second, here the +SC trait operates in conjunction with the other SCCC property, namely the -C condition associated with this case. As we know, and unlike the example of the previous section, here the case is classified as a low-cost situation. Such low-cost property renders the floating conclusion as skeptically acceptable information. More precisely, since there is a low cost attached

to the deliberation as to whether Brygt Rykkje likes ice skating, such conclusion can be taken as good enough information, *even evaluated through a skeptical standard*. For example, we know that the sources of information conflict with one another. Nevertheless, and regardless of which line of reasoning we choose, we can rest assured that it leads to the same conclusion. This taken in conjunction with the fact that there is little cost associated with the information entailed by the conclusion, provides us with a reason to take it as a skeptically enough conclusion. That is, the sole condition of conflicting sources fails to render the floating conclusion as skeptically unacceptable. In such a vein, the extent to which we are prone to impose restraints on inferences produced by +SC and +C conditions is conspicuously different in cases in which we solely rely on +SC to trigger those high skeptical standards.

One could contest such a stance due to the fact that there being a low cost associated to ascribing the interest of ice skating to Brygt Rykkje is not a good enough reason to render such floating conclusion as skeptically acceptable. Nevertheless, and as we pointed out in the previous section, it was precisely the conjunction of conflicting sources with the significant cost of a financial investment that made us restrain ourselves from endorsing information that would make us act like we would inherit a large sum of money. This situation takes the same thread, but states that the absence of a high cost renders the floating conclusion as skeptically acceptable information. That is, it should be taken as information that withstands a sufficient and reasonable amount of scrutiny. In our case, the acceptance of the floating conclusion is based on the stance that despite the existence of a conflict within the sources, we would not, in being information-hungry, be self-undermining to any serious extent. In the same vein, the indirect version of skeptical inference is already a skeptical attitude towards information extraction, since it only admits information that is contained in every single line of reasoning. Thus, this version of skeptical inference is far from a credulous attitude. Given the above, and in the face of cases such as the one addressed in this subsection, the indirect stance can be taken as a sufficiently skeptical reasoning policy.

Opposing the aforesaid stance would imply moving towards a potentially undesired and impractical threshold of skeptically acceptable inferences. Such impractical threshold would deter us from inferring that Brygt Rykkje likes ice-skating just because on the one hand he might be Norwegian due to his Norwegian name, but on the other, he might be Dutch due to his birthplace, while knowing that no matter how the conflict is resolved, we end up with the same conclusion, in conjunction with the low-cost nature of the situation. The aforementioned would amount to taking the stance that a conclusion being contained in every single line of reasoning is not a skeptical enough attitude, which a case like Brygt Rykkje's interest for ice-skating seems to contravene as we have argued.

Thus, our stance is that, like the previous section, the +SC condition alone fails to render the floating conclusion as skeptically unacceptable information. Moreover, the combination of conflicting sources with a low-cost situation provides sufficient grounds to justify the admission of the floating conclusion, even by a skeptical standard. In this sense, not only the ice-skating example suggested by Prakken (2002) provides a good case in favour of indirect skepticism, but it does so in virtue of the +SC and -C traits of the case study. In such a vein, we agree that in situations with such features, the acceptance of floating conclusions withstands the threshold of a skeptical stance towards information processing. Based on the above, the indirect stance towards skeptical inference is to be regarded as a better approach.

8.2.3 Prescriptive Stance (+SC -C)

In this final subsection, we revisit the main ideas developed in the previous subsections removing ourselves from the specific details and circumstances of the case study we inquired into and proceeding to a general outlook over the nature of debate and problem itself. The above being said, we still operate under the SCCC framework in which we have conflicting sources, but with low cost associated with the situation. We construct our stance towards the acceptance/rejection of floating conclusions and the direct/indirect skepticism debate under the aforementioned circumstance of the SCCC we analysed.

First, and as previously noted, the indirect approach towards floating conclusions is better in +SC but -C cases. Furthermore, we do not hold the view that this should be read as floating conclusions being skeptically acceptable information *per se*. It is the combination of conflicting sources (+SC) with low cost (-C) that renders floating conclusions as skeptically acceptable, favouring the indirect stance in this case. We agree with the view that the presence of conflicting sources does produce the need for a restrained and conservative attitude towards inferences made under such circumstances. Nevertheless, the low cost associated with such situations entails that the admission of floating conclusions does not transgress the restrained attitude contained in the indirect stance towards information. That is, the indirect version of skeptical inference is already a skeptical attitude towards inferences and information processing that, in the face of low-cost

situations renders itself as a skeptical enough stance.

Second, the direct stance towards skeptical inferences in the context of conflicting sources, but with low cost associated to them falls short of endorsing conclusions that, as we have extensively argued seem reasonable and acceptable enough *under the aforementioned SCCC circumstances*. That is, in +SC but -C situations, floating conclusions are to be held skeptically acceptable information, since they already satisfy one particular form of skeptical scrutiny, namely the indirect criteria of skeptical scrutiny which, as we have argued, is a satisfactory threshold in such circumstances. The aforementioned renders the direct approach as ill-equipped to account for inferences with such attenuating factors. In such a vein, at this stage and under the aforesaid specifications, the direct version of skeptical inference should be declined in favour of the indirect stance.



Figure 8.2.3 Floating Conclusions Stance and the +SC -C condition

Figure 8.2.3 summarises our stance. More precisely, it highlights that floating conclusions ought to be regarded as skeptically acceptable information in cases involving conflicting sources, but with a low cost associated with these situations. The above naturally leads to upholding the view that the indirect rather the direct version to skeptical inference should be taken *in cases with the previously specified circumstances*.

8.3 Non-Conflicting Sources with High Cost

In this subsection, we consider cases of non-conflicting sources but with a high cost to assess the acceptance/rejection conditions of floating conclusions under such circumstances. That is, we inquire into -SC and +C cases involving floating conclusions. This section aims to provide a detailed assessment of cases involving non-conflicting sources and a high cost, using the SCCC framework. Lastly, we provide a normative account for the acceptance/rejection of floating conclusions under the aforesaid conditions.

8.3.1 Case Study

In the last two sections, we used the SCCC to assess the cases discussed in the literature. As we extensively reviewed throughout the last section of the fifth chapter, all cases involved conflicting sources. That is, all cases were classified into the +SC category. Thus, we cannot rely on cases discussed in the literature to examine the -SC situations as there are none. Therefore, in this subsection, we construct a specific case that complies with the non-conflicting sources, but high cost conditions that we are interested in assessing.

Based on the case of our own making, we inquire into two specific aspects. First, we assess whether the floating conclusion involved in the case study is skeptically acceptable information in light of the required SCCC conditions. Second, we analyse whether such a case, with the previously noted features, is best accounted by either the direct or the indirect version of skeptical inference. This approach provides us with a strategy to conduct our inquiry in the absence in the literature of a specific case study with the desired traits.

To address the task at hand, we build a case involving a medical decision. We imagine a given patient, which we refer to as P_1 seeks medical care at a given facility. In such vein, we assume P_1 is seen by three different practitioners, which we denote as D_1 , D_2 and D_3 . Lastly, we also work on the assumption that P_1 seeks medical care due to a life-threatening condition, which we refer to as C_1 . In the aforementioned context, the three practitioners individually and independently endorse the prescription of X_1 for the purposes and needs of P_1 with respect to C_1 . In light of the above, we have the following prognoses:

- D_1 : To address C_1 , P_1 needs to take X_1 , because it will reduce and alleviate the distress to the liver which has been placed under industrious work.
- D_2 : To address C_1 , P_1 needs to take X_1 , because it will increase his mobility which will facilitate engagement in physical activity which will help accelerate recovery.
- D_3 : To address C_1 , P_1 needs to take X_1 , because it will facilitate releasing a vital neurotransmitter which is currently blocked.

Having provided the basic details surrounding our case study we review why it complies with the properties we are looking for. First, and perhaps the clearest property of the example is that it involves non-conflicting sources. More precisely, although the lines of reasoning pursued by each practitioner are, *prima facie*, unrelated, they do not conflict with each other. For example, despite the lines of reasoning pursued by D_1 and D_3 are unrelated, they do not undermine each other. The same can be said about D_1 and D_3 and the other combinations. Second, and concerning the cost associated with the situation, it is safe to state that it renders itself as a high-cost case. Given that P_1 's condition is life-threatening, the consequence associated with the decision-making process, in the sense of the consequences, should be considered of a very high cost. In light of both of the previous aspects, we can proceed knowing that our particular case study does comply with the - SC and +C conditions that concern us in this section. That is, the example adheres as a case with non-conflicting sources and a high cost.

Having stated the above, the example, not only presents itself as a -SC and +C situation, but more importantly, it is a case involving a floating conclusion. More precisely, all of the *prima facie*, unrelated and independent lines of reasoning, point to the same conclusion. In particular, D_1 , D_2 and D_3 all converge on the idea that P_1 ought to take medication X_1 to address life-threatening condition C_1 . First, D_1 grounds his expert judgment on X_1 's positive effect on the liver, which based on the condition C_1 has been put through arduous work and might accelerate, unless reversed, the progression of the condition. Second, D_2 grounds his medical judgment on X_1 's role in facilitating P_1 's physical mobility allowing him to engage in more physical activity, which will help P_1 fight back C_1 . Third, D_3 supports his recommendation of P_1 taking X_1 based on what he knows will facilitate the release of an essential neurotransmitter which will address P_1 's condition. Thus, no matter which route one chooses, it is clear that all of the involved medical experts converge on the fact that given P_1 's medical condition C_1 , it is imperative for P_1 to start taking X_1 . In such a vein, the conclusion that P_1 needs to take X_1 to address C_1 is a floating conclusion since

it is endorsed by all of the medical experts, although it is only supported through different and unrelated routes of reasoning.

8.3.2 Analysis

In this section, we inquire into the case study we have constructed in the preceding subsection to ground our stance with respect to the acceptance/rejection of floating conclusions, under cases involving non-conflicting sources, yet high-cost conditions.

First of all, and to briefly recall the overall thread of the case under scrutiny, we have three different medical practitioners who all independently and separately agree on a patient's need to take a given medication based on a life-threatening condition. Furthermore, all of the practitioners agreed on the patient's mandatory requirement to take a specific medication, but only in virtue of different lines of reasoning. Thus, we have a case involving a floating conclusion since such a conclusion is supported by every line of reasoning, but not by a common argument shared by every source of information. Moreover, the involved lines of reasoning are not in conflict since they do not oppose or mutually undermine each other. Finally, we noted that the situation itself is of a high cost, given that the medical condition of the patient is life-threatening.

Having provided a brief outlook of the main aspects of the case under scrutiny, we present our stance favouring an indirect account of skeptical inference rather a direct approach. To support the aforesaid, we use the SCCC not only as a classification tool, but rather as a framework for understanding the attenuating circumstances that shape the floating conclusion acceptance conditions. More precisely, we argue in favour of the view that the particular -SC and +C features are what renders the floating conclusion as skeptically acceptable information. There are two important reasons as to why we take such a position.

First, and as we have pointed out, the situation involves a high cost attached to the decision-making process due to its high-risk nature. Furthermore, the situation also provides enough information to state that the different lines of reasoning are non-conflicting sources of information. Despite the cost associated with the case study being high (i.e., +C), we have an absence of conflict among the sources (i.e., -SC). Thus, even though the decision-making process is a high-cost process, we have the certainty that the sources or lines of reasoning are not at odds among them. In light of the

above, we have a strong enough reason to support the admission of the involved floating conclusion, even by a conservative standard. More precisely, given that there is no conflict whatsoever associated to the situation, the fact that all of the involved practitioners agree on the view that the P_1 should take X_1 to address C_1 provides sufficient reasons to uphold such conclusion as acceptable information, even when evaluated through a skeptical standard. To put it in another way, the sole trait of a high cost attached to the example would be a rather *weak* reason to dismiss the floating conclusion given that the different lines of reasoning support the floating conclusion without undermining each other. Therefore, the extent to which we restrain from information extraction such as with direct approaches would seem to be counterproductive given that the -SC serves as a strong reason to consider floating conclusions as skeptically acceptable information.

Second, the example lacks a commonly shared non-conflicting line of reasoning. There is no single argument supporting the need for the patient to take the same medication that is common to all of the practitioners since all of them endorse the same conclusion, but only through different arguments. Nevertheless, the absence of such property cannot be held as an undermining reason in detriment of the floating conclusion given that the various lines of reasoning are not at odds with each other. That is, the presence of conflicting sources would provide a legitimate reason to consider elevating the requirement of a unique and common shared argument in high-cost cases. Nevertheless, such circumstances are not present in the case under scrutiny.

Based on the previous two aspects, the floating conclusion that supports the patient taking the medication endorsed by all of the involved practitioners should be accepted rather than dismissed. Furthermore, it can be held as an inference drawn by restrained and sufficiently skeptical standards. Given that the various sources are not at odds among each other, the sole existence of a high cost attached to the case or the lack of a commonly shared path seems too weak of a reason to dismiss the floating conclusion. Having stated the above, we take the stance that this particular case study does provide supporting evidence in favour of the indirect approach towards floating conclusions.

8.3.3 Prescriptive Stance (-SC + C)

In this subsection, we merge the ideas developed throughout the previous subsection. More precisely, we look into the stance we have constructed towards floating conclusions and the direct/indirect skepticism debate in cases involving non-conflicting sources but of a high cost.

First, we have that the indirect approach towards floating conclusions is the best account of skeptical inference in -SC and +C conditions. In such a vein, and as we noted in the previous sections, this is not to be taken as a position stating that floating conclusions are to be held as skeptically acceptable information *simpliciter*. By contrast, our stance is that the specific combination of high-cost conditions (+C) but no-conflicting sources (-SC) makes floating conclusions skeptically acceptable. This ultimately renders the indirect stance as a more appropriate account of skeptical inference under such circumstances. More precisely, the fact that such situations are of a high cost does trigger the need for a restrained and conservative attitude towards inferences made under such circumstances. Nevertheless, and despite there being a high cost associated with such cases, these situations involve non-conflicting sources. This last feature is what enables the admission of floating conclusions, in such a way that it does not transgress the skeptical benchmark that is already entailed in the indirect stance towards inferences and information extraction such that, in the face of non-conflicting sources, it is skeptical enough to admit floating conclusions.

Second, and concerning the direct account of skeptical inference in the context of high-cost situations but with non-conflicting sources, we claim that such approach fails to endorse conclusions that, as we have thoroughly argued, seems more than reasonable to take as legitimate, *under the aforementioned SCCC circumstances*. That is, in situations with the -SC and +C traits, floating conclusions ought to be taken as skeptically acceptable information, since they already comply with one particular form of skeptical scrutiny, namely the indirect criteria of skeptical scrutiny which as we have argued is a satisfactory mechanism in such circumstances. In light of the above, the direct approach is unsuited to account for inferences involving such attenuating factors, and therefore, it should be declined in favour of the direct stance.



Figure 8.3.3 Floating Conclusions Stance and the -SC +C condition

Figure 8.3.3 depicts our stance. It highlights that floating conclusions should be regarded as skeptically acceptable information in high-cost situations, which involve non-conflicting sources. This naturally leads to the stance that the indirect rather the direct version to skeptical inference should be taken *in cases with the previously specified circumstances*.

8.4 Non-Conflicting Sources with Low Cost

In this section, we consider cases involving non-conflicting sources with low cost. We analyse the acceptance/rejection of floating conclusions under such circumstances. More precisely, we address the last of the SC and C value combination that we have been assessing throughout the previous subsections. This section has three objectives. First, we aim to consider cases of non-conflicting sources and a low cost. Second, and more importantly, we aim to analyse these low-cost circumstances involving non-conflicting sources using the SCCC as a framework. Third, we want to provide a normative account for the acceptance/rejection of floating conclusions under the aforesaid conditions.

8.4.1 Case Study

Here we construct a particular case study to conduct our assessment and present our stance concerning the problem of floating conclusions. In particular, we build a case study which involves non-conflicting sources (-SC) and low-cost conditions (-C). In such a vein, and instead of constructing a case from scratch, we draw upon the case we previously considered in the preceding section and modify it to fit the purposes of this section.

Like in the previous section, we imagine a patient, which we refer to as P_2 , seeking medical care at a given facility. Likewise, we also assume that the aforementioned patient is seen by three different practitioners, referred to as D_1 , D_2 and D_3 . Nevertheless, and unlike in the previous section, here P_2 does not have a potentially life-threatening condition, but rather she wants to enhance her overall physical fitness for optimising her performance at a recreational sports league in which she is leisurely involved, but has not been enjoying because of a mild level of underperforming in contrast with her peers of a similar age range, health and lifestyle conditions. In such a vein, we refer to this lack of physical fitness for the aforementioned purposes as C_2 . In light of the above, each practitioner has individually and independently considered the prescription of X_2 for the purposes and needs of P_2 for C_2 . In this context, we assume we have the following prognoses:

D_1 :	To address C_2 , P_2 needs to take X_2 , because it will enhance
	her cardiovascular capacity.

- D_2 : To address C_2 , P_2 needs to take X_2 , because it will enhance her stamina and endurance.
- *D*₃: To address C_2 , P_2 needs to take X_2 , because it will facilitate muscular recovery.

In light of the above, we now have the basic details that compose our case study. Furthermore, we have constructed a case that complies with the -SC and -C values we are pursuing at this stage. We immediately proceed to discuss why this case suits our needs and how it complies with the -SC and -C conditions.

First, and like in the previous section, one of the clearest traits is that the case involves nonconflicting sources. That is, all lines of reasoning, independently pursued by each practitioner are not at odds with each other. For example, the practitioner who claims that P_2 needs to take X_2 because it enhances P_2 's cardiovascular capacity might not be aware at all about X_2 's effect on a person's stamina and endurance. Nevertheless, and although potentially unrelated, they certainly do not undermine each other. Thus, we can say that this situation is one containing non-conflicting sources.

Second, and concerning the cost associated with the example, we can safely say that this is the point in which this version disagrees with the one presented in the previous section. That is, this situation seems to be a case of low cost. In particular, given that P_2 's "condition" is one that only alludes to her desire to enhance general physical fitness, within the context of recreational sports, we can very safely assume, all things considered, such medical concerns as being of low cost. That is, such condition does not represent a life-threatening risk to her. Thus, her concern being just a social/recreational inquiry, although legitimate, is certainly, all things considered, a low cost situation.

In light of the previous points, we can rest assured that the case we have constructed, complies with the SC and C features we are looking for the purposes of this section. In particular, this case study presents us with the conditions of being a situation involving non-conflicting sources and low-cost.

Proceeding to the key aspect that underlies our inquiry at this stage, it is important to mention that this case not only involves non-conflicting sources and low-cost circumstances, but it also contains a floating conclusion. That is, the previous example depicts the case in which the same conclusion is reached by all lines of reasoning, however they do not have a commonly shared argument.

In particular, and like in the previous case study, in this medical facility, all of the involved practitioners, D_1 , D_2 and D_3 converge on what this new patient ought to take to address her condition. In particular, they all agreed that P_2 ought to take medication X_2 to address C_2 . First, D_1 grounds his expert judgment on X_2 's effect on the cardiovascular system, which will help enhance P_2 's performance by lasting longer with less trouble with her physical activity. Second, D_2 bases his stance on X_2 's known properties in high performing athletes for increasing overall stamina and endurance, which will certainly help a recreational sportsperson like P_2 with a mild case of

substandard performance. Finally, D_3 grounds his judgment on the basis that X_2 will facilitate muscular recovery, which will reduce P_2 's experience of fatigue and will enhance her ability to engage more frequently in the physical activities organised by the recreational league which will help her overall fitness increase. Thus, regardless of which expert opinion we are inclined to take, we end up endorsing that P_2 ought to take X_2 to address her condition. That is, all medical experts separately and independently reach the same conclusion that P_2 should take X_2 to address C_2 , but only through different lines of reasoning. In light of the above, the stance that P_2 should take X_2 in order to address C_2 is to be regarded as a floating conclusion, since such conclusion is supported by each of the medical experts, but through separate and unrelated lines of reasoning, and not through a commonly shared argument.

8.4.2 Analysis

In this subsection, we look into the case study previously constructed to ground our position concerning the acceptance/rejection of floating conclusions and the direct/indirect skepticism debate, under the -SC and -C circumstances. To address the aforementioned, we recall the main aspects of our working example.

In the case under consideration, we have three different medical practitioners, which all agree on a given patient's need to take a specific medication not based on an urgent or life-threatening condition, but rather on the patient's desire to enhance her performance in the context of socially engaged sports. Furthermore, each of the three practitioners agree on the patient's need for the same specific medication following separate and independent lines of reasoning. Thus, we have a situation involving floating conclusions given that the same conclusion is reached by every source of information, but only through different lines of reasoning rather than through an argument shared by all of the involved sources. Moreover, this case study is composed of lines of reasoning that do not conflict among each other, and it can be classified as low cost, given that the medical condition of the patient is not life-threatening.

Here we persist, despite the modification of the case study, and state that this new version still favours an indirect approach towards skeptical inference rather than a direct approach. To support our stance, we make use of the SCCC not only as a classification tool, but as a framework of analysis to understand the features surrounding such cases and shaping the acceptance conditions

of floating conclusions that consequently favours the indirect approach to skeptical inference. In such a vein, we claim that the -SC and -C features are precisely what renders the aforementioned situations as better addressed by indirect skepticism since this version of skeptical inference justifiably endorses floating conclusions as skeptically acceptable information. Like in the previous section, we focus on two key features as to why we adopt such a stance.

First, and as with the case of the previous section, here we also have a situation in which the various lines of reasoning do not conflict among each other. Furthermore, not only do the sources not conflict, but we also have it that the case is of low cost. In light of the above, we have even more than enough reasons to take the stance that the floating conclusion should be admitted, and by doing so, still comply with a strict attitude towards information extraction. Moreover, we have a far stronger position here than in the previous case, which despite having non-conflicting sources, it was of high cost. That is, if our argument in the previous conditions holds, now it holds even stronger when the high-cost condition is swapped with a low-cost condition. In support of the above, neither of the potentially worrying features such as a high-cost or conflicting source are present. Thus, since none of the practitioners disagrees with any other, nor do they undermine each other, in conjunction with the fact that the situation itself has a low cost associated with it as it is not a life-threatening situation, we can assume that the floating conclusion that stipulates that P_2 should take X_2 to treat her C_2 condition can be taken as good enough information, even evaluated through a skeptical standard.

Second, like in the previous section, this case lacks a commonly shared argument. That is, there is not a single line of reasoning supporting the need for the patient to take the same medication that is prescribed by all of the involved practitioners. To put it another way, all the practitioners reach the same conclusion, but only through different arguments. Nevertheless, given that, the involved non-commonly shared sources are non-conflicting sources, such trait cannot be held as a reason to take the position that the floating conclusion should be dismissed. Likewise, if we follow the argumentative thread, we note that this trait was a weak and non-conclusive reason towards the dismissal of floating conclusions even when non-conflicting sources involved a high cost condition. Thus, here it should be even a stronger case as to why this same reason is still a weak reason to dismiss the floating conclusion when not only the sources are not mutually undermining but also the cost condition is low.

Based on the previous two aspects, the floating conclusion that the patient should take the

medication endorsed by all of the practitioners should be taken as a sufficiently skeptical conclusion. In particular, given that not only the lines of reasoning are not at odds among each other, but there is no cost attached to the case, we have a stronger argument in favour of endorsing the floating conclusion. Furthermore, the absence of a common line of reasoning shared among all of the practitioners is a weak reason taken into consideration the previous aspects of the case. In such a vein, we argue that this case supports and serves as evidence in favour of the indirect approach towards floating conclusions.

8.4.3 Prescriptive Stance (-SC -C)

In this third and final subsection, we revise the main theses of the previous subsection in such a way that we shift from the particular circumstances of the analysed case to provide a more general outlook. Furthermore, we recall that we are considering non-conflicting and low-cost situations. We summarise our stance towards floating conclusions and the direct/indirect skepticism debate, based on the aforementioned features.

First, and based on our assessment, the indirect stance towards floating conclusions is a better approach to skeptical inferences in -SC and -C conditions. Nevertheless, it is important to emphatically state that the aforementioned needs not to be taken as a position stating that floating conclusions are to be held as skeptically acceptable information simply in virtue of themselves. To the contrary, the key aspects that ground our stance in the debate revolves around the specific combination of non-conflicting sources (-SC) and low-cost (-C) conditions. It is based on the aforesaid SC and C value combination that we take floating conclusions as skeptically acceptable. This, in turn, leads us to take the indirect stance as a more appropriate account of skeptical inference in this type of situations, given that such an approach to skeptical inference is already a sufficiently skeptical account of inferences and information extraction.

Second, and concerning the direct account of skeptical inference in cases with non-conflicting sources and low-cost conditions, we take the stance that such an approach fails to endorse conclusions that, as we have argued, seem reasonable to take as legitimate even by a conservative and skeptical standard, *under the aforementioned SCCC circumstances*. That is, in cases with the -SC and -C traits, floating conclusions are skeptically acceptable information. In particular, floating conclusions need to be treated as skeptically acceptable since they already follow one
particular mechanism of skeptical evaluation, namely indirect skepticism which as we have argued is a sufficiently skeptical mechanism in the aforesaid circumstances. The above considerations render the direct approach as unsuited to account for inferences with such features.



Figure 8.4.3 Floating Conclusions Stance and the -SC -C condition

Figure 8.4.3 depicts the stance we have constructed throughout this section. More precisely, it highlights that floating conclusions need to be held as skeptically acceptable information in cases involving non-conflicting sources and a low cost. The aforesaid leads us to take the indirect rather the direct version to skeptical inference *in cases with the previously specified circumstances*.

8.5 SCCC and Inheritance Networks

Through this chapter, we have evaluated the acceptance or rejection of floating conclusions through the SCCC. We have inquired into the circumstances where floating conclusions seem to satisfy a sufficient skeptical standard. Nevertheless, we have not reflected on the extent to which the SCCC could be combined with the formalism described in the first chapter. Here we reflect on the feasibility of implementing the SCCC into a logical formalism such as inheritance networks.

As a preliminary remark, it is important to highlight that the Source Conflict Cost Criterion operates with semantic information, in the sense of Floridi (2019). Thus, notions such as cost and conflict as well as skeptically acceptable conclusions do not belong to the same level of analysis as nodes, paths and extensions. As such, the SCCC as a decision matrix is located at a different level of analysis than Inheritance Networks as a logical formalism. Nevertheless, we will reflect various ways in which both the SCCC and a non-monotonic formalism can be combined.

Although it is tempting to draw an immediate connection of the SCCC with Inheritance Networks insofar as this was the chosen formalism to explore the problem of floating conclusions, we wish to step back and assess the connection between our framework and non-monotonic logic and non-monotonic reasoning in general.

The first point to be made is that there seems to be a natural progression between what has been done in the context of non-monotonic logics. As evidence of this claim we just need to go back to the very core problem that produced floating conclusions, i.e. the problem of deciding what to infer in the context of a network (or defeasible theory in general) when one can draw with equal legitimacy (two unrelated but viable lines of reasoning) opposing or conflicting information. This was the foundational problem upon which floating conclusion later emerges.

As we have extensively mentioned, the two major solutions to such a problem are the credulous and skeptical approaches. That is, we either draw as much as possible, even at the expense of conflicting conclusions or we draw the absolute minimum to avoid conflicting information. This much we already know, but the point to be made is that this problem itself and its proposed solutions already reveals an important connection with our approach to the more complex problem of floating conclusions.

More precisely, both the credulous and skeptical approaches are very rudimentary and crude reasoning policies to address the very simple problem represented by situations like the Nixon Diamond. That is, both attempts to solve the problem rely on a general and abstract intuition that answers the question of what conclusion is one warranted to draw in the face of conflicting lines of reasoning. Both skeptical and credulous strategies attempt to say something relevant regarding this question, and they do so through the use of a very general rule of thumb which is later translated back into the logical formalism as a reasoning policy.

In the same vein, the SCCC resembles both the credulous and skeptical reasoning policy in being exactly that, a general and abstract strategy that attempts to answer a question within a logical formalism, i.e. what conclusions should we draw when different lines of reasoning all support the same conclusion despite not having a commonly shared argument in support to the conclusion. What we have constructed is certainly more complex than the simpler credulous and skeptical reasoning policies, but the question it answers is equally more complex, but only as a matter of degree rather than principle. Nevertheless, and despite differences, the SCCC has with the credulous and skeptical reasoning policies, they share the same underlying principle of attempting to construct a general and abstract reasoning policy for a problem that emerges in a non-monotonic formalism. As such, there is an important continuity which connects our work with the type of work that has been done in the field.

Despite the resemblance in principle between our work and other solutions to the simpler problem in non-monotonic logics, it is without question that our solution is far more nuanced than what has been done previously to the simpler problems. In particular, our approach to the question raised by floating conclusions takes into account not only internal properties of the logical formalism, but it takes into account cues from the environment in which the problem represented by the problem might elicit. For example, the two central features of cost and source conflict not only relates to the internal structural arrangement of a network, but it also concerns themselves with the environment in which the problem represented occurs and the assessment a cognitive agent makes of such environment and how it relates to the problem represented by the logical formalism.

The above suggests that there is an additional dimension taken into account by our approach to floating conclusions, or at least there are an important series of factors and features which are at play in our understanding of such phenomenon. In this sense, the extent to which the SCCC can be implemented back into a non-monotonic formalism like inheritance networks is constrained only by the ability to construct a satisfactory formal representation of the important features contained in the SCCC, like cost and source conflict.

The aforesaid can be done employing various methodological approaches such as computational models that capture the aforesaid features and integrate them back into the logical formalism. Nevertheless, the integration between both approaches is not entirely straightforward as Bader & Hitzler point out:

Intelligent systems based on logic programming on the one hand, and on artificial neural networks (sometimes called connectionist systems) on the other, differ substantially. Logic programs are highly recursive and well understood from the perspective of knowledge representation: The underlying language is that of first-order logic, which is symbolic in nature and makes it easy to encode problem specifications directly as programs. The success of artificial neural networks lies in the fact that they can be trained using raw data, and in some problem domains the generalization from the raw data made during the learning process turns out to be highly adequate for the problem at hand, even if the training data contains some noise. Successful architectures, however, often do not use recursive (or recurrent) structures. Furthermore, the knowledge encoded by a trained neural network is only very implicitly represented, and no satisfactory methods for extracting this knowledge in symbolic form are currently known. (2004, 273-274)

Nevertheless, there has been ample efforts devoted to reconciling both approaches. For example, Hatzilygeroudis & Prentzas (2015) discuss the integration of rule-based systems (such as logical formalisms) with computational models such as neural networks. Such efforts fall into the research programme known as Knowledge-Based Neural Networks (KBANN), which "use a core of propositional rules to construct an initial neural network and then use empirical knowledge to train the network." (Hatzilygeroudis & Prentzas, 2015, 4596). Thus, if such computational methods were to be the best mechanisms to account for complex features such as cost and source conflict, KBANNs would provide the necessary means to connect this back to the non-monotonic formalism.

Having stated the above, we do not attempt to implement the SCCC into a non-monotonic formalism, as such effort falls out of the scope of our research, which is understanding and accounting for the problem of floating conclusions in its epistemological dimension. Nevertheless, this section serves the purpose of evaluating the ways and strategies that can be employed in attempting to construct an implementation of our framework of analysis into a non-monotonic formalism.

8.6 Chapter Summary

In this chapter, we addressed the task of assessing the problem of floating conclusions with the SCCC as the framework of analysis. In particular, we reviewed situations involving floating conclusions in the context of conflicting and non-conflicting sources as well as high and low-cost conditions. The above has been examined using paradigmatic cases found in the relevant literature and in the absence of examples complying with specific SCCC nuances; we have constructed cases to address such needs.

Based on our assessment, we have put forward a prescriptive stance towards both the acceptance and rejection conditions of floating conclusion and the direct/indirect modes of skeptical inference. We addressed all the possible SC and C value combinations and placed such review in the context the acceptance/rejection of floating conclusions and the direct/indirect skepticism debate. Moreover, we addressed the main objectives underlying our research. In the next chapter, we discuss the most important aspects and problems we draw from the work undertaken throughout the previous chapters.

Chapter 9 Synthesis

Throughout the previous eight chapters, we have undertaken the review, analysis and construction of an approach to better understand floating conclusions. Furthermore, we framed floating conclusions not only as a technical problem associated with inheritance networks but also as an epistemological problem in the context of defeasible reasoning. In this last chapter, we reflect on some lines of research that connect with our work. In particular, we discuss: (a) the extent to which the SCCC aligns with the heuristics research programme and its view of human cognition, (b) the congruence of the SCCC and some foundational theories of the ethics of belief, and (c) the role empirical philosophy can play in addressing some of the shortcomings raised against the methodology to understand floating conclusions. The aim of this chapter is to reflect on the problems and work we have done in the previous chapters and how they connect with other relevant lines of inquiry in epistemology and philosophy in general. Having stated the above, it is important to state that our work here leans towards the more speculative end of the spectrum as opposed to the stricter line we have decided to follow in the previous chapters.

9.1 SCCC and Cognition

In this section, we reflect on an issue that has been overlooked in the previous chapters, which is the connection between our approach to floating conclusions and its relationship with the heuristics research programme. We discuss the extent to which we undertake the same approach to rationality and human cognition as heuristics does. Furthermore, we suggest a way on which this approach has been mildly discussed but not fully developed, and the way such discussion could continue.

9.1.1 Heuristics

Heuristics is a research programme within psychology and cognitive science that is placed as a theory of how human cognitive agents draw inferences in complex environments. It is a relatively recent effort which can be traced back to the work of Herbert Simon, but it has been developed more systematically by Gerd Gigerenzer. Heuristics offers an alternative approach to human cognition and more importantly, a nuanced account of human rationality.

The heuristics approach consists on understanding inferences and decision-making in complex environments by looking into simple strategies that take into account a reduced number of features rather than resorting to complex processes that take into account numerous variables associated to the problem or the environment in which the problem emerges. Thus, to account for various reasoning processes, heuristics focuses on the most relevant features of the problem through which a simple and satisfactory strategy can be followed to address the problem at hand. The best way to understand the heuristic approach to human cognition is through an example:

A man is rushed to a hospital in the throes of a heart attack. The doctor needs to decide whether the victim should be treated as a low risk or a high risk patient. He is at high risk if his life is truly threatened, and should receive the most expensive and detailed care. Although this decision can save or cost a life, the doctor must decide using only the available cues, each of which is, at best, merely an uncertain predictor of the patient's risk level. Common sense dictates that the best way to make the decision is to look at the results of each of the many measurements that are taken when a heart attack patient is admitted, rank them according to their importance, and combine them somehow into a final conclusion, preferably using some fancy statistical software package. (Todd & Gigerenzer, 2000, 727)

Contrary to what certain approaches to rationality dictate, in this situation, there is a simple decision tree that can classify heart attack patients taking into account only three variables: age, systolic blood pressure, and the presence of sinus tachycardia. By following a very simple series of yes/no questions, doctors can very quickly and accurately classify heart attack patients, without having to process numerous variables in a complex fashion. This decision-making process is made through what is referred to as a "fast and frugal strategy". That is, the reasoning process is driven by a set of minimal variables and cues from the environment and the problem itself and completely ignores a multitude of other information. In doing so, not only it offers a straightforward strategy to address these types of scenarios, but evidence has shown that it is not just satisfactory, but it can render a superior strategy to classifying heart attack patients than more complex methods (Todd & Gigerenzer, 2000). This is one of the central tenets of the heuristics research programme, that such fast and frugal strategies can be as good and even superior to more complex forms of inferences, such as probabilistic reasoning.

As Gigerenzer points out throughout his work, one central thesis regarding human cognitive agents revolves around their intrinsic limitations, like memory, computing power, time, and knowledge about its environment. More importantly, the view of human cognition and rationality is profoundly determined by such intrinsic features. Contrary to moving away from the aforesaid 'shortcomings', heuristics paints a picture of rationality and cognition that have such shortcomings as a core feature (Todd & Gigerenzer, 2000). Moreover, these strategies that can account for important human reasoning process work as descriptive theories as well as providing the necessary building blocks for artificially intelligent systems. As Todd & Gigerenzer explain:

[...] heuristics can be seen as models of the behavior of both living organisms and artificial systems. From a descriptive standpoint, they are intended to capture how real minds make decisions under constraints of limited time and knowledge. From an engineering standpoint, these heuristics suggest ways to build artificially intelligent systems – artificial decision-makers that are not paralyzed by the need for vast amounts of knowledge or for extensive computational power. (Todd & Gigerenzer, 2000, 728)

The above underscores that in the context of this research programme, there is not a strong differentiation between the descriptive and normative dimensions of such models given that a heuristic can serve both, the purpose of describing human cognition as well as serving as a blueprint to construct a machine that performs similar inferences.

The above is a significantly compressed characterisation of the heuristic research programme. It highlights the main tenets, ideas and assumptions of human cognition and their place in the world. Hopefully, the connection between this research programme and the Source Conflict Cost Criterion is clear, although we will devote the next subsection to reflect on such explicit connection.

9.1.2 SCCC and Heuristics

As we mentioned in chapter seven, the Source Conflict Cost Criterion is the framework we propose to address the problem of floating conclusions. Nevertheless, in that chapter, we were silent about the theoretical underpinnings of our framework and our assumption of cognitive agents facing floating conclusions. Here we want to explain the idea that our framework of analysis takes on board the most important features and assumptions of the heuristics research programme.

A central idea of heuristics is the concept of *bounded rationality*, which has two important components: (a) the methods and strategies employed to solve problems and (b) the environment in which such methods and strategies are employed. An example of bounded rationality is Simon's notion of *satisficing*, which Todd & Gigerenzer explain as:

[...] a method for making a choice from a set of alternatives encountered sequentially when one does not know much about the possibilities in advance. In such situations, there may be no optimal method for stopping searching for further alternatives – for instance, there would be no optimal way of deciding when to stop looking for prospective marriage partners and settle down with a particular one (see Ch. 13 for more on satisficing in mate search). Satisficing takes the shortcut of setting an aspiration level and ending the search for alternatives as soon as one is found that exceeds the aspiration level (Simon 1956b; 1990), for instance leading an individual with Jack-Sprat- like preferences to marry the first potential mate encountered who is over a desired width. (Todd & Gigerenzer, 2000, 730)

The SCCC shares the same assumptions of bounded rationality and the idea of satisficing. On the one hand, one central assumption made by the SCCC is that cognitive agents, when facing scenarios involving floating conclusions do not need to have access to a vast number of features associated with the problem like the reliability of the sources, the probabilities associated with each line of reasoning or the quality of the evidence involved. In this sense, the SCCC offers a decision-making strategy such that it does not need to take into account numerous variables for us to assess and determine the admissibility of a floating conclusion in a specific situation. By focusing on two simple yet fundamental features of the problem (cost and source. conflict) agents can determine whether a floating conclusion is admissible or not. In this sense, the SCCC advocates for a picture of decision-making in which the agents have limited resources, such as time, computing power and knowledge about all the variables associated with the problem itself. On the other hand, the SCCC offers a decision-making strategy that can address the problem with sufficient confidence and adequacy (as we argued in chapter seven) about the admission or rejection of floating conclusions in a given scenario in the absence of an optimal strategy. It is highly debatable the existence of both: (a) all the relevant variables and (b) a known computing

mechanism; such that an optional solution to the admission of floating conclusions by means of using (b) to process (a) is feasible. Moreover, even if we grant the above, the problem itself might not even be tractable. As such, the SCCC operates on the idea that instead of working on an optimal solution to the admission of floating conclusions, it offers a decision-making strategy that satisfices the problem with a sufficient level of adequacy.

The above underscores in general terms how is it that our framework of analysis to account for floating conclusions is framed in terms of a heuristic. Nevertheless, the previous description is a general and abstract sketch of our proposal as a whole. In the next subsection, we dive deeper into how the Source Conflict Cost Criterion can take the role of a heuristic and the theoretical basis for such thesis.

9.1.3 Pragmatic Modulation

One of the rare and only attempts to assess the problem of floating conclusions outside the technical literature in non-montonic logics is from Bonnefon (2004). In his work, Bonnefon discusses the extent to which the Mental Model Theory (MMT) can account for the various reasoning strategies, i.e. credulous and skeptical approaches, and where does the MMT fall with respect to such reasoning policies. Here we want to point out how Bonnefon's work on the credulous and skeptical reasoning policies in the context of the MMT highlight the heuristic nature of the Source Conflict Cost Criterion.

There is a vast amount of work done on the MMT, and a fair share of debate around it. Briefly described, the MMT states that "when individuals understand discourse, they construct a simulation of the possibilities consistent with what the discourse describes" (Khemlani & Johnson-Laird, 2013, 2). MMT proposes that human reasoning processes, such as logical reasoning, use mental representations to construct a model of the cognitive task at hand and perform a series of operations on such models to conduct such tasks, such as making inferences.

Here we will not go in detail as to how the theory works as this is not what is at stake at the moment. Rather, it will suffice to say that Bonnefon's work on the MMT with respect to the credulous and skeptical reasoning policies associated with non-monotonic logics was able to locate the MMT as one that falls "in between credulous and skeptical" (Bonnefon, 2004, 629).

Furthermore, and what is truly meaningful for our purposes is that the MMT infers floating conclusions, which Bonnefon points out as being the object of ample debate insofar as such conclusions can be 'counter-intuitive' at times (Bonnefon, 2004).

Bonnefon identifies two possible routes to change the conclusions the MMT draw in the context of defeasible reasoning. The first possibility is to change the procedures within the theory to prevent it from inferring floating conclusions. The second possibility, which is the one Bonnefon takes as a more sensible approach, would "refine the pragmatic dimension of the theory so that it can differentiate between situations where floating conclusions or reinstated arguments are acceptable and situations where they are not." (Bonnefon, 2004, 630). Bonnefon explains that this second option would resort to understanding the specific nuances surrounding the cases in which floating conclusions seem unwarranted, as they are not always undesirable conclusions (Bonnefon, 2004). Of course, Bonnefon is unable to point exactly when, how or why such conclusions are unwarranted, but he acknowledges that they should not be dismissed altogether and therefore, we should have alternative means to make such distinction. More precisely, Bonnefon points to what is referred to in the MMT literature as the 'Principle of Pragmatic Modulation', which states that:

The principle of pragmatic modulation: The context of a conditional depends on general knowledge in long-term memory and knowledge of the specific circumstances of its utterance. This context is normally represented in explicit models. These models can modulate the core interpretation of a conditional, taking precedence over contradictory models. They can add information to models, prevent the construction of otherwise feasible models, and aid the process of constructing fully explicit models. (Johnson-Laird & Byrne, 2002, 658)

This Principle of Pragmatic Modulation (PPM) is what Boneffon sees as a potential solution for determining the adequacy and acceptance of floating conclusions. Following this route, the basic operational rules of the MMT are preserved and floating conclusions are still inferred, and the problem of determining the acceptance of floating conclusions is shifted to this new pragmatic dimension. Nevertheless, Boneffon recognises that this is no minor task when he states that:

The very general nature of this principle makes it difficult to decide how exactly it should apply to a given problem. In its present state, the principle of pragmatic modulation cannot help to address the issues I have raised here. Nevertheless, if a

solution to those issues is to be found, it will most certainly emanate from a systematic development of the pragmatic, interpretative component of MMT. It is likely that we will consider the principle of pragmatic modulation to have been the first step in this development. (Bonbon, 2004, 630)

Although our framework of analysis is not built in any way, fashion or form around the MMT, we do recognise that there is a bridge between the idea contained in the PPM as a strategy to account for the phenomena of floating conclusions in the context of defeasible reasoning. In a very liberal fashion, we can state that the pragmatic elements contained in the Source Conflict Cost Criterion seem to fall in line with the type of features one would expect in a pragmatic account for the problem of floating conclusions. As we have argued in chapter eight, both the source conflict and the cost shape the admissibility of floating conclusions. Since both notions have pragmatic elements, they can provide the grounds for better articulating the PPM in the context of defeasible reasoning. That is, in the absence of an adequate or fully developed account of the inner workings of the PPM, the SCCC can provide the building blocks as to how the PPM could work in the context of certain problems that emerge in defeasible reasoning.

Having stated the above, the full adequacy of the SCCC for the purposes of the PPM is left open to further research. At this stage, what we do wish to point is that there is a significant way in which the SCCC could provide a route to construct a more accurate notion of pragmatic modulation for defeasible reasoning. In this sense, our framework of analysis and strategy to address the problem of floating conclusions seem to fall in line with a contemporary thread in cognitive science which rethinks the role of rationality and human cognition.

9.2 Ethics of Belief

In this section, we consider the problem of floating conclusions in the context of the larger debate known as the ethics of belief. In particular, we discuss the connection of floating conclusions with some topics on responsible belief contained in the evidentialist and pragmatists approaches. We also highlight how some key concepts and stances might cohere with some ideas contained in the SCCC and how the SCCC might shed light into some of the problems and discussions in the ethics of belief. Our discussion and assessment are general as addressing all the finer details will require independent research on its own. Nevertheless, our work here is presented as a reflection and a

guide of how these lines of research can be conducted.

9.2.1 Evidentialism and Floating Conclusions.

The basic principle of the evidentialist approach to the ethics of belief is referred to as Clifford's Principle, according to which "It is wrong always, everywhere, and for anyone, to believe anything upon insufficient evidence." (Clifford, 1879, 183). Although there is ample discussion on the extent and limitation of this very basic idea, we will take it as endorsing the claim that it is only evidence what matters when debating whether we ought to espouse our beliefs. As Adler & Hicks puts it "[...] the only considerations relevant to whether one ought or ought not to believe are epistemic, and thus that one ought only to believe what one's evidence (or, more generally, epistemic reasons) establish." (2013, 140). Thus, we will take the position that as cognitive agents our epistemic duty is to ground our beliefs on the evidence one has to support or refute such belief, as doing otherwise is epistemically blameworthy.

The above poses a significant connection to the problem of floating conclusions insofar as the very nature of such conclusions is to be supported by every line of reasoning in a given situation but only through different arguments. In this context, one can legitimately raise the question as to whether such conclusions transgress the epistemic duty set out in Clifford's principle, and ultimately if endorsing such conclusions is a form of responsible belief.

To assess the above, let us consider the case proposed by Horty (2002) the commander of a military unit is considering launching an attack to an enemy which holds a strong defensive position. The commander is constantly pressed by a superior to proceed with an attacking strategy as soon as possible. Thus, he sends two scouts, Henry and Mark, to investigate the enemy's defences position. On the one hand, Henry reports that the enemy's defensive position is very weak as the main body of the enemy forces can be seen retreating to the mountains and that a diversionary group was sent out as well to make it seem like they were retreating along the river. On the other hand, Mark reports that the enemy's defensive position is very weak, but unlike Henry, Mark reports that the main body of the enemy forces can be seen retreating along the river and that a diversionary group was sent out as well to make it seem like they were retreating along the river and that a diversionary group was sent out as well to make it seem like they were retreating along the river and that a diversionary group was sent out as well to make it seem like they were retreating along the river and that a diversionary group was sent out as well to make it seem like they were retreating along the river and that a diversionary group was sent out as well to make it seem like they were retreating along the river and that a diversionary group was sent out as well to make it seem like they were retreating to the mountains.

As we have extensively discussed, the question is whether the commander should believe that the

enemy's position is indeed weak and launch his attack or withhold believing what the scouts have reported based on the fact that both sources of information undermine each other. Direct skepticism tells us that drawing the conclusion that enemy's position is week is unwarranted as both spies support it but only through different reasons, whereas the indirectly skeptical approach tells us that such conclusion is warranted insofar as it is reached by both spies, regardless if they came to the conclusion through different lines of reasoning. This much we already know. Here we want to examine the way either skeptical approach transgresses or aligns with our epistemic duty of responsible belief according to Clifford's Principle.

Framing the problem as we have done here, seems to highlight the question of what counts as sufficient evidence in such scenarios. For example, is the commander in our previous example in possession of sufficient evidence to launch his attack on the enemy at the potential expense of his own soldiers' lives? It seems that all things considered, he is indeed in possession of evidence to launch the attack on the enemy's position, based on the fact that he has two sources of information supporting such belief. Nevertheless, his evidence may not be entirely straightforward insofar as his sources of information are at odds with one another, regardless that they support the same conclusion. Thus, the question is whether the commander can espouse the belief in a responsible fashion?

To better understand floating conclusions in the context of the evidentialist dictum, we follow Feldman (2006) in his use of Chisholm's notion of *counterbalanced evidence* to understand Cliffford's Principle. In particular, Feldman states that:

[...] one's evidence concerning a proposition can either support the proposition, support the negation of that proposition, or equally support the proposition and its negation. When the evidence is equally divided in this way, I will say, following Chisholm (1989, p. 9), that the evidence is *counterbalanced*. I take Clifford's claim to be that when one's evidence supports the negation of a proposition or is counterbalanced, then one has insufficient evidence to believe that proposition. (2006, 20)

In light of the above, it seems that the problem of floating conclusions is precisely that the evidence in favour of the floating conclusion is counterbalanced. Nevertheless, one might be tempted to think that the evidence is not equally divided for and against the floating conclusion because all the evidence points towards such conclusion, in which case the problem is not having equal evidence in favour or against the floating conclusion. In our example, all the sources of information support the conclusion that the enemy is holding a weak defensive position. However, despite both spies supporting the same conclusion, the fact of the matter is that not only they do so through different lines of reasoning, but more importantly, their arguments undermine each other. That is, each account in support of the conclusion regarding the weakness of the enemy's position undermines the alternative account. Thus, in these very specific circumstances, the feature of having mutually undermining lines of reasoning seem to render floating conclusions as supported through counterbalanced evidence.

If the above assessment is correct, then it seems that an evidentialist approach to the ethics of belief would align with the directly skeptical view in the context of defeasible reasoning. That is, the intuitions underlying direct skepticism regarding floating conclusions, as discussed extensively in chapter six, revolve around the idea of such conclusions being supported only dubiously through less than secure reasons. Although the evidentialist approach is clearly not worded in this way, it does contain the same preoccupations towards what is it that we ought to believe. In this sense, admitting floating conclusions would not be a responsibly espoused belief.

The above summarises how the problem of floating conclusions can be framed in the context of the evidentialist approach to the ethics of belief, and how this specific problem of defeasible reasoning supports or undermines evidentialism. We are very much aware that there are numerous topics that we have not addressed. Nevertheless, addressing all the features of the evidentialist approach regarding floating conclusions is something that is out of scope.

9.2.2 Pragmatism and Floating Conclusions

In contrast with Clifford's Principle as a guiding benchmark to determine the acceptance or dismissal of our beliefs as appropriate or blameworthy beliefs to endorse, William James depicts a different set of standards as to how we should make such determination. In particular, James writes:

Our passional nature not only lawfully may, but must, decide an option between propositions, whenever it is a genuine option that cannot by its nature be decided on intellectual grounds; for to say, under such circumstances, "Do not decide, but leave the question open," is itself a passional decision,–just like deciding yes or no,–and is attended with the same risk of losing the truth. (1911, 11)

Like in the previous subsection, the extent and limitation of James' view will not be pursued to its fullest reach, as this is the object of an ample debate in the philosophical literature and out of scope in our research. However, we will mention a few important distinctions highlighted by Feldman (2006) regarding James's pragmatism. An *option* is the choice to adopt the belief of certain propositions. Second, the beliefs in question are said to be *living* if they have grounds to be adopted, i.e. they are plausible. An option is *forced* if not choosing between the alternatives is not a possibility, i.e. a decision must be made no matter what. Moreover, if the choice of adopting the belief is significant the option is referred to as *momentous*. Thus, when an option is *living*, *forced*, and *momentous* the option is referred to as *genuine* (Feldman, 2006).

Situations in which our choice regarding certain beliefs are not decided on *intellectual grounds* despite the evidence for such believe being counterbalanced are what lies at the centre of the pragmatist dictum concerning the beliefs we ought to endorse. To better understand the above in the context of the problem of floating conclusions we will use the same example of the commander of a military unit facing the decision of whether to launch an attack on an enemy's which holds a strong defensive position, based on the accounts of two independent scouts of his own military unit. It is significant to highlight the resemblance between the notion of a 'genuine option' and what is portrayed by the military example. In this context, various aspects are important to highlight.

First, the choice between both beliefs, i.e. that the enemy is holding a weak defensive position and the contrary is *living* in the sense that there seem to be grounds to believe as well as not believing such conclusion. For example, both scouts support the same conclusion but their accounts for supporting the conclusions undermine each other. Thus, if we opt for the idea of both reaching the same conclusion, we can endorse the belief of the enemy's weak defensive position whereas if we focus on the fact that their reasons undermine each other we might refuse to endorse such conclusion. This falls in line with what we discussed in the previous subsection regarding the evidence for the conclusion being counterbalanced. Second, given that the commander in the example *must* make a decision given that he is being pressed by his superiors to take action, we can state that the decision is *forced*, i.e. not espousing a choice with respect to the belief is not

possible. Finally, given that the decision that the commander will adopt can carry significant consequences for his life and the lives of the soldiers of his unit it is very uncontroversial to state that such decision is *momentous*.

Based on the above, we can conclude that the choice between endorsing the belief that the enemy holds a weak defensive position is *genuine* in the sense previously described, which by itself does not tell us much nor it solves the situation. Nevertheless, given that the situation qualifies as a genuine choice, then the pragmatist's approach would hold, at least in principle, that endorsing a belief in such circumstances based on non-intellectual grounds should not be blameworthy. Of course, outsourcing the justification of the belief to the vague notion of non-intellectual grounds can call into question the standing of such belief. However, the pragmatist's view, contrary to what is entailed by Clifford's principle gives room to endorse and support beliefs which fall out of the scope of strict epistemic virtues.

In line with the above, one could state that the pragmatist's approach to the ethics of belief could fall in line with the approach to indirect skepticism. Nevertheless, the connection between one and the other is fundamentally different insofar as direct skepticism is thought as a sufficiently skeptical standard by which information can be endorsed whereas the pragmatist's potential acceptance of floating conclusions can be held on non-intellectual grounds. The extent to which these two different stances are non-trivially and satisfactorily congruent is a matter that exceeds our research at this moment.

9.2.3 Ethics of Belief and the SCCC

In the previous two subsections, we have examined the extent to which the problem of floating conclusions aligns with different views concerning the question of what we ought to believe. Nevertheless, we have yet addressed how this question should be weighed with respect to our own solution to the problem of floating conclusions, i.e. the Source Conflict Cost Criterion. To address this matter, we consider the extent to which the SCCC agrees or disagrees with the two approaches to the ethics of belief discussed in the previous subsections.

As we mentioned before, there is one sense in which scenarios involving floating conclusions point to the fact that such conclusions emerge only as a byproduct of counterbalanced evidence. As we discussed in the subsection regarding the evidentialist stance, we stated that this approach seems to favour the directly skeptical stance towards floating conclusions. Nevertheless, our framework of analysis has pointed out something that has been greatly overlooked, if not entirely ignored, which is the fact that there are cases of floating conclusions which do not involve conflicting lines of reasoning. For example, let us recall the scenario discussed in section 8.3.1 concerning the patient PI seeking medical advice from three practitioners to address a life-threatening condition CI., for which we have the following prognoses:

- *D1*: To address *C1*, *P1* needs to take *X1*, because it will reduce and alleviate the distress to the liver which has been placed under industrious work.
- D2: To address C1, P1 needs to take X1, because it will increase his mobility which will facilitate engagement in physical activity which will help accelerate recovery.
- D3: To address C1, P1 needs to take X1, because it will facilitate releasing a vital neurotransmitter which is currently blocked.

Notice that, unlike the example of the commander of the military unit receiving conflicting evidence from his scouts, here the patient receives *different* evidence to support the same conclusion. That is, in cases in which we have different lines of reasoning supporting the same conclusion but only through different arguments *that do not conflict with one another* instantiates a situation in which floating conclusions are not supported through counterbalanced evidence. For example, in the case of the patient seeking for medical advice, one would be pressed to find a reason to reject the floating conclusion as we have argued in chapter eight, as one cannot expect that our conclusions will always be supported by the same reason. Different arguments can converge on the same conclusion without this being a reason to undermine the conclusion itself.

Thus, although we previously stated that in some sense floating conclusions appeal to evidence that is counterbalanced, the SCCC highlights that there are instances of floating conclusions in which there is no counterbalanced evidence, namely, cases involving non-conflicting sources, like the medical example previously discussed. As we argued in the previous chapter, in such scenarios, floating conclusions seem to have grounds to be endorsed information, even from a skeptical standard. Thus, floating conclusions in the context of non-conflicting situations seem, according to the evidentialist approach, responsible beliefs insofar as these situations do not involve counterbalanced evidence. In this sense, the evidentialist approach seems to follow align with what we have proposed through the SCC. Moreover, there can be an overlap between the notion of source conflict contained in the SCCC and the notion of counterbalanced evidence in the evidentialist stance. Nevertheless, the main point to be drawn here is that the SCCC seems to highlight an overlooked condition in the debate over floating conclusions.

Having stated the above, the extent to which the SCCC coheres with the evidentialist stance to the ethics of belief is something that falls out of scope in our research. However, the aforesaid is a legitimate line of inquiry that might help better understand both the problem of floating conclusions and its approach to the ethics of belief.

On the other hand, we think that there is an important and non-trivial connection between the pragmatist approach to the ethics of belief and the underlying assumptions of the SCCC. In particular, we think that unlike the evidentialist approach, pragmatism seems to have a strong and prima facie recognition for non-intellectual grounds on which a belief might be espoused, or what James refers to as the 'passional' realm.

Feldman describes such passional basis or non-intellectual grounds as referring to a given individual's desires and preferences (Feldman, 2006), and as we argued in the previous subsection, adopting a belief 'on a passional basis' is not wrong insofar as the choice is genuine, i.e. of great significance and inescapable. (Feldman, 2006). Nevertheless, what we want to underscore here, is the recognition of what James, and more importantly the initial pragmatists view refers to as 'passional basis', and the role it plays in our decision-making processes and belief-endorsing acts.

In our construction of the SCCC as a framework of analysis to address floating conclusions, we undertook various nuances associated with a given decision-making process: (a) the source conflict and (b) cost. We think there is room to draw an important and non-trivial argument as to why (b) encompass the type of non-intellectual features contained in the pragmatist approach to the ethics of belief. For example, the notion of cost appeals to a subjective assessment as to why a given decision-making process represents a significant burden to the agent. As we noted in chapter seven, a situation can be determined as high cost due to a series of practical matters, like legal, medical or some other form of burden. Nevertheless, it does not strike us as straightforward as how the notion of source conflict might encompass the same type of non-intellectual basis. That being said, translating the somewhat archaic notion of a decision made on a passional basis following Feldman as decisions made on desires and preferences, one can certainly see how features employed

extensively in the SCCC are aligned within such 'non-intellectual' features of our decision-making process. Again, exploring the coherence between the SCCC and the underlying assumption of the pragmatist approach to the ethics of belief seems like a potential line of inquiry that will not be pursued in its full extent here. Nevertheless, it strikes us that the SCCC seem to include important aspects and features endorsed by both the evidentialist and pragmatist stance.

We believe that the aforesaid discussion not only highlights the connection between the SCCC and other debates and problems in epistemology which can both shed light into the problem of floating conclusions, but also such debates and problems can be framed in terms of our framework of analysis to better understand them.

9.3 Empirical Basis

In this section, we consider the methodological problem raised in the literature that discusses the problem of floating conclusions. In particular, we focus on the methodological shortcomings which both sides of the debate agree upon, and how they can be addressed with the aim of the methods, and strategies proposed by the field of empirical philosophy. More precisely, we propose the idea that neither the debate nor the way it has been conducted so far has to be abandoned, even if one concedes the legitimacy of the methodological critique raised against the debate. In contrast, we suggest, even if the criticism stands, an empirical approach can address a significant part of such problems. The aim of this section is to reflect on how empirical philosophy can refine and enhance the debate over floating conclusions. That being said, our assessment is general and abstract and we do not construct experimental design, nor go in full detail as to what specific methods would be best, but rather how empirical philosophy as a whole might addressed some of the substantial criticisms of the methodological approach to floating conclusions.

9.3.1 Methodological Inadequacy

As we mentioned in chapter six, the problem of floating conclusions has been addressed unsatisfactorily. However, we are not alone in this assessment of the discussion, as one key aspect that both sides of the debate seem to agree on, is that the methodology employed to address the problem has been addressed quite inadequately. In particular, the literature converges on the idea that relying on examples to support either version of skeptical inference is a shortcoming of the debate. In this sense, two specific points have been raised. First, these examples often carry hidden assumptions that are not made explicit. Second, there is no agreement of different author's intuitions on these examples.

Concerning the first point, one of the authors to note such problem was Schlechta (1993) who pointed out that the examples that were constructed and discussed carried a significant amount of 'unexpressed information', which render the strategy of relying on examples as less than desirable. That is, depending on how examples are crafted, they might contain more or less information that in turn could lead one's intuitions astray. For example, in assessing the case of the large inheritance which can be used to purchase a luxury item, Prakken states that:

[...] it is very easy to read an additional default principle into the example, viz. that people tend to speak the truth about their intentions, and that it is also very easy to mistake this additional information for the intuition that the reasoning pattern is invalid. Note that this additional default principle is undercut as soon as it turns out that a person has told conflicting things about his or intentions to different persons. Now if both the additional principle and its undercutter are made explicit, there will also be conclusion sets where neither father nor mother lets Bob inherit anything, so that the issue of floating conclusions does not arise. (2002, 6)

Regarding the second point, Prakken highlights that the intuitions of the various authors sometimes conflict. Moreover, that the intuitions that are being considered are of persons with a heavy acquaintance with the problem and the methods of logic. Thus, Prakken resolves to the idea that "in modelling defeasible reasoning it is very dangerous to rely on intuitions in concrete examples." (Prakken, 2002, 6). Following such thread, and the strongest stance towards the role of intuitions in the context of floating conclusions and how it has been addressed comes from a relatively recent paper from Maier & Nute in which they state that:

Acceptance of floating conclusions is not universal, however. Defeasible logic rejects them, as do other proponents of direct skepticism such as Horty (2002) and Antonelli (2005). [....] Prakken (2002) opposes this and writes that reliance on intuition pumps of the above sort can be a very dangerous affair, and that intuitions in many other examples support floating conclusions. He also writes that in at least some of Horty's examples there is implicit information which, when made explicit,

would prevent one from drawing the floating conclusion. [....] We agree that one may be misled by poor examples or faulty intuitions, but notice that opponents on both sides of this issue depend on examples and intuitions to support their positions. We have not found in the literature disciplined arguments either for or against floating conclusions that do not rely on these intuition pumps. (2010, 247-248, italics added)

Despite the above, the debate has relied exclusively on examples and case studies to argue in favour of one position or the other, all while admitting at the same time that relying on the intuitions elicited by specific examples is a suboptimal approach to address the subject matter. Moreover, what it is more often criticised is how people's intuitions regarding the various examples are to be assessed, if they are to be taken into consideration at all.

Thus, the debate surrounding floating conclusions is quite paradoxical insofar as it has proceeded following the strategy everyone agrees as a substandard approach to the problem. Moreover, there exists a common and unfortunate consensus that not only there has been no progress whatsoever in our understanding of the problem, but that the method we have employed to address the problem seems to perpetuate such state of affairs.

With the aforesaid background, it is only natural to reflect on the accuracy of such consensus. In particular, we wish to question the assumption that the way the debate has been conducted is intrinsically mistaken. For this matter, in the next subsection, we expand our discussion regarding the role of intuitions and what they can tell us about the problem of floating conclusions.

9.3.2 The Role of Intuitions

One of the most common problems of the role of intuitions in the context of floating conclusions has to do with the very controversial fact that different authors disagree on each other's assessment of the same examples. Regarding this matter, Prakken explains that "Horty's examples are intriguing and his discussion of them is very insightful. Yet I disagree with the conclusions he draws from them. I shall argue that his examples do not demonstrate the invalidity of the reasoning patterns." (2002, 1). Thus, the problem seems not to be the use of examples or hypothetical scenarios to probe our attitudes to the acceptance or rejection or floating conclusions, but the lack

of consensus from such examples.

As we mentioned in the previous subsection, one crucial point is that researchers are highly biased in such a way that their assessment of the examples might not reflect a layperson's intuitions on the same scenarios. As Prakken notes "Then the problem often arises that their answers reveal a lacking understanding of the reasoning patterns; however, teaching them about these reasoning patterns infects them with the theoretical bias we were hoping to avoid." (Prakken, 2002, 2). Nevertheless, this can be solved by outsourcing the evaluation of the various cases to people without significant training or knowledge on the subject matter. This has somehow been hinted in the literature. For example, and regarding the case of inheriting a large sum of money discussed in section 6.2.3, Horty states that:

In this situation, then, there is a vivid practical difference between the two skeptical options. If I were to reason according to the first, I would not be justified in concluding that I am about to inherit half a million dollars, and so it would be foolish for me to place a deposit on the yacht. If I were to reason according to the second, I would be justified in drawing this conclusion, and so it would be foolish for me not to place a deposit. Which option is correct? *I have not done a formal survey, but most of the people to whom I have presented this example are suspicious of the floating conclusion, and so favor the argument intersection option. Most do not feel that the initial information from the default theory would provide sufficient justification for me to conclude, as the basis for an important decision, that I will inherit half a million dollars. Certainly, this is my own opinion—I believe the example shows, contrary to the widely held assumption, that it is at least coherent for a skeptical reasoner to withhold judgment from floating conclusions. (2012, 182, italics added)*

An even more problematic and perplexing situation is that more often than not, these authors disagree with other people's intuitions regarding such hypothetical cases. For example, when reflecting on the Ginsberg Diamond scenario John Horty's states that:

There is no need to labor the point by fabricating further examples in which floating conclusions are suspect. But what about the similar cases, exemplifying the same pattern, that have actually been advanced as supporting floating conclusions, such as Ginsberg's political extremist [...] *I have always been* surprised that this particular example has seemed so persuasive to so many people. (2002, 69, italics added)

Thus, what we can draw from this and the previous subsection is that the literature agrees that a serious shortcoming associated with the debate over floating conclusions consists of relying on hypothetical scenarios. Nevertheless, we believe that this is an instance when researchers seem to be 'throwing the baby out with the bathwater'. That is, doing without examples and hypothetical scenarios entirely might be the wrong way to frame the debate. To make our point, we suggest in the next subsection a strategy that can remedy the shortcomings, namely recurring to experimental philosophy to find an empirical basis to layperson's intuitions on the discussed examples.

9.3.3 Empirical Approach to Floating Conclusions

Putting aside the assumed inadequacy of relying on hypothetical examples and our intuitions they elicit, an important line of research has been significantly overlooked, i.e. the possibility of finding an empirical basis for our intuitions when confronted with such examples through the research programme known as empirical philosophy.

The extent to which a given method or line of inquiry aligns with the experimental philosophy research programme is the object of discussion, as there are important milestones early in the second half of 20th century philosophy (Mortensen & Nagel, 2016). Nevertheless, as Williamson defines it, experimental philosophy:

[...] refers to a more specific kind of philosophically motivated experimental inquiry, in which verdicts on hypothetical cases relevant to some philosophical question are elicited from significant numbers of subjects, sometimes under controlled conditions, and hypotheses are tested about the underlying patterns. (2016, 22)

One of the aspects that experimental philosophy focuses is on examining with great detail both philosophers and untrained layperson's assessments of hypothetical cases. Experimental philosophy is usually critical of philosophers' intuitions and more inquisitive on untrained layperson's judgments of common philosophical problems. Thus, employing methods common to

social sciences such as statistics, surveys, etc. experimental philosophy tries to test the adequacy of philosophers' views with respect to a given problem.

Nevertheless, here we do not wish to describe in full detail the research programme of experimental philosophy and the challenges it raises to common philosophical discourse, but rather we wish to point out how this research programme offers a legitimate solution to the main concerns and preoccupations expressed by the various authors regarding the problem of floating conclusions and its methodological approach.

For example, it may very well be the case that an empirical study conducted on an appropriately selected demographic presented with examples containing floating conclusions might show a consensus on a given case or important classes of cases. Furthermore, through the methods employed in empirical philosophy, we might come to better understand what features or factors are at play when considering floating conclusions. We might come to realise that intuitions of untrained persons disagree with what the literature has taken for granted, or we might find that they somehow align with one or other side of the debate.

The aforesaid is highly speculative, and it is important to point out that we are not proposing an experimental study nor the general sketches of its design. The point we think is important to be made is that some concerns raised in the literature might have grounds. This should not imply that the strategy that has been employed to address the problem of floating conclusions nor any aspiration to better understand the problem should be abandoned, as there are legitimate venues through which the theoretical and metatheoretical dispute can be addressed.

In line with the above, it also needs to be stated that we do not advocate to outsourcing the problem of floating conclusions to experimental philosophy, as the potential findings there need not completely dictate the theories constructed or defended for non-monotonic reasoning. Following this line of thought, Mortensen & Nagel point out that:

Armchair philosophers are not obliged to construct theories that will capture or accommodate all our intuitions: they can take intuitions elicited by the case method as a defeasible source of evidence concerning the ultimate targets of their inquiry. However, cases matter enough to the armchair that if there is empirical work that can also help to defend the case method from challenges, remedy its faults, or extend its reach, this work would clearly qualify as armchair-friendly. (2016, 57)

The above discussion although speculative wishes to address what we think is a widespread distrust in the discussion of floating conclusions and the way it has been conducted. Although it may not solve the problem entirely, it is a line of research that deserves significant attention as results in the area with the methods previously described can shed light into the nature of the problem.

9.4 Chapter Summary

In this chapter, we addressed the task of reflecting how our work relates to various fields in cognitive science, epistemology and philosophy in general. In particular, we drew connections between our work and the heuristics research programme of cognitive science, the ethics of belief and experimental philosophy. We discussed both how various discussions in those fields can contribute to a better understanding of our framework of analysis, and how our framework of analysis can shed light in some of those discussions. Although there is a significant amount of matters, we did not include, we have pointed out how those discussions can be further pursued. In the next chapter, we put forward the most important conclusions we draw from the work undertaken throughout all of the previous chapters.

Chapter 10 Conclusions

Throughout the previous eight chapters, we have undertaken the review, analysis and construction of an approach to better understand floating conclusions. Furthermore, we framed floating conclusions not only as a technical problem associated with inheritance networks, but also as an epistemological problem in the context of defeasible reasoning. In line with our work, we have addressed four main objectives which we immediately proceed to consider.

First, and foremost, our research has addressed the need for a thorough and systematic assessment of the problem of floating conclusions not only as a problem associated with inheritance networks but also as a problem that highlights the need for a better epistemological understanding. In light of the above, an overt relevance has been given to the epistemological debate based on the view that the technical discussion by itself has been unable to provide a proper understanding of the problem so far. In such a vein, our work has tried both, to overcome the intrinsic shortcomings surrounding the debate over floating conclusions (such as its dispersed nature), and to place the focus on the epistemological debate rather than on the technical subtleties of such problem.

Second, and based on our assessment of floating conclusions and the surrounding epistemological debate, but breaking from what has been attempted so far, we have constructed a unified and general framework of analysis to tackle such conclusions, namely the *source conflict cost criterion*. The framework of analysis of our making, although general, is based on two straightforward and simple cues: (a) conflict among sources or lack thereof and, (b) the cost associated with the decision-making process. In the same vein, and despite having constructed our framework of analysis based on a pair of very general notions (i.e. source, conflict and cost), we have simplified the framework in terms of a binary model that retains its utility without scarifying any of the theoretical intuitions that grounds it.

Third, and through the application of our two-component framework, we undertook the analysis of the technical problem of floating conclusions as well as the epistemological debate surrounding such problem. In particular, we examined how the various conditions regarding the conflict of sources (or lack thereof) as well as the various conditions regarding cost shape our stance in the

debate. Moreover, our framework of analysis identifies and accommodates the salient traits involved in the various examples surrounding the debate of floating conclusions.

Fourth, and based on the *source conflict and cost criterion*, we have provided a normative stance for the acceptance and rejection conditions of floating conclusions. In particular, we have taken the stance that the only situation in which floating conclusions render themselves as skeptically unacceptable information is where there exists conflict among the sources coupled with a high cost associated to the decision-making process. In the same vein, we have also stipulated that all other situations render floating conclusions as sufficiently skeptical information.

As we have stated in the third chapter concerning the extent and limitations of inheritance networks and their role in the context of our research, it should be clear that the present work does not solve nor address a myriad of still open questions concerning such a knowledge representation framework for defeasible reasoning. Furthermore, the question of how our approach to floating conclusions fits into the philosophical and epistemological problems associated with such type of networks is unaddressed. However, this simply exceeds the limits of our work. In this sense, our research can be said to bypass several of the foundational problems underlying the research programme of inheritance networks. For example, in the second chapter, we noted the various discussions and debates surrounding the extent and limitation of the concept of nodes and their multiple interpretations and we followed suit with the notion of links. The above can be said to be a substantive issue inasmuch as it results in different approaches to inheritance networks as well as the various implementation-specific frameworks. In this sense, our research does not provide answers to such types of questions; nevertheless, this is a deliberate choice on our behalf.

In such a vein, it needs to be stated that we take on board such criticism. Nevertheless, it also needs to be understood that despite the legitimacy of such questions, they are only indirectly associated with our research since they are debates attached to the knowledge representation framework, but do not belong to the core problem undertaken by our research. Our work addresses a specific problem associated with a basic framework of inheritance networks and some of the basic conflict-resolution mechanisms. To this extent, it satisfies the objective it has set itself to address, but recognises the need for further research into other problems. This, in itself, does not represent a limitation, but it does produce certain features of our work that can be held as such. In particular, we see the above as having three particular reasons that render it *prima facie* of more noteworthy merit.

First, we have addressed the problem of floating conclusions through a unified and general framework of analysis. That is, we have provided the theoretical grounds for a comprehensive approach to the phenomenon of floating conclusions. Moreover, such a comprehensive framework has the conditions to account for floating conclusions and has its place in the debate between direct and indirect skepticism without aligning itself with either approach. This has been the precise shortcoming pervading all of the previous attempts to address the problematic nature of floating conclusions.

Second, we have addressed the previous through an inherently implementation-agnostic approach to the problem of floating conclusions. That is, we have constructed our framework of analysis in a way that supersedes the set of idiosyncratic or implementation-specific networks. In such vein, we have not committed ourselves to a fixed implementation of inheritance networks or any of the involved components, but rather we have looked into the essential features associated with the problem and proposed our contribution based on these traits rather than on an implementationdependent version of the problem.

Finally, the previous pair of advantages led to the fact that by being (a) a general and unified framework of analysis and (b) an implementation-agnostic approach, we end up having an understanding of the problem of floating conclusions, its extent and limitation in a way that can account for both the technical and epistemological aspects of the problem. That is, our approach to the problem satisfies both ends of the spectrum and sacrifices neither for the other. This has allowed us to ground our approach in a simple framework that has not restricted our understanding of the problem nor our stance over the debate of such problem.

The previous traits inherent in our approach strike us exactly as a trade-off which one way or another one is compelled to take. Despite the underlying predicament, we not only are aware of its existence, but we have taken a deliberate stance towards it, and find that it also reveals an underlying set of methodological and theoretical features that enhance our research.

Having stated the above, we conclude by highlighting that the problem of floating conclusions is an intricate technicality that has somewhat gone under the radar (with a few exceptions) in the field of defeasible reasoning inasmuch as its treatment has been peripheral. In our view, much of this failure to have a comprehensive understanding emerges from the interdisciplinary nature of the field in which the problem lies. Thus, much philosophical or conceptual effort has come from logicians and computer scientists who have technical preoccupations as their primary interest, and epistemologist have not addressed the problem to any extent because it is buried within many technical subtleties. Our effort has been to synthesise both realms to achieve a better understanding of the problem. In this sense, this is not a thesis on epistemology, nor a thesis on logic, but rather our work takes a technical problem located within a formal framework, examines the epistemological dimension of the problem, and offers a comprehensive solution.

Furthermore, the problem by itself is one that exceeds the realm of logic or epistemology. As it has been emphasised in the context of our research, floating conclusions have important and farreaching implications for real-world situations concerning information processing and decisionmaking. In such a vein, our research is a first, but comprehensive step in the direction of having a proper understanding of such a phenomenon. Far from aiming to be the last word on the subject matter, our earnest hope is that it only serves the purposes of producing a debate, but one in the right direction, with the proper tools and means to tackle the task at hand, one which we are confident has been satisfactorily hereby started.

Bibliography

- Adler, J., & Hicks, M. (2013). Non-Evidential Reasons to Believe. *The Aim of Belief*, 116, 141–166.
- Anderson, M. L., Gomaa, W., Grant, J., & Perlis, D. (2013). An Approach to Human-Level Commonsense Reasoning. In K. Tanaka, F. Berto, E. Mares, & F. Paoli (Eds.), *Paraconsistency: Logic and Applications* (pp. 201–222). Springer Netherlands. <u>https://doi.org/10.1007/978-94-007-4438-7_12</u>

Antonelli, A. (2005). Grounded Consequence for Defeasible Logic. Cambridge University Press.

- Antoniou, G. (2006). Defeasible Reasoning: A Discussion of Some Intuitions. International Journal of Intelligent Systems, 21(6), 545–558. <u>https://doi.org/10.1002/int.20147</u>
- Bader, S., & Hitzler, P. (2004). Logic programs, iterated function systems, and recurrent radial basis function networks. *Journal of Applied Logic*, 2(3), 273–300. <u>https://doi.org/10.1016/j.jal.2004.03.003</u>
- Bastiaanse, H., & Veltman, F. (2016). Making the right exceptions. *Artificial Intelligence*, 238, 96–118. <u>https://doi.org/10.1016/j.artint.2016.05.005</u>
- Batens, D. (2009). The Need for Adaptive Logics In Epistemology. In S. Rahman, J. Symons, D.
 M. Gabbay, & J. P. van Bendegem (Eds.), *Logic, Epistemology, and the Unity of Science* (pp. 459–485). Springer Netherlands. <u>https://doi.org/10.1007/978-1-4020-2808-3_22</u>
- Bonnefon, J.-F. (2004). Reinstatement, floating conclusions, and the credulity of Mental Model reasoning. *Cognitive Science*, 28(4), 621–631. <u>https://doi.org/10.1016/j.cogsci.2004.03.002</u>
- Brachman. (1983). What IS-A Is and Isn't: An Analysis of Taxonomic Links in Semantic Networks. *Computer*, *16*(10), 30–36. <u>https://doi.org/10.1109/MC.1983.1654194</u>

- Brachman, R. J. (1977). What's in a concept: Structural foundations for semantic networks. International Journal of Man-Machine Studies, 9(2), 127–152. <u>https://doi.org/10.1016/S0020-7373(77)80017-5</u>
- Brachman, R. J. (1979). On the Epistemological Status of Semantic Networks. In N. V. Findler (Ed.), Associative Networks (pp. 3–50). Academic Press. <u>https://doi.org/10.1016/B978-0-12-256380-5.50007-4</u>
- Brachman, R., & Levesque, H. (2004). *Knowledge Representation and Reasoning* (First Edition). Morgan Kaufmann.
- Brewka, G. (1991). Nonmonotonic Reasoning: Logical Foundations of Commonsense (Vol. 12). Cambridge University Press.
- Brewka, G., Niemelä, I., & Truszczyński, M. (2008). Nonmonotonic Reasoning. In F. van Harmelen, V. Lifschitz, & B. Porter (Eds.), *Foundations of Artificial Intelligence* (Vol. 3, pp. 239–284). Elsevier. <u>https://doi.org/10.1016/S1574-6526(07)03006-4</u>
- Cantwell, J. (2009). Conditionals in reasoning. *Synthese*, *171*(1), 47–75. https://doi.org/10.1007/s11229-008-9379-6
- Carnap, R., & Bar-Hillel, Y. (1952). An Outline of a Theory of Semantic Information (TR-247; Technical Report (Massachusetts Institute of Technology. Research Laboratory of Electronics)). Research Laboratory of Electronics, Massachusetts Institute of Technology. https://dspace.mit.edu/handle/1721.1/4821
- Carpenter, B., & Thomason, R. (1990). Inheritance Theory and Path-Based Reasoning: An Introduction. In H. E. Kyburg, R. P. Loui, & G. N. Carlson (Eds.), *Knowledge Representation and Defeasible Reasoning* (pp. 309–343). Springer Netherlands. <u>https://doi.org/10.1007/978-94-009-0553-5_13</u>
- Clifford, W.K. (1879). The Ethics of Belief. Originally printed in *Contemporary Review* (1877) and reprinted in Clifford's *Lectures and Essays* London: MacMillan.

- Costa, H. A. (1999). Epistemic Context, Defeasible Inference and Conversational Implicature. International and Interdisciplinary Conference on Modeling and Using Context, 15–27.
- Feldman, R. (2000). The ethics of belief. *Philosophy and Phenomenological Research*, 60(3), 667–695.

Feldman, R. (2006). Clifford's principle and James's options. Social Epistemology, 20(1), 19-33.

Floridi, L. (2004). Information. In *The Blackwell Guide to the Philosophy of Computing and Information* (pp. 40–61). John Wiley & Sons, Ltd. <u>https://doi.org/10.1002/9780470757017.ch4</u>

Floridi, L. (2011). The Philosophy of Information. Oxford University Press.

- Floridi, L. (2016). Semantic information. In L. Floridi (Ed.), *The Routledge Handbook of Philosophy of Information* (pp. 60–65). Routledge.
- Floridi, L. (2019). Semantic Conceptions of Information. In *The Stanford Encyclopedia of Philosophy* ((Winter 2019 Edition)). Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/win2019/entries/information-semantic/
- Gabbay, D. M. (1985). Theoretical Foundations for Non-Monotonic Reasoning in Expert Systems. In *Logics and Models of Concurrent Systems* (pp. 439–457). Springer.
- Gabbay, D. M., & Schlechta, K. (2010). An Analysis of Defeasible Inheritance Systems. In D.
 M. Gabbay & K. Schlechta (Eds.), *Logical Tools for Handling Change in Agent-Based* Systems (pp. 251–293). Springer. <u>https://doi.org/10.1007/978-3-642-04407-6_9</u>
- Gabbay, D. M., & Schlechta, K. (2016). Defeasible Inheritance. In D. M. Gabbay & K. Schlechta (Eds.), A New Perspective on Nonmonotonic Logics (pp. 75–90). Springer International Publishing. <u>https://doi.org/10.1007/978-3-319-46817-4_3</u>
- Gabbay, D. M., & Woods, J. (2008). Resource-origins of Nonmonotonicity. *Studia Logica*, 88(1), 85–112. <u>https://doi.org/10.1007/s11225-008-9100-2</u>

Gigerenzer, G. (2005). I think, therefore I err. *Social Research: An International Quarterly*, 72(1), 195–218.

Ginsberg, M. L. (1987). Readings in Nonmonotonic Reasoning. Morgan Kaufmann Publishers.

- Grégoire, É., Lagniez, J.-M., & Mazure, B. (2014). A General Artificial Intelligence Approach for Skeptical Reasoning. In B. Goertzel, L. Orseau, & J. Snaider (Eds.), *International Conference on Artificial General Intelligence* (pp. 33–42). Springer International Publishing. <u>https://doi.org/10.1007/978-3-319-09274-4_4</u>
- Hatzilygeroudis, I., & Prentzas, J. (2015). Symbolic-neural rule based reasoning and explanation.
 Expert Systems with Applications, 42(9), 4595–4609.
 <u>https://doi.org/10.1016/j.eswa.2015.01.068</u>
- Hendrix, G. G. (1979). Encoding Knowledge in Partitioned Networks. In NICHOLAS V. Findler (Ed.), Associative Networks (pp. 51–92). Academic Press. <u>https://doi.org/10.1016/B978-0-12-256380-5.50008-6</u>
- Horty, J. F. (1990). A Skeptical Theory of Mixed Inheritance. In J. M. Dunn & A. Gupta (Eds.), *Truth or Consequences: Essays in Honor of Nuel Belnap* (pp. 267–281). Springer Netherlands. <u>https://doi.org/10.1007/978-94-009-0681-5_17</u>
- Horty, J. F. (1994). Some Direct Theories of Nonmonotonic Inheritance. In Handbook of Logic in Artificial Intelligence and Logic Programming (vol. 3): Nonmonotonic Reasoning and Uncertain Reasoning (Vol. 3, pp. 111–187). Oxford University Press, Inc.
- Horty, J. F. (2001). Argument construction and reinstatement in logics for defeasible reasoning. *Artificial Intelligence and Law*, 9(1), 1–28. <u>https://doi.org/10.1023/A:1011288805074</u>
- Horty, J. F. (2002). Skepticism and floating conclusions. *Artificial Intelligence*, *135*(1), 55–72. https://doi.org/10.1016/S0004-3702(01)00160-6

Horty, J. F. (2012). Reasons as Defaults. Oxford University Press.

- Horty, J. F., & Thomason, R. H. (1990). Boolean extensions of inheritance networks.*Proceedings of the Eighth National Conference on Artificial Intelligence*, 1, 633–639.
- Horty, J. F., & Thomason, R. H. (1988). Mixing strict and defeasible inheritance. *Proceedings of the Seventh AAAI National Conference on Artificial Intelligence*, 427–432.
- Horty, J. F., Thomason, R. H., & Touretzky, D. S. (1990). A skeptical theory of inheritance in nonmonotonic semantic networks. *Artificial Intelligence*, 42(2), 311–348. <u>https://doi.org/10.1016/0004-3702(90)90057-7</u>
- James, W. (1911). 'The Will to Believe.' Reprinted in *The Will to Believe and Other Essays in Popular Philosophy*. New York: Dover Publications, 1956.
- Khemlani, S., & Johnson-Laird, P. N. (2013). The processes of inference. *Argument & Computation*, 4(1), 4–20.
- Knobe, J. (2016). Experimental Philosophy Is Cognitive Science. In A Companion to Experimental Philosophy (pp. 37–52). John Wiley & Sons, Ltd. <u>https://doi.org/10.1002/9781118661666.ch3</u>
- Koons, R. (2017). Defeasible reasoning. In *The Stanford Encyclopedia of Philosophy* (Winter 2017 Edition). Metaphysics Research Lab, Stanford University. <u>https://plato.stanford.edu/archives/win2017/entries/reasoning-defeasible/</u>
- Kraus, S., Lehmann, D., & Magidor, M. (1990). Nonmonotonic Reasoning, Preferential Models and Cumulative Logics. *Artificial Intelligence*, 44(1–2), 167–207.
- Kuipers, B. (1979). On Representing Commonsense Knowledge. In NICHOLAS V. Findler (Ed.), Associative Networks (pp. 393–408). Academic Press. <u>https://doi.org/10.1016/B978-0-12-256380-5.50018-9</u>
- Lehmann, D., & Magidor, M. (1992). What does a conditional knowledge base entail? *Artificial Intelligence*, 55(1), 1–60.

- Lehmann, D., & Magidor, M. (1990). Preferential logics: The Predicate Calculus Case. Proceedings of the 3rd Conference on Theoretical Aspects of Reasoning about Knowledge, 57–72.
- Levesque, H., & Mylopoulos, J. (1979). A Procedural Semantics for Semantic Networks. In NICHOLAS V. Findler (Ed.), Associative Networks (pp. 93–120). Academic Press. <u>https://doi.org/10.1016/B978-0-12-256380-5.50009-8</u>
- Maier, F., & Nute, D. (2010). Well-founded semantics for defeasible logic. *Synthese*, *176*(2), 243–274. https://doi.org/10.1007/s11229-009-9492-1
- Makinson, D. (2005). How to Go Nonmonotonic. In D.M. Gabbay & F. Guenthner (Eds.), Handbook of Philosophical Logic (pp. 175–278). Springer Netherlands. <u>https://doi.org/10.1007/1-4020-3092-4_3</u>
- Makinson, D., & Schlechta, K. (1991). Floating conclusions and zombie paths: Two deep difficulties in the "directly skeptical" approach to defeasible inheritance nets. *Artificial Intelligence*, 48(2), 199–209. <u>https://doi.org/10.1016/0004-3702(91)90061-N</u>
- McCarthy, J. (1980). Circumscription—A Form of Non-Monotonic Reasoning. *Artificial Intelligence*, *13*(1–2), 27–39.
- McDermott, D., & Doyle, J. (1980). Non-monotonic logic I. *Artificial Intelligence*, *13*(1–2), 41–72.
- Meheus, J., Straßer, C., & Verdée, P. (2013). Which style of reasoning to choose in the face of conflicting information? *Journal of Logic and Computation*, *26*(1), 361–380.
- Mortensen, K., & Nagel, J. (2016). Armchair-Friendly Experimental Philosophy. In A Companion to Experimental Philosophy (pp. 53–70). John Wiley & Sons, Ltd. <u>https://doi.org/10.1002/9781118661666.ch4</u>
- Nute, D. (1988). Defeasible Reasoning: A Philosophical Analysis in Prolog. In Aspects of Artificial Intelligence (Fetzer, James H., pp. 251–288). Springer.
- Nute, D. (2003). Agents, Epistemic Justification, and Defeasibility. *Invited Address, 5th Augustus de Morgan Workshop*.
- Pollock, J. L. (1987). Defeasible Reasoning. *Cognitive Science*, *11*(4), 481–518. <u>https://doi.org/10.1207/s15516709cog1104_4</u>
- Pollock, J. L. (1995). Cognitive Carpentry: A Blueprint for how to Build a Person. MIT Press.
- Prakken, H. (2002). Intuitions and the Modelling of Defeasible Reasoning: Some case studies. Proceedings of the 9th International Workshop on Non-Monotonic Reasoning (NMR'2002). <u>https://arxiv.org/abs/cs/0207031v1</u>
- Quillian, M. R. (1967). Word concepts: A Theory and Simulation of some Basic Semantic Capabilities. *Behavioral Science*, 12(5), 410–430. <u>https://doi.org/10.1002/bs.3830120511</u>
- Rankin, T. L. (1988). When is Reasoning Nonmonotonic? In J. H. Fetzer (Ed.), Aspects of Artificial Intelligence (pp. 289–308). Springer Netherlands. <u>https://doi.org/10.1007/978-94-</u>009-2699-8 10
- Reiter, R. (1980). A Logic for Default Reasoning. *Artificial Intelligence*, *13*(1), 81–132. https://doi.org/10.1016/0004-3702(80)90014-4
- Reiter, R. (1987). Nonmonotonic reasoning. Annual Review of Computer Science, 2(1), 147-186.
- Schlechta, K. (1993). Directly Sceptical Inheritance Cannot Capture the Intersection of Extensions. *Journal of Logic and Computation*, *3*(5), 455–467.
- Schubert, L. K., Goebel, R. G., & Cercone, N. J. (1979). The Structure and Organization of a Semantic Net for Comprehension and Inference. In NICHOLAS V. Findler (Ed.), *Associative Networks* (pp. 121–175). Academic Press. <u>https://doi.org/10.1016/B978-0-12-256380-5.50010-4</u>
- Shapiro, S. C. (1981). What do Semantic Network Nodes Represent. *Conference on* Foundational Threads in Natural Language Processing SUNY at Stony Brook, 1–15.

- Simonet, G., & Ducournau, R. (1994). On Stein's paper: Resolving ambiguity in nonmonotonic inheritance hierarchies. *Artificial Intelligence*, 71(1), 183–193. <u>https://doi.org/10.1016/0004-3702(94)90065-5</u>
- Sperber, D., & Wilson, D. (1995). *Relevance: Communication and Cognition* (Second Edition). Blackwell Publishers.
- Stalnaker, R. (1994). What is a Nonmonotonic Consequence Relation? *Fundamenta Informaticae*, *21*(1, 2), 7–21.
- Stein, L. A. (1989). Skeptical Inheritance: Computing the Intersection of Credulous Extensions (Technical Report CS-89-14). Department of Computer Science, Brown University.
- Stein, L. A. (1991). *Resolving ambiguity in nonmonotonic reasoning* [PhD Thesis]. Brown University.
- Stein, L. A. (1992). Resolving ambiguity in nonmonotonic inheritance hierarchies. Artificial Intelligence, 55(2–3), 259–310. <u>https://doi.org/10.1016/0004-3702(92)90057-5</u>
- Strasser, C., & Antonelli, G. A. (2018). Non-monotonic Logic. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2018). Metaphysics Research Lab, Stanford University. <u>https://plato.stanford.edu/archives/sum2018/entries/logic-nonmonotonic/</u>
- Thomason, R. H., & Horty, J. F. (1989). Logics for inheritance theory. In M. Reinfrank, J. de Kleer, M. L. Ginsberg, & E. Sandewall (Eds.), *Non-Monotonic Reasoning* (pp. 220–237). Springer. <u>https://doi.org/10.1007/3-540-50701-9_30</u>
- Thomason, R. H., Horty, J. F., & Touretzky, D. S. (1987). A Calculus for Inheritance in Monotonic Semantic Nets (CMU-CS-86-138). Carnegie-Mellon. University of Pittsburg. Computer Science Department.
- Todd, P. M., & Gigerenzer, G. (2000). Précis of "Simple heuristics that make us smart". *Behavioral and Brain Sciences*, 23(5), 727–741.

Touretzky, D. S. (1986). The Mathematics of Inheritance Systems. Morgan Kaufmann.

- Touretzky, D. S., Horty, J. F., & Thomason, R. H. (1987). A Clash of Intuitions: The Current State of Nonmonotonic Multiple Inheritance Systems. *Proceedings of the 10th International Joint Conference on Artificial Intelligence - Volume 1*, 476–482.
- Touretzky, D. S., Thomason, R. H., & Horty, J. F. (1991). A Skeptic's Menagerie: Conflictors, Preemptors, Reinstaters, and Zombies in Nonrnonotonic Inheritance. *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence (IJCAI-91)*, 478–485.
- Truszczyński, M. (2006). Nonmonotonic Logics and Their Algebraic Foundations. In Z. Ésik (Ed.), Computer Science Logic (pp. 58–71). Springer. <u>https://doi.org/10.1007/11874683_4</u>
- Williamson, T. (2016). Philosophical Criticisms of Experimental Philosophy. In A Companion to Experimental Philosophy (pp. 22–36). John Wiley & Sons, Ltd. <u>https://doi.org/10.1002/9781118661666.ch2</u>
- Wilson, D., & Sperber, D. (2004). Relevance Theory. In *The Handbook of Pragmatics* (pp. 607–632). Blackwell Publisher.
- Witteveen, G., & van der Hoek, W. (1997). A General Framework for Revising Nonmonotonic Theories. *International Conference on Logic Programming and Nonmonotonic Reasoning*, 258–272.
- Woods, W. A. (1975). What's in a Link: Foundations for Semantic Networks. In D. G. Bobrow & A. Collins (Eds.), *Representation and Understanding* (pp. 35–82). Morgan Kaufmann. https://doi.org/10.1016/B978-0-12-108550-6.50007-0