

Case-based reasoning in the law

A formal theory of reasoning by case comparison

Abraham Cornelis Roth

Case-based reasoning in the law

A formal theory of reasoning by case comparison

PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit Maastricht,
op gezag van de Rector Magnificus, Prof. dr. A.C. Nieuwenhuijzen Kruseman,
volgens het besluit van het College van Decanen,
in het openbaar te verdedigen
op woensdag 26 november 2003 om 14.00 uur

door

Abraham Cornelis Roth

Promotor:

Prof. dr. H.F.M. Crombag

Co-promotor:

Dr. H.B. Verheij

Beoordelingscommissie:

Prof. mr. N.H.M. Roos

Prof. dr. K.D. Ashley (University of Pittsburgh, PA, USA)

Dr. F.M.P. van Dun

Dr. mr. J.C. Hage

Prof. dr. H.J. van den Herik

Preface

It has been a great adventure for me to carry out the research reported in this book. Having a background as a physicist, I started working in the field of Artificial Intelligence and Law, on a research project about legal case-based reasoning. Although I never regretted my step into a completely different field of research, it has been an enormous challenge for me to overcome the difficulties that such a change can bring about.

The aim of the research project was to arrive at a logical model of case-based reasoning in the law, and my work on the project has resulted in a formal theory of that reasoning strategy. In the theory that I developed, case comparison occupies a central place as a method to decide whether or not a settled case can be followed.

In research, progress tends to be made only slowly, and the present work is no exception to that rule. My ideas on case-based reasoning in the law have evolved only gradually, and they did not mature until the final year of my research appointment. In that last year they finally took shape, thanks to the close co-operation with my co-advisor, Bart Verheij. I thank him for his support and for the many fruitful and inspiring discussions that we had.

I also want to thank Jaap Hage for giving me the opportunity to become an Artificial Intelligence and Law researcher. He was responsible for the original research proposal, and he coached me during the first years of my appointment.

I am especially grateful to my advisor Hans Crombag, who has always supported me at crucial moments. He encouraged me to continue my work, and his reassurance helped me to cope with the difficulties that I encountered in the course of my research.

I thank the staff of the Maastricht Faculty of Law and of the Department of Metajuridica, for giving me the opportunity to finish off my work. In particular, I want to mention Nico Roos and Frits Stroink here.

I thank the secretaries Elly Peters and Argentine Lankkamp of the Department of Metajuridica, for doing administrative work, such as making hotel and travelling arrangements necessary for attending conferences.

I thank all members of the Department of Metajuridica for providing a warm and pleasant working atmosphere. In particular, I want to mention Sef Janssen, Henk Koning, Alex Jettinghoff, Tanja van der Meer, Remco van Rhee and Louis Berkvens.

I thank my former room mate Lyana Francot for our discussions on life in research, and my present room mate Lode Vereeck for his positive attitude and the reassurance that he gave me at the harder moments.

I also want to thank all the junior researchers of my generation at the Faculty of Law, who made my research a somewhat less solitary activity by exchanging experiences, good or bad, with me. I thank them also for each of the nice lunches that we occasionally enjoyed.

I thank my parents for being there when I needed counsel. Their faith in me has helped me to overcome my most serious doubts, and to face the challenges posed by my research. I thank my sister Liesbeth and my brother Pieter for being there at a distance, and for their willingness to stand by my side during the defence of this thesis.

Finally and most importantly, I thank my wife Corrie for her constant loving support and her encouragement at the hardest moments, and my children Lisa and Arthur for making me see the relativity of it all.

THE UNIVERSITY OF CHICAGO
PRESS
to my father

Chapter 1: Introduction	1
1.1 Case-based reasoning in Artificial Intelligence and Law.....	1
1.2 Research questions.....	4
1.3 Outline of the thesis	5
Chapter 2: Case-based reasoning in the law.....	7
2.1 The stare decisis principle.....	8
2.2 Rule extraction and case comparison.....	9
2.3 Two questions: Which conclusions follow and how do they follow?.	14
2.4 Case-based reasoning as a means of conflict resolution	15
2.5 Two perspectives on reasoning by analogy	18
2.6 Which case features are relevant in case comparison?	20
2.7 Settled cases versus problem cases	21
2.8 Departures from stare decisis.....	22
2.9 Starting points of the present research	23
Chapter 3: Case-based reasoning by case comparison: a formal theory	27
3.1 Example with analysis	28
3.1.1 An extended example.....	28
3.1.2 Dialectical arguments	30
3.1.3 An analysis of the example.....	33
3.2 Support and attack.....	37
3.3 Representing cases	40
3.4 The background.....	41
3.5 Reasoning with support and attack	42
3.6 Case comparison	46
3.6.1 Comparing cases.....	46
3.6.2 The comparison basis.....	50
3.6.3 The entangled factor hierarchy	53
3.6.4 Deriving factors from non-factors	55
3.6.5 Which factor statements are relevant?	56
3.6.6 Comparison outcomes.....	63
3.6.7 Summary.....	68
3.7 Which settled cases are relevant?.....	68
3.7.1 Problem cases, settled cases and precedents.....	69
3.7.2 The case base	71
3.7.3 When do precedents conflict?.....	72
3.7.4 Summary	76
3.8 Reasoning by case comparison and reasoning with support and attack	77
3.9 Applications of the theory	86
3.9.1 Analogising and distinguishing	86

3.9.2	Downplaying and emphasising	92
3.10	Discussion	98
3.10.1	Reasoning with hypothetical cases	98
3.10.2	Fact types in cases	99
3.10.3	Overruling precedents.....	99
3.10.4	How to choose the comparison basis?	101
Chapter 4:	Related research	103
4.1	Ashley's HYPO: adversarial case-based reasoning	103
4.1.1	Factors, arguments and counterarguments.....	104
4.1.2	Dimensions and hypotheticals	108
4.1.3	On pointness, case comparison and argument evaluation	110
4.1.4	Reasoning with factors in HYPO: one-step or multi-step?.....	113
4.2	Aleven's CATO: argumentation in case comparison.....	115
4.2.1	The Factor Hierarchy	116
4.2.2	Downplaying and emphasizing distinctions	118
4.2.3	Heuristic criteria	122
4.3	Prakken and Sartor's formal dialogue game	124
4.3.1	Representing cases.....	126
4.3.2	Reasoning with cases as premise introduction.....	129
4.3.3	Case comparison and on pointness	131
4.4	Bench-Capon and Sartor's approach incorporating theories and values	134
4.4.1	Reasoning with values	135
4.4.2	Which conclusions follow?.....	139
4.5	Branting's model of ratio decidendi.....	140
4.5.1	Warrants and explanations.....	140
4.5.2	Case comparison.....	142
4.5.3	Output	145
4.6	Hage's reason-based account of reasoning with cases.....	146
4.7	Henderson and Bench-Capon's dynamic evolution of case law	149
4.8	Other research	151
4.9	Concluding remarks	153
Chapter 5:	Evaluation	157
5.1	Results.....	157
5.2	Evaluation of the model	158
Summary		167
Samenvatting.....		170
References		174
Index		180

Chapter 1: Introduction

1.1 Case-based reasoning in Artificial Intelligence and Law

This book is about case-based reasoning in the law. Case-based reasoning (occasionally abbreviated CBR) is a technique to draw conclusions about cases, by comparing them to cases already settled (cf. Ashley 1992, pp. 113f.).

Case-based reasoning is a broad field of research that covers many different ways of reasoning with cases (Kolodner 1993, pp. 73f.). Much research on case-based reasoning can be positioned within the vast field of Artificial Intelligence (AI), especially where it meets psychology in the interdisciplinary field of cognitive science (Kolodner 1993, p. 29). Legal case-based reasoning is a main stream of research in the field of Artificial Intelligence and Law (AI & Law), a sub field of AI. This sub field has connections with legal theory as well, in particular with the theory of legal reasoning with precedents (MacCormick and Summers 1995; Cross 1977; Llewellyn 1989).

As said much research on case-based reasoning can be positioned within Artificial Intelligence (AI). In AI one attempts to build intelligent agents to perform a wide range of tasks for which intelligence is required (Russell and Norvig 2003, pp. 1f.), from playing chess to producing legal arguments or decisions (HYPO, Ashley 1990; CATO, Alevén 1997; Prakken and Sartor 1998; Herik 1991). The field of AI and its sub-field AI & Law can be divided roughly into two branches. The first involves research that aims at building systems that mimic human intelligence computationally, for instance for the purpose of enhancing our understanding of the human mind (this use of AI has been criticised by Crombag 1994). Research in the second branch of AI focuses on an idealised standard of intelligence, called rationality (cf. Russell and Norvig 2003, p. 2). Legal case-based reasoning can accordingly be modelled in AI & Law with roughly two purposes in mind. One can attempt to copy case-based reasoning as it actually takes place in legal practice, and to this end do an empirical study, for example, of written judicial opinions or case law reports. One can also develop an idealised model of case-based reasoning, producing the conclusions for problem cases that are best by some standard of rationality. To arrive at such 'rationally best' conclusions one will have to take into account, for instance, all relevant settled cases that are available. On the whole the present work is biased more to the second, rationalist branch of AI research than to the first. Its aim was to define the conclusions that follow by case comparison, given a technique to make explicit the dialectical structure of cases (dialectical arguments, Section 3.1). The examples were inspired by an empirical study of judges' decisions in the field of Dutch dismissal law.

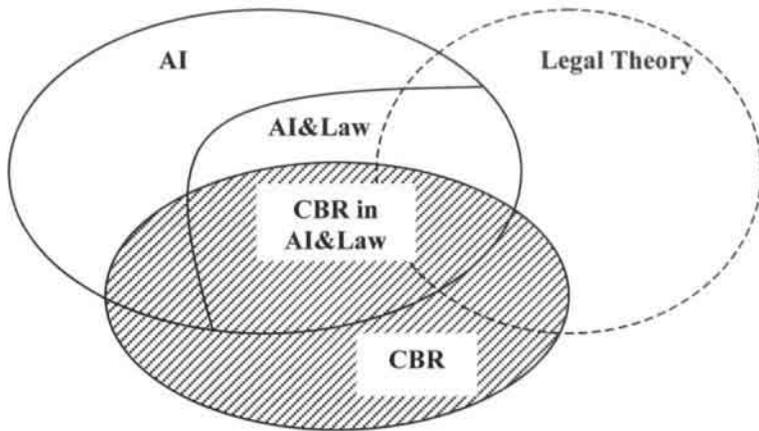
Case-based reasoning in the law can also be studied from a legal-theoretical stance. In legal theory one attempts to gain insight into problems like "What is law?", "How does legal obligation relate to moral obligation?" (Hart 1994, p. 13), or "What is binding in a precedent?" (Peczenik 1997, pp. 461f.). Various answers to questions like these have been given. The legal theorist Hart

(1994, 1983), for instance, took the view that the question what the law is can be answered by so-called rules of recognition (1994, pp. 100f.). Other legal theorists have laid more emphasis on the role of fundamental rights as the ultimate source of law (Dworkin 1977, Ch. 7; Alexy 1985, Kap. 10).

Legal case-based reasoning is one of the main streams of research within the field of Artificial Intelligence and Law (AI & Law). In the field there are roughly two other main streams of research. First, there is research that aims at developing formal models of legal argument, and this kind of research is often theoretically oriented. Some of this work focuses on non-monotonic logics for defeasible legal reasoning (*e.g.*, Hage 1997; Hage 1996; Prakken 1993; Verheij 1996; Verheij 1999b). In other theoretical work on legal argument one attempts to model argumentation formally by means of dialogues (Lodder 1998; Gordon 1995; Prakken 1997; Prakken and Sartor 1996; Hage, Leenes and Lodder 1994).

A second main stream of research in AI & Law aims at investigating the practical use of AI techniques in the law, by developing (prototype) applications of these techniques (see Oskamp and Lodder 2002 for an overview). An example is the system developed by Vey Mestdagh (1997), which operates in the field of Dutch environmental permit law. Another prototype system is ARMOR (Matthijssen 1999, Ch. 5), a system which provides an intelligent interface between lawyers and legal knowledge. A practical application is the Split-Up system that supports parties involved in divorce (Zeleznikow and Stranieri 1995; 1997). Examples of applications tested in legal education are CATO (Aleven 1997) and LITES (Span 1992; 1993). Practically oriented work is also done on computer assisted drafting and testing of legislation (Debaene *et al.* 1999). Recent developments are alternative dispute resolution (Lodder and Huygen 2001) and argument assistance (Verheij 1999a). Practically oriented work is also done on the development of legal ontologies (*e.g.*, Valente 1995), for instance to improve information retrieval (Winkels *et al.* 2000).

In the following figure the position of case-based reasoning research is sketched relative to AI and its sub-field AI & Law, and relative to the field of legal theory. The present work can be positioned within the research to CBR in AI & Law, while there is some overlap with legal theory as well.



A number of different paradigms of case-based reasoning can be distinguished in the field of AI (cf. Ashley 1992, p. 127). In the model-based approach (Koton's Casey 1988), cases are examples explained in terms of a theoretical domain model. Given a new case, the case-based reasoning system determines whether a past explanation applies. In the planning-oriented approach (Hammond's CHEF 1989), cases are records of solutions in past problems, which are used as templates that are mapped onto new problems. In the exemplar-based approach (Bareiss' PROTOS 1988), cases are exemplars that extensionally supplement the definition of vague concepts.

From the perspective of the present book the most important paradigm of case-based reasoning is the precedent-based paradigm, in which a conclusion is argued for by putting forward a similar case where this conclusion was drawn. In the field of AI & Law many researchers have modelled this paradigm of reasoning with cases (*e.g.*, Rissland and Ashley 1987; Ashley 1990; Alevén 1997; Prakken and Sartor 1998; Branting 2000; Skalak and Rissland 1992; Brüninghaus and Ashley 1999a; 1999b).

Much research on case-based reasoning in AI & Law has focused on the question which reasoning patterns (or argument moves) are possible when comparing cases, such as analogising and distinguishing (Ashley 1990; Alevén 1997). Other AI & Law research on case-based reasoning has dealt with the question which conclusions can be drawn from settled cases (Prakken and Sartor 1998; Hage 1997).

The author of this book has continued the research on case-based reasoning with an eye on reasoning patterns in case comparison (Roth 2000b; 2001a; 2001c). The work reported in the book is a continuation and extension of this work on case comparison, in the sense that case comparison is now formalised, and the conclusions that follow are defined. Moreover, some reasoning patterns in case comparison are now formalised, such as analogising and distinguishing.

The work reported here was carried out within the logic-oriented research group at the university of Maastricht, the Netherlands (cf. Verheij's DEFLOG

2000; Hage's Reason-Based Logic 1997; Verheij 1996). A strong point of this group is that it aims at clarifying the concepts typical of the domain of legal reasoning. In Hage's (1997, pp. 115f.) work, for instance, a clear conceptual distinction is made between legal rules and principles. The experience of the Maastricht group has shown that a formal modelling approach is a fruitful research methodology to arrive at a clear conceptualisation of legal reasoning. In line with this, the present work presents a formal theory of reasoning by comparison with settled cases, in which key phenomena like case comparison are made formally precise.

1.2 Research questions

Case-based reasoning will be formalised as a method of adhering to decisions on the basis of a comparison with settled cases. If a settled case is – in some sense – analogous to the problem case, then one can follow the settled case and draw the conclusion. The question thereby is how good the analogy between the cases actually is, and to answer this question the cases have to be compared.

The aim of the present research is to formalise this method of adhering to decisions. The main issue thereby is which conclusions follow on the basis of case comparison, while a second issue is along which patterns of reasoning they follow. As will be seen the comparison makes use of the dialectical structure of cases, in terms of the statements occurring in them that support or attack certain *conclusions*. *The criterion for following a settled case will be that the problem case provides at least as much support for the conclusion as the settled case.* Then by a variant of reasoning a fortiori, the settled case can be followed and the conclusion may come to hold. But how can the dialectical structure of cases be used in their comparison (Section 3.6)? Which settled cases are to be treated as relevant precedents (Section 3.7)? Which conclusions follow by comparison with settled cases, and what if precedents give rise to contradictory conclusions (Section 3.8)? Along which patterns of reasoning do the conclusions follow (Section 3.9)?

Another issue is how the method of case comparison relates to existing approaches to case-based reasoning (Chapter 4). In this connection it is an interesting question how the comparison of cases is formalised in these other approaches. A further question is whether an account is given of the conclusions that follow from settled cases, and of the reasoning patterns along which they follow.

The work described in this book aims at finding answers to the following groups of research questions:

- How can case comparison be formalised as a criterion to determine whether a settled case can be followed? In particular, how can the dialectical structure of cases be used in their comparison?

- Which conclusions can be drawn by comparison with settled cases, and along which patterns of reasoning? In particular, what if precedents contradict each other?
- How does the method of case comparison relate to other approaches to case-based reasoning? In particular, how is case comparison formalised in these approaches? Is there an account of the conclusions that follow?

1.3 Outline of the thesis

The outline of the thesis is as follows. First some important issues concerning reasoning by case comparison are addressed informally (Chapter 2), both from a legal theoretical and from an AI & Law perspective. One of these issues is, for instance, which case features are relevant for the purpose of case comparison. Two methods of adhering to decisions are described, namely the rule extraction method and the case comparison method.

Then a formal theory is developed for reasoning by case comparison (Chapter 3). The theory includes formal definitions of the outcomes of case comparison, and it specifies the conditions under which a settled case can become relevant as a precedent. As will be seen the comparison of cases involves their dialectical structure, that is, statements supporting or attacking conclusions. Moreover, in the proposed model it can also be supported or attacked that a statement supports or attacks a conclusion. In its account of case comparison, the model explicitly acknowledges that in the law it is contingent which case features are relevant for comparing cases. It is also shown how one can accommodate well-known reasoning patterns in case comparison, such as analogising and distinguishing.

After the formal theory of case comparison has been presented, it is investigated how the model relates to other approaches to case-based reasoning (Chapter 4). As will be seen the proposed model offers a generalisation of the account of case comparison given by these approaches. Furthermore, in contrast to other work with an emphasis on case comparison, the proposed model precisely specifies the conclusions that follow by comparison with settled cases. It is also shown that in other research it is not acknowledged explicitly that cases are compared relative to the case features that are deemed relevant.

Finally, the theory is critically evaluated (Chapter 5). A number of reasoning patterns are discussed which are not captured by the proposed model. Among these one has, for instance, patterns of reasoning that are normally left to judicial discretion, such as deviating from undesirable conclusions that may follow from precedents. The evaluation also addresses the question how such patterns of reasoning may be dealt with within the proposed model. In addition a number of subjects for future research are recommended, like modelling reasoning with hypothetical cases.

In summary, the outline of the thesis is as follows:

- Some important issues concerning reasoning by case comparison are addressed informally, both from a legal theoretical and from an AI & Law perspective (Chapter 2).
- A formal theory is developed as a model of reasoning by case comparison (Chapter 3), which among others includes a formalisation of case comparison (Section 3.6).
- The conclusions are specified that follow from a comparison with all settled cases. In particular, the situation of conflicting precedents is dealt with (Section 3.7 and 3.8).
- The model is compared to other approaches to case-based reasoning (Chapter 4). The focus will be on the way case comparison is dealt with, and on the question whether and how the conclusions that follow are accounted for.
- The model is critically evaluated and possible points of criticism are discussed (Chapter 5). In particular, the evaluation addresses ways of reasoning not captured by the model, and possible deviations from conclusions drawn from settled cases.

Chapter 2: Case-based reasoning in the law

Case-based reasoning is a common type of argumentation in the law, in which legal conclusions are supported with decided cases. If some decided case is sufficiently similar to the case at hand, then under the doctrine of *stare decisis* one should not depart from that decision, and the same conclusion should hold. Case-based reasoning is a widespread practice in areas of common law, but it is becoming more and more popular under statutory law as well (MacCormick *et al.* 1997, pp. 11-12; Wiarda 1999, pp. 125-127).

This chapter puts forward two important questions in relation with case-based reasoning. The first question is which conclusions can be established on the basis of this reasoning strategy. In this connection one primarily thinks of conclusions that are copied from decided cases as authoritative examples.

The second question concerning case-based reasoning is along which patterns of reasoning one arrives at the conclusions. Patterns of reasoning typically found in case-based reasoning are 'analogising' and 'distinguishing'. The former involves pointing out similarities between a decided case and the case at hand, while in the latter the relevant differences between the cases are stressed.

There are at least two different methods to derive conclusions by adhering to prior decisions, and these are called the 'rule extraction method' and the 'case comparison method,' respectively. They will be described informally in this chapter, as well as the typical patterns of reasoning that they involve.

In modelling case-based reasoning one should address both questions indicated above. This means that one should not only be able to account for the conclusions that can be established by case-based reasoning, but also for the patterns of reasoning that are used to this end. Regarding these patterns of reasoning the present work will take the case comparison method as a starting point, because this method is considered more typical of the actual practice in case-based reasoning.

The outline of this chapter is as follows. First the principle of *stare decisis* is discussed, and in particular the reasons for adopting this principle (Section 2.1). Next the methods of rule extraction and case comparison are described in an informal way (Section 2.2). After that two questions are discussed that must be addressed by any model of case-based reasoning (Section 2.3). The first question is which conclusions can be drawn on the basis of decided cases, while the second is along which patterns of reasoning they follow. Then the discussion focuses on the case comparison method. The resolution of conflicts is presented as the key subject of application of the method (Section 2.4). It is explained that if reasoning by analogy is to be a sound way of following decisions, then in effect it comes down to reasoning a fortiori (Section 2.5). Moreover, the establishment of an analogy is shown to be relative to the case features that are deemed relevant for the purpose of case comparison (Section

2.6). It is also indicated that in case comparison the relevant distinction between cases is not so much that between past and present cases, but rather whether or not a certain issue has been settled by a judicial authority (Section 2.7). After that possible departures from *stare decisis* are briefly discussed (Section 2.8). The last section of the chapter (Section 2.9) summarises the main points and sets forth the starting points of the present research.

2.1 The *stare decisis* principle

The principle of *stare decisis* holds that in settling new cases judges tend to adhere to prior decisions, in the sense that once a question of law is decided, the decision is normally not departed from afterwards. There are a number of reasons to adhere to earlier decisions. First, courts are often formally bound by decisions made by courts higher up in the hierarchy and by their own decisions (cf. Cross 1977, pp. 103f.). In this connection one sometimes recognises several degrees in which a precedent can be binding (cf. Peczenik 1997, pp. 462f.). Apart from formal requirements, decisions are often considered as binding *de facto*, because decisions against precedents are likely to be reversed on appeal (Hage *et al.* 2001, pp. 50-51). As a result, lower courts usually adhere to decisions by higher courts, and this is a reason why decisions by higher courts are often assumed to carry more authority (Burnham 1995, pp. 66).

A second and more fundamental reason for adhering to decisions is that a departure from settled cases would make judicial decisions uncertain and unpredictable, so that one could never be sure about the status of one's legal claims in court. This unpredictability would of course be undesirable from a perspective of legal security, and by adhering to prior decisions the courts promote legal security within their jurisdictions (Bankowski 1997, pp. 486-487; Golding 2001, p. 99). Moreover, it would often require a complete reform of whole parts of case law to depart from settled points of law. Many judges are reluctant to introduce such far-reaching reforms, because it is usually considered the legislature's task to enact them (Raz 1979, pp. 200-201).

A third and closely related reason for adhering to decisions is that it is considered fair and just to treat all individuals equally under the law, and to make no difference between people in similar cases. By adhering to prior decisions the courts contribute to the ideal of legal equality (Bankowski 1997, pp. 487-488).

A fourth reason to adhere to decisions is that for courts this turns out to be an economic and efficient practice (Bankowski 1997, p. 490). By adhering to the prior decision one can reuse the reasoning underlying the earlier result, instead of having to invent new lines of reasoning for each new case. The practice of adhering to decisions of course still requires reasoning, namely case-based reasoning. However, for the judiciary the practice does have the gain that unnecessary work is avoided.

2.2 Rule extraction and case comparison

Above a number of reasons were given why courts tend to adhere to earlier decisions, and in the following it will be discussed how this can be done. Different methods to adhere to decisions have been described in the literature (Llewellyn 1960, pp. 77-84), and two of them will be discussed next. The first method involves the extraction of rules from decided cases, while the second relies on factual comparisons among cases (cf. Summers 1997, p. 387).

Both mechanisms take as a starting point the judge's deliberations, which can be found, for instance, in published law reports. In the first method the judge's reasoning behind the decision is isolated from all that is not relevant. Once the reasons underlying the conclusion have thus been isolated, the reasons and conclusion are generalised (or 'universalised', cf. Hage 1997, p. 47; Hare 1963, Chapter 2; MacCormick 1987, pp. 162f.) into a rule that can explain the outcome.¹ Such rules are often identified with what is called the 'ratio decidendi' of a case, while the unessential parts are usually referred to as 'obiter dicta' (MacCormick 1987, pp. 156-157). As an example, suppose that a steel welder named John is dismissed on the grounds that he has caused considerable damage, and that he has committed a serious act of violence. Suppose, moreover, that a judge has decided that these facts indeed constitute a pressing ground for dismissal. Then this decision can be generalised into the rule that if some employee causes considerable damage and commits a serious act of violence, then there is a pressing ground to dismiss him.

Once such rules have been obtained from decided cases, they can be applied to other cases. In the example just given, for instance, it is obvious that the extracted rule also applies to the case of a bank employee instead of a steel welder who causes considerable damage and commits a serious act of violence. Since this mechanism in effect comes down to extracting rules from concrete decided cases, it is called the 'rule extraction method' for short.

In the field of AI & Law many formalisations of case-based reasoning seem to be based explicitly or implicitly on the rule extraction method (Prakken and Sartor 1998, Bench-Capon, Branting 1991, 2000, Henderson and Bench-Capon 2001). Among legal theorists it is widely accepted as a suitable representation of the practice of case-based reasoning (MacCormick 1987, pp. 169f.; Cross 1977). However, rule extraction has also met some criticism in legal theory (cf. Stone 1968, pp. 267f. who speaks of ratio decidendi as a 'category of illusory reference'). This holds especially for those theorists influenced by legal realism (such as Oliphant 1928).

¹ Sometimes a whole series of cases is 'synthesised' into one rule, and in the United States such rules are often published in so-called 'restatements'. See Henderson and Bench Capon 2000, 2001 in *Artificial Intelligence and Law*, and Burnham 1995, p. 69 and Burton 1985, p. 64 for a legal perspective. Cf. also Dworkin 1977, p. 115-118, who speaks of the construction of a '(...) coherent justification for all (...) precedents' (p. 116, emphasis added).

The second method of adhering to a decision is by assigning the corresponding legal conclusion to the new case as an authoritative example (cf. Ashley 1990, pp. 11-12; Alevén 1997, pp. 58; Oliphant 1928, pp. 144f.). If the decided case is 'sufficiently similar' to the case at hand, then one can argue by analogy and judicial authority that the decided case should be followed and that the same decision should be taken once more. Suppose again, for instance, that a judge has decided that there is a pressing ground to dismiss a steel welder who has caused considerable damage and who has committed a serious act of violence. Suppose, moreover, that in the case at hand the same facts appear, except that the dismissed person is a bank employee instead of a steel welder. Then one can argue that this difference does not matter for the issue whether the dismissal can be voided, and that the case at hand is therefore analogous to the decided case in this respect. Furthermore, one can argue on the basis of the analogy that the decision should therefore be adhered to, so that a pressing ground for dismissal also exists in the case at hand.

When one relies on analogies between cases the problem arises how 'good' the analogy between the cases actually is. To answer this question the two cases need to be compared, and this mechanism is therefore called the 'case comparison method', for short². To establish an analogy one can point out important similarities between the cases, and this is often called 'analogising' (Alevén 1997, pp. 58f.; cf. also Prakken and Sartor 1998, pp. 242-243). To attack the analogy the relevant differences between the cases are stressed, and this is usually called 'distinguishing' (Ashley 1990, pp. 14-15 and 26, Alevén 1997, pp. 58f.). Suppose, for instance, that in a past dismissal case the judge decided that there was a pressing ground for dismissal, because the employee had caused considerable damage to company property and had committed a serious act of violence. Suppose, moreover, that in a new dismissal case the employee only caused considerable damage to company property, and did not commit a serious act of violence. Then one can analogise by pointing out that in both cases considerable damage was done to company property. Distinguishing is possible by pointing out that the employee has committed a serious act of violence in the past case but not in the new one. As a result of this distinction, one cannot adhere to the decision that there is a pressing ground for dismissal (see Section 2.4).

Summarising this description of case comparison, two different parts can be distinguished. First an analogy between a decided case and a new one must be established by case comparison. If the two cases are indeed found to be analogous, then one can adhere to the decision and draw the corresponding legal conclusion. Accordingly, two types of conclusion must be distinguished in the

² Note that case comparison is used here merely to answer the question whether there is an analogy between a given decided case and the case at hand. It does not serve as a heuristic for retrieving similar decided cases from a case base, as is done, for instance, in HYPO (Ashley 1990), CATO (Alevén 1997), and GREBE (Branting 2000).

case comparison method. Conclusions of the first type concern the presence of an analogy between two cases, while the legal conclusions drawn by analogy are of the second type. To distinguish between both types, the formal model will refer to the first type as *outcomes* of case comparison. It can be an outcome of case comparison, for instance, that a settled dismissal case is analogous to a problem case. A legal conclusion may be, for example, that the employee's dismissal can be voided.

In the field of AI & Law the case comparison method is taken as the basis of some prominent contributions to the research on case-based reasoning (Ashley 1990; Aleven 1997; Skalak and Rissland 1992.) The case comparison method underlies some theoretical accounts of case-based reasoning for educational use (Levi 1949, parts I and II; Burton 1985, Chapter 2; Burnham 1995, pp. 70f.). Case comparison is also acknowledged in legal theory as a method of adhering to decisions (Golding 2001, p. 102f.; cf. also Oliphant 1928, p. 152, who speaks of the 'case method' as an objective method for operating *stare decisis*). In accordance with the ordinary view on *stare decisis* (Cross 1977, p. 106), however, the legal theoretical accounts tend to contain elements of rule extraction as well (Raz 1979, pp. 183f.; Cross 1977, pp. 182-188).

The methods of rule extraction and case comparison can be presented schematically as a list of reasoning steps. Of course different analyses resulting in different lists are possible, but presently the following reasoning steps are considered essential in both methods. The first step in rule extraction involves the extraction of one or more rules from decided cases. Among other things this requires a choice of the relevant facts upon which the conclusions were based. This is similar to the first step in case comparison, where a selection is made of the facts relevant for the purpose of comparing cases. The second step in rule extraction is showing that the conditions of an extracted rule are satisfied by the facts of the new case. In case comparison the second step involves the establishment of an analogy between the decided case and the new one. This is comparable to the second step in rule extraction, because it shows that the decision can in principle be adhered to, and that the corresponding conclusion can hold. As a last step in rule extraction one can either apply a rule, or point out an exception to it. In the former situation the rule's conclusion holds, while in the latter it normally does not. This is comparable to the last step in case comparison. This step can either be that the decided case is followed and that its conclusion holds, or that the conclusion does not hold because the decided case is distinguished.

In the figure below the steps involved in the two methods of case-based reasoning are listed. The figure illustrates that there are strong relations between both methods, in the sense that they involve comparable reasoning steps. The second step is different in both methods, however. The establishment of an

analogy is crucial in case comparison, while analogy does not play a role in the rule extraction method as depicted here.³

Rule extraction method		Case comparison method	
(1) Extracting rules from decided cases		(1) Selecting relevant case facts	
(2) Showing that rule conditions are satisfied		(2) Establishing an analogy between cases	
(3a) Applying extracted rules to the case at hand	(3b) Pointing out exceptions to extracted rules	(3a) Following decided cases in the case at hand	(3b) Distinguishing decided cases from the case at hand

From a modelling perspective, in both methods the first step is the most problematic one, and for both methods two difficulties can be distinguished in connection with this step. The first difficulty has to do with the selection of case parts relevant for the decision, while the second is at which level of generality the result of the selection is to be stated.

In rule extraction, for instance, the difficult part is to make a distinction between those parts of the reasoning in a case that were relevant for the final outcome, and those parts that were not. In this connection the relevant parts are traditionally called 'ratio decidendi' (MacCormick 1987, pp. 156-170; Branting 2000, pp. 27f.), while the unessential parts are labelled 'obiter dicta' (MacCormick 1987, pp. 156; Burnham 1995, p. 68). There has been some discussion among legal theorists about the issue how the ratio decidendi of a case should be ascertained. Goodhart (1930; 1959) has taken the view that the ratio decidendi of a case can be determined by identifying the facts that the court considered material in reaching its decision. This view has met serious criticism (Simpson 1959, pp. 29f.), however, especially since different selections of facts deemed material are possible (Stone 1959, p. 608). Given the case of a steel welder who was dismissed for having committed a serious act of violence, for instance, one can wonder whether the person's profession played any role in reaching the decision. If it did the statement of the profession is ratio decidendi, otherwise the statement is part of obiter dicta.

The problem of isolating the ratio decidendi has its counterpart in case comparison, because there a selection must be made between those case features

³ This does not mean that analogy could not play a role in reasoning with rules. Its role has been studied, for instance, in the context of rule application (Hage 1997 p. 189-195; Verheij and Hage 1994).

which are relevant for the comparison result, and those which are not (see Burton 1985, pp. 31-40; Burnham 1995, p. 70; see also Section 2.5 below).

The second difficulty with rule extraction is that even if the essential parts of a judge's reasoning can be isolated, then usually there can still be different rules underlying the same decision (Goodhart 1930, pp. 161f.). In that situation it is not clear whether a rule can be identified as the 'best' one fitting with a judge's actual reasoning, and in particular it is often hard to determine the level of generality with which such a rule is to be stated. The reason for this is that given a suitable rule, it is usually possible to state a more general one that can explain the same outcome (cf. Branting 2000, pp. 30-31; Oliphant 1928, pp. 144-145). Suppose, for instance, that a rule holds that there is a pressing ground to dismiss male employees who causes considerable damage and who act violently. A suitable generalisation of this rule says that if an employee causes considerable damage, then this is enough to constitute a pressing ground for dismissal.

The problem of generality has its counterpart in case comparison, where the central question is what counts as a good analogy between cases. An important question in this connection is at which level of generality the facts of a case are to be described (Stone 1968, pp. 264-274). Since this in turn determines the 'scope' of a decision in terms of new analogous cases, this in fact comes down to determining to what extent the decision can be generalised.⁴ In the example just given, for instance, the question could be whether there is an analogy between a new case and a decided one, if in the decided case a serious act of violence was committed while in the new case the worker merely insulted a colleague. Thus described the cases are different. However, both cases can also be characterised in a more general way, namely by saying that the employee 'offended' someone. And if the facts are stated in this way, then there is arguably an analogy between the cases.

For the purpose of modelling case-based reasoning the present work has chosen the case comparison method, because the reasoning patterns involved in it are considered more typical of the actual practice in case-based reasoning. The reason for this is that case comparison sticks to cases as the primary subject of interest, while in rule extraction rules are introduced. This is why this method seems to be a less direct way to deal with case-based reasoning than case comparison.

⁴ The scope of a decided case is often in part determined by later decisions. See Branting 2000 p. 34f.; Raz 1979 p. 186; Stone 1959, Section VII; cf. also Levi 1949, pp. 9-27 on the gradual development of liability for dangerous products.

2.3 Two questions: Which conclusions follow and how do they follow?

In modelling case-based reasoning one should deal with two questions. The first is which legal conclusions follow on the basis of decided cases, and the second is how these conclusions follow.

The present work has concentrated on the case comparison method, and to answer the first question in relation to this method one has to distinguish between two types of conclusion. The first type concerns the presence of an analogy between two cases, while the legal conclusions drawn by analogy are of the second type. Suppose, for instance, that the conclusion of case comparison is that a new dismissal case is analogous to a decided one, so that the decision can be adhered to. Suppose, moreover, that the legal conclusion was that the dismissal could be voided. Then one will normally adhere to the decision on the basis of the analogy, and the same legal conclusion holds in the new case as in the decided one. In other words, the dismissal can be voided in the new case too. However, case comparison can also lead to the conclusion that no analogy can be established between a new case and a decided one. Then the legal conclusion does not follow on the basis of the decided case, and it may be that the dismissal cannot be voided.

The second question regarding case-based reasoning is how conclusions follow on the basis of decided cases. In other words, the question is which reasoning patterns are actually used to arrive at conclusions. The present work has focused on case comparison, and two typical reasoning patterns involved in that method are 'distinguishing' and 'analogising'. The former involves pointing out relevant differences between cases, while in the latter the similarities are stressed. Suppose, for instance, that a dismissed employee argues that his dismissal should be voided, and suppose that he does this by comparison with an analogous dismissal case where this decision was indeed taken. Then the employee will point at relevant similarities, such that in both cases the employee was highly esteemed as a colleague. The employer will instead emphasise the relevant differences between the cases, such as that a serious act of violence was committed in the new case and not in the decided one. In rule extraction other typical patterns of reasoning are involved, such as applying the rule or pointing out an exception to it, where a possible rule is that the dismissal of highly esteemed employees can be voided.

Ideally a model of case-based reasoning should answer both questions in a satisfactory way. It should not only account for the conclusions following from decided cases, but also for the patterns of reasoning along which the conclusions are reached. Indeed, both questions are addressed by the existing work on case-based reasoning that takes the rule extraction method as a starting point (Prakken and Sartor 1998; Bench-Capon 1999). For defining conclusions this work tends to involve logical or dialectical models, while the typical reasoning patterns with cases are captured by corresponding patterns involving rules (*e.g.*, Prakken and Sartor 1998, pp. 261-263). In the work on case comparison, however, the reasoning patterns are emphasised while the conclusions that

follow are not, or hardly addressed (*e.g.*, Ashley 1990; Aleven 1997). This work tends to involve working computer models that mimic the reasoning as it actually takes place, and do not indicate whether or not a conclusion follows.

The present model aims to provide satisfactory answers for both questions, and as indicated above the case comparison method will thereby be the starting point. To deal with this method the model will specify the conclusions that it produces, and the conditions under which these conclusions follow. The model will also account for the reasoning patterns along which case comparison typically leads to conclusions, such as distinguishing and analogising.

It is the purpose of case comparison to determine whether a problem case is analogous to a settled case, such that the settled case can be followed. However, one often speaks of analogy only if the differences between the cases are such that the settled case cannot be followed directly (Burton 1985, p. 26; Prakken and Sartor 1998, p. 242). In accordance with this view on analogy it suffices to have one significant similarity between cases (*cf.* Ashley 1990, pp. 25f.), and the corresponding analogising move is dealt with in Section 3.9.

However, in this view on analogy no conclusions can be attached directly to settled cases. The reason for this is that the conclusions that follow depend on an additional substantive judgement on the relative importance of differences and similarities (*cf.* Prakken and Sartor 1998, p. 266; Burton 1985, pp. 31f.). If the similarities are more important than the differences, then the settled case can be followed. If the differences are more important, however, then the settled case cannot be followed. In other words, in this view on analogy the similarities would have to be metaphorically 'weighed' against the differences.

The present section aims at defining a slightly different notion of analogy, which allows for following settled cases without having to weigh similarities against differences. As a result, it becomes possible to define the conclusions that follow on the basis of a comparison of cases without resorting to additional contingent weighing information.

2.4 Case-based reasoning as a means of conflict resolution

When reasoning with cases, one is nearly always confronted with conflict situations. The reason for this is that litigants normally have opposite interests, and this holds for past cases and new cases alike. See Ashley on the 'adversarial' nature of legal reasoning (Ashley 1990, p. 3), and Prakken and Sartor on the 'dialectical' structure of cases (Prakken and Sartor 1998, pp. 237-238). An obvious example is the case of a dismissed employee who challenges the dismissal in court. The employee will argue that the dismissal should be voided, for instance on the ground that he is highly esteemed as a colleague. The employer will instead argue for the opposite, for instance by stating that the employee has committed a serious act of violence.

Since litigants have opposing interests, legal decision-making often comes down to settling a conflict in one way or another. As a consequence, decided

cases tend to provide information on how their conflicts have been resolved. Suppose, for instance, that in a decided dismissal case an employee who was highly esteemed as a colleague committed a serious act of violence. Suppose, moreover, that on the basis of these facts the judge decided that the dismissal could be voided. Then apparently the fact that a serious act of violence was committed, counted for less in the eye of the judge than the fact that the dismissed person was highly esteemed as a colleague. This kind of information may help resolve other similar conflicts, and in the present work case-based reasoning is seen as the key mechanism for exploiting this conflict resolving potential.

In the case comparison method the conflict resolving information can be left implicit in the way the case was ultimately decided. The reason for this is that if an analogy with a new conflict is established by case comparison, then this is enough to follow the decided case as an authority and to resolve the new conflict accordingly (Burton 1985, p. 29; Burnham 1995, p. 70). As a result, regarding the resolved conflict it suffices to know which conclusion was ultimately drawn, so that the resolving effect does not have to be made explicit in some form. Suppose, for instance, that a new dismissal case contains precisely the same facts as a decided one, namely that an employee who was highly esteemed as a colleague, was dismissed for having committed a serious act of violence. Suppose, moreover, that in the decided case it was concluded that the dismissal could be voided. Then the two dismissal cases are analogous, and from this it follows immediately that the dismissal can be voided in the new case as well.

If in case comparison the conflict resolving information is left implicit, then this does not mean that the facts that are the source of the conflict can be hidden, for instance by omitting reasons against the conclusion. On the contrary, the factual situations resulting in conflicts are important in the comparison of cases, and the case representations should therefore allow for the occurrence of such conflicts. Suppose in the example above, for instance, that in the decided dismissal case the employee had not committed a serious act of violence, so that the decided case in fact did not give rise to any conflict regarding the issue whether the dismissal could be voided. If a serious act of violence was committed in the new case, then this case is no longer analogous to the decided one.

If one employs the rule extraction method for dealing with decided cases, then the conflict resolving potential of these cases must be made explicit in some form. The reason for this is that according to this method cases are not simply followed, but are transformed into rules that can apply to new cases. Accordingly, conflicts appearing in cases are reduced to conflicts between rules, which can occur in other cases. Suppose again, for instance, that an employee was dismissed for having committed a serious act of violence, and that this employee was highly esteemed as a colleague. Suppose, moreover, that the judge decided that if an employee has committed a serious act of violence, then

as a rule this pleads against the conclusion that the dismissal can be voided, and that being highly esteemed generally pleads for the opposite. Then in effect one has two conflicting rules, and if the judge decided that in the presence of both facts the dismissal can be voided, this provides information on how the conflict between these rules is to be resolved.

Such conflict resolving information can be accounted for in different ways. One way to do this is to reformulate the reasoning into one new rule which takes into account all conditions of the original conflicting rules, and which produces the 'right' conclusion, that is, the conclusion ultimately attached to the decided case. In the example above the reformulated rule would state that if a dismissed employee is highly esteemed as a colleague, then his dismissal can be voided even if he has committed a serious act of violence. In effect the conflict resolving information is then left implicit in the newly stated rule.

Another way to account for the decision is to restate the resolved conflict explicitly as 'priority information' (Prakken and Sartor 1998, p. 256; Bench-Capon 1999). Suppose as above, for instance, that there are two conflicting rules, the first of which says that if one is highly esteemed as a colleague, then one's dismissal can in general be voided. The second rule states the opposite conclusion under the condition that the employee has committed a serious act of violence. Then if the decision held that the dismissal can be voided, this can be explained by asserting that the first rule had priority over the second. The explicit information thus obtained can be used in new cases where the same conflicting rules apply. In a new dismissal case with a highly esteemed employee who commits a serious act of violence, for instance, this priority information would once more give priority to the first rule just stated, so that the dismissal can be voided again.

If one employs case comparison the conflict resolving information does not have to be made explicit, but this way of using the information is not precluded by this method. It is possible, for instance, to treat the conflict resolving effect as a kind of weighing information on sets of reasons for and against a conclusion. If this is done, then case comparison comes down to comparing the sets of reasons involved (cf. Hage 1997, pp. 185-187 in combination with pp. 203-204). Sometimes this allows for arguments *a fortiori* to the effect that one set of reasons outweighs the other. Suppose again, for instance, that a new dismissal case is compared to a decided one where a highly esteemed employee had been dismissed for having committed a serious act of violence, and where the decision held that the dismissal could be voided. Suppose, moreover, that in both cases the same facts appear, except that in the new case the employee has always done an impeccable job. Then the decision can be interpreted such that being highly esteemed is a reason for the conclusion that the dismissal can be voided, while having committed a serious act of violence is a reason against it. Moreover, the former fact obviously outweighs the latter as a reason regarding the issue whether the dismissal can be voided. If the fact that one has always done an impeccable job is interpreted as another reason for the conclusion that

the dismissal can be voided, then it is obvious that in the new case the reasons for the conclusion outweigh those against.

In sum, cases often represent conflict situations and case-based reasoning is a technique for resolving them with the help of decided cases. The conflict resolving potential of decided cases can be exploited in two ways, depending on the method of case-based reasoning that is used. The present model will be about the case comparison method, and it will formalise this method in such a way that the conflict resolving effect does not have to be made explicit. The reason for this is that in the present model one simply follows a decided case if the case at hand is – in some sense – ‘analogous’ to it.

But when is the case at hand ‘analogous’ to a decided case, in the sense that the decided case can be followed in the problem case? This is the subject of the next section.

2.5 Two perspectives on reasoning by analogy

It is the purpose of case comparison to determine whether a decided case can be followed in the case at hand. If the case at hand is – in some sense – analogous to the decided case, then the decision can be adhered to, and the conclusion can follow.

There are at least two possible perspectives on analogy as a basis for adherence to decisions. According to the first perspective one should determine the significant factual similarities and differences between the case at hand and the decided case (Burton 1985, pp. 29f.; Cross 1977, pp. 182f.). The cases are analogous, then, if the similarities between them are more important than the differences (Burton 1985, p. 31). As a result, in this view the establishment of an analogy requires a step in which the significant similarities and differences are ‘weighed’ against each other (cf. Prakken and Sartor 1998, p. 267). Suppose, for instance, that a dismissed person cites a dismissal case for the conclusion that the dismissal could be voided. Suppose further that there is one significant similarity between the cases, namely that the employee is highly esteemed as a colleague. Suppose finally that there is one significant difference between the cases too, namely that in the decided case the dismissal had not affected the working atmosphere, while it had in the case at hand. Then one must decide which is more important: The difference that the working atmosphere had been affected in the decided case and not in the case at hand, or the similarity that in both cases the employee is highly esteemed as a colleague? If the former is deemed to be more important than the latter, then the decided case cannot be followed. If the latter is deemed more important, though, then the decided case can be followed, and one can draw the conclusion that the dismissal can be voided.

In AI & Law the research that takes this view on analogy tends to lay more emphasis on the reasoning patterns in case comparison than on the conclusions that follow (Ashley 1990, see Section 4.1; Alevén 1997, see Section 4.2; Branting 2000, see Section 4.5). Accordingly, in this research the intermediate

step of weighing similarities against differences is not taken. Instead the similarities and differences are used to account, for instance, for argument moves in case comparison (distinguishing or analogising, Ashley 1990, pp. 159f.; Alevén 1997, pp. 58f.).

The second perspective is that reasoning by analogy comes down to reasoning a fortiori. In this view there must not be any significant difference between the decided case and the case at hand. A significant difference is a difference, then, that tends to make the decided case stronger for the conclusion for which it is cited than the case at hand (Ashley 1990, pp. 66-68 and 70-71; Alevén 1997, p. 60). Suppose, for instance, that a decided dismissal case is cited for the conclusion that the dismissal can be voided. Suppose, moreover, that in the decided case the dismissal had not affected the working atmosphere, while this cannot be said of the case at hand. Then if the conclusion is at issue whether the dismissal can be voided, the statement on working atmosphere is a significant difference between the cases. In the view on reasoning by analogy as a kind of reasoning a fortiori, there may be differences between the cases that are not significant. Such differences typically make the case at hand stronger for the conclusion than the cited case. Suppose, for instance, that in a dismissal case the employee is highly esteemed as a colleague, and suppose this is not true of a decided case that is cited for the conclusion that the dismissal can be voided. Then this difference between the cases is not a significant one, because it tends to make the case at hand stronger than the cited case for the conclusion that the dismissal can be voided.

In the perspective on reasoning by analogy as reasoning a fortiori, there must not be differences that tend to make the decided case stronger for its conclusion than the case at hand. In other words, the case at hand then provides at least as much support for the conclusion as the decided case. As a result, the conclusion can then be drawn immediately by a comparison with the decided case. This in contrast with reasoning by analogy according to the first perspective sketched above, which requires an intermediate step in which the significant similarities are weighed against the significant differences. From a modelling perspective it is convenient that this intermediate weighing step can be skipped, because then no additional information is needed concerning the relative significance of similarities and differences. In other words, conclusions can follow then immediately by a comparison with decided cases. This explains why some work in AI & Law deals with case-based reasoning exclusively as a kind of reasoning a fortiori (Hage 1997, pp. 185-187; Bench-Capon 1999).

The proposed model will also treat reasoning by analogy as reasoning a fortiori. This is reflected in the model's definition of precedents, that is, decided cases that are relevant with respect to some case at hand (Definition 12, Section 3.7). According to this definition, a decided case can only be relevant as a precedent for a conclusion, if the case at hand provides at least as much support for the

conclusion. Then by reasoning a fortiori, the conclusion can follow in the case at hand as well.

But when does the case at hand provide at least as much support for a conclusion as the decided case? It is clear that if the two cases are identical in all 'relevant' aspects, then the decision can be adhered to and the conclusion can follow. However, this raises the question how these 'relevant' case features are to be determined. This is the subject of the next section.

2.6 Which case features are relevant in case comparison?

Cases contain facts, and case comparison involves comparing sets of facts. Not all facts in cases play a role in the comparison with other cases (see Burton 1985, pp. 31-40 and Burnham 1985, p. 70 on the 'problem of importance' in connection with analogy; cf. in this connection also Ashley 2002, pp. 172f. on relevance criteria). An employee's sex, for instance, is normally not relevant for the presence of a pressing ground for dismissal. Accordingly, it will usually not play a role either in the comparison of dismissal cases. However, it may sometimes be better to make a distinction between women and men in case of dismissal, for instance from the perspective of women's emancipation. If this distinction is indeed made, the employee's sex can play a role in the comparison of dismissal cases. As a consequence, it may happen that a women's dismissal case is not considered comparable to the case of a male employee, even if both cases are the same save for the person's sex.

The comparison results depend on the particular selection of case facts that is made to this end. Suppose, for instance, that a comparison is made between a new dismissal case and a past one in which the dismissal could not be voided. Furthermore, let the two cases be identical, except that in the new case a superior was insulted, and not in the decided one. Then the question is whether the two cases are analogous with respect to the conclusion that the dismissal can be voided. If regarding this conclusion it is deemed relevant whether a superior was insulted, then the two cases are obviously not analogous and the decided case cannot be followed. However, if it is considered irrelevant whether a superior was insulted, then the two cases are analogous with respect to the conclusion that the dismissal can be voided. As a result, the decided case can then be followed and its conclusion holds again.

It follows that the comparison of cases presupposes a selection of case facts deemed relevant for this purpose.⁵ Moreover, if case comparison is to be a method that can be equally applied to all cases, then the selection criteria should be uniform and independent of the cases being compared. Suppose, for instance,

⁵ It must be noted that in practice the situation is often reversed, that is, what is relevant is determined by what judges consider to be analogous cases. See Raz 1979 p. 186; cf. also Levi 1949 p. 3-6, who speaks of the 'creation of rules' out of analogies. Whether two cases are analogous then depends on what in the eye of the judge is a reasonable and fair outcome of the case at hand.

that insulting a superior is considered relevant for comparing two dismissal cases, because this generally supports the conclusion that there is a pressing ground for dismissal. Then it must be considered relevant for comparing all other cases where dismissal is at stake.

If case comparison presupposes a uniform set of criteria for selecting case facts, this does not mean that the set itself would be beyond debate. On the contrary, the contents of the set will often be debatable, and this holds for the results of comparison as well. In principle it can always be questioned, for instance, whether it is relevant for dismissal that a superior was insulted. If the relevance is indeed disputed, then it can be supported, for instance, by considering that offences negatively affect the working atmosphere.

To make explicit that case comparison involves selection criteria that are in principle subject of debate themselves, one can state the comparison results and their conclusions relative to this set. Take the example above, for instance, where two cases are compared with respect to the issue whether a pressing ground for dismissal exists. The two cases are analogous relative to a selection that does not include the fact that a superior was insulted. There is no analogy between the cases, however, if the fact that a superior was insulted is considered relevant for the purpose of case comparison.

2.7 Settled cases versus problem cases

In the foregoing discussion of the case comparison method, a distinction was made between 'new cases' on the one hand and 'decided cases' on the other, whereby these decisions were taken in the past. Moreover, case comparison was presented as a means to resolve conflicts in new cases on the basis of past and decided cases. To this end a past case was compared with a new case and if an analogy between the cases could thus be established, then one could adhere to the decision as a legal authority and resolve the conflict accordingly.

In general, however, the picture is not so simple. A complication can be that it may be that some possible issues have not been addressed in a past case. This can happen if these issues were not raised by the parties, and were not relevant for the final decision either. In that situation one may wonder afterwards how these issues would have been resolved if they had actually been raised in court. Suppose, for instance, that a highly esteemed employee is dismissed for having caused considerable damage to company property. Suppose, moreover, that the main issue was whether the dismissal could be voided, but that parties did not raise the issue whether the employer is entitled to compensation for the damage. In situations like this one could ask how such unsettled issues could have been resolved, and one could use case comparison in an attempt to answer this question. In effect case comparison is then used to resolve issues in past cases which are only raised afterwards. Although these past cases themselves are of course not decided afterwards by means of case comparison, one can say that they are – in a sense – being reinterpreted in the light of other cases.

While some issues may have been left unsettled in past cases, conversely it is conceivable that some of the issues raised in a new case have already been settled in an authoritative way by the judge. It may be, for instance, that if an employee gives rise to a dismissal dispute for having caused considerable damage, there is already agreement on the issue whether the employer should receive compensation for this.

In case comparison past and new cases can thus appear in two roles. The first role is that of a case where some conflict still is to be settled, and the second is the role of a case as a legal authority which helps to resolve other cases. For the mechanism of case comparison it apparently does not matter whether a case is a past or a new one, but only whether or not some possible issue was resolved. It is possible, for instance, that a dismissal could not be voided because considerable damage had been done, while the issue remained open whether the employee was liable to compensate for this damage.

This suggests that the relevant distinction⁶ in case-based reasoning is not between 'past' and 'new' cases. The relevant distinction is between cases that have not yet been settled on a point, and those that have. More precisely, for case comparison the distinction is always relative to the issue that is at stake. Moreover, one should distinguish between cases where an issue is still problematic, and cases where the problem is settled. Shortly, the difference is thus between 'problem cases' and 'settled cases', all relative to an issue. If the issue is whether an employee's dismissal can be voided, for instance, then the cases involved will be categorised into cases where this issue is problematic, and cases where the issue has been settled. This distinction between settled cases and problem cases is made formally precise in the model proposed, as will be seen in the next chapter.

2.8 Departures from *stare decisis*

Courts normally stick to the principle of *stare decisis*. If a decision can be adhered to, then the corresponding conclusion is normally drawn. A court may have good grounds to decide against the principle of *stare decisis*, however, for instance if this seems desirable in the light of values or goals promoted by the system of law. In this connection at least two possibilities can be distinguished that are relevant for the present research.

First, it may happen that a court expressly or implicitly overrules a decided case (cf. Cross 1977, pp. 126-127), with the intention to enforce a change among the case features that must be treated as relevant for case comparison. As an example of such a change, suppose that from the perspective of women's

⁶ Sometimes 'hypothetical' cases are distinguished as a separate category of cases useful for case-based arguments (cf. the 'hypotheticals' in HYPO, Ashley 1990 p. 35-36). A hypothetical case can play a role, for instance, in arguments showing that a certain outcome would have absurd and unacceptable consequences for other cases ('consequentialist reasoning', MacCormick 1997 pp. 129f.).

emancipation, a court believes that in dismissal cases the employee's sex should be treated as relevant, even though in earlier decisions it was not considered as such. Then if the court has the appropriate authority, it could simply rule that an employee's sex should always be taken into consideration in all dismissal cases (cf. Raz's distinguishing, 1979, pp. 185f.). Note that in such situations a court acts almost like a legislature.

Overruling in this sense should be distinguished from deciding a new case in ignorance of a precedent (that is, deciding *per incuriam*, Cross 1977, pp. 143f.). On the contrary, it is a deliberate action taken by the court (p. 124), with the intention not only to avoid an undesirable conclusion, but also to introduce a change among the case features relevant for case comparison. The procedural authority to introduce such changes will typically be attributed to the highest courts only. Accordingly, such enforced changes often have consequences for all other legal decision makers.

There is a second, and more restricted, way in which a judge can depart from *stare decisis*. This happens if the judge wants to make a single exception to the principle, for instance because it is believed that the conclusion would be undesirable in the concrete case at hand. It may be, for instance, that the judge decides that if a dismissed employee is near retirement, then the consequences of a dismissal would be disproportionate to the goal served by it. Then the judge may want to set aside a precedent to the effect that the dismissal cannot be voided. In situations like this the decision is not intended to set a new precedent, while moreover it is not questioned which case features should be treated as relevant in case comparison. Accordingly, exceptions like this have little or no consequences for future decisions. In other words, they have at most an undermining effect on the precedent that is set aside (Cross 1977, pp. 127f.).

In sum, several departures are conceivable from the *stare decisis* principle. One possibility is that the court aims to enforce some desirable change among the case features that are to be treated as relevant for case comparison. Such a departure has consequences for future decisions. Another possibility is that one exception to *stare decisis* is made in a single case by setting aside a precedent, without putting at issue which case features are relevant for case comparison.

2.9 Starting points of the present research

In the discussion above a number of observations were made on legal case-based reasoning. In this section the main parts of the foregoing are summarised, and the starting points are set forth that were adopted in carrying out the present research.

There are roughly two methods to adhere to decisions in accordance with the *stare decisis* principle, namely by rule extraction and by case comparison. The primary subject of the present research is the case comparison method, of which a formalisation is presented in the next chapter.

A starting point thereby is that two important questions should be addressed in relation with case-based reasoning. The first is which legal conclu-

sions eventually follow on the basis of decided cases, while the second is along which patterns of reasoning this is achieved. Two typical reasoning patterns occurring in the case comparison method are analogising and distinguishing.

Another starting point is that case-based reasoning is particularly useful for resolving conflicts in cases, with the help of decided cases where a conflict has been settled. If there is no conflict in a case with respect to some defensible claim, then in the present model the corresponding conclusion will be derivable without having to resort to decided cases.

To account for the conclusions derived with case comparison, a formal theory is presented that defines a number of concepts necessary to deal with reasoning by case comparison, such as 'case' and 'precedent'. For instance, cases will be defined as sets of facts that can be expressed in a formal language. For reasons indicated above the conceptualisation makes a distinction between problem cases and settled cases, while new and past cases do not appear as separate concepts.

A central starting point of the formal theory is that case comparison involves establishing analogies between cases, and that legal conclusions follow by adhering to decisions in analogous decided cases. The model will therefore make a distinction between on the one hand the question whether two cases are analogous for a claim, and on the other the conclusions that can be attached to the presence of an analogy. Case comparison can lead to the conclusion that two cases are analogous relative to a certain claim, and legal conclusions follow on the basis of such analogies.

The establishment of analogies on the basis of case comparison will therefore occupy a central place in the formal theory of reasoning by case comparison, and the conditions will be specified under which such an analogy is obtained. Apart from this, the theory will indicate the legal conclusions that can be attached to the presence of an analogy.

Reasoning by analogy will be modelled as reasoning *a fortiori*. This kind of reasoning is only possible if there are no significant differences between the problem case and the settled case, that is, differences that tend to make the settled case provide more support for the conclusion than the problem case. As a result, conclusions can follow immediately by comparison with settled cases, without having to weigh significant similarities against significant differences. In accordance with this, in the proposed model a settled case can only be relevant as a precedent with respect to a settled case if the problem case provides at least as much support for the conclusion as the settled case.

For the purpose of reasoning *a fortiori*, the question must be answered whether the case at hand provides at least as much support for the conclusion as the decided case. The answer to this question can be obtained by comparing the cases involved, and the outcome of this comparison depends on the case features that are deemed relevant for the purpose of case comparison. The dependence of comparison outcomes on relevant case features is explicitly acknowledged in the proposed model. In particular, it is possible to state the

relevant case features at different levels of abstraction. It is possible, for instance, to count as relevant that an employee caused considerable damage to company property, but one may also focus on more concrete things, such as that the employee cracked up a company car.

Besides the conclusions that follow by case comparison, a number of patterns of reasoning pertaining to it are accounted for. Analogising and distinguishing, for instance, are accommodated by defining the case features with which one can analogise or distinguish a settled case from a problem case, respectively.

Finally, reasoning by case comparison is but one method to adhere to decisions. Moreover, it is conceivable that case-based arguments are overruled on the basis of other considerations, such as, for instance, arguments appealing to judicial authority. Some related work on case-based reasoning is therefore discussed and compared to the proposed model, and possibilities are addressed to deviate from conclusions derived by case comparison.

Chapter 3: Case-based reasoning by case comparison: a formal theory

The main objective of this chapter is to specify precisely the conclusions that follow on the basis of the case comparison method, and to this end a formal theory of the method is formulated. A secondary aim is to account for reasoning patterns involved in case comparison, such as analogising and distinguishing.

The case comparison method is a way of adhering to decisions on the basis of analogy. If an analogy between a problem case and a settled case can be established, then on the basis of judicial authority one can argue that the same conclusion should hold. The method can be used for resolving conflicts of reasons in cases, by comparison with settled cases where a conflict was resolved by a judge.

A formal definition is given of what counts as an analogy between a problem case and a settled one. Intuitively this definition holds that there is an analogy if the problem case provides at least as much support for the disputed conclusion as the settled one. By a variant of reasoning *a fortiori*, the conclusion can then hold in the problem case as well.

To determine whether the problem case indeed provides at least as much support for the disputed conclusion as the settled case, a comparison is made of the complete argumentation relevant for the conclusion. This argumentation includes the statements of fact upon which the conclusion can be founded, but also intermediate conclusions. Besides this, the argumentation allows for conflicting reasons, so that case comparison can be used as a means of conflict resolution. In other words, it does not only involve reasons for conclusions, but also reasons against. Moreover, in contrast to other work on case comparison, the model does not take for granted which statements are reasons for or against conclusions. Accordingly, it can itself be supported and attacked that a statement is a reason for or against a conclusion.

An important feature of the present model is, that it is acknowledged explicitly that in the law it depends on a choice which case features are relevant in comparing cases, and that this choice is to some extent contingent. In accordance with this, the relevant case features are given by what is called a *comparison basis*, and the comparison outcomes are stated relative to this comparison basis.

To deal with case comparison along the lines sketched, a number of concepts have to be clarified. For instance, what is a case? When is a case a problem case? What does it mean that a case is settled? What does it mean if the problem case provides at least as much support for the disputed conclusion as a settled case? Which case facts are relevant for this question? When is a settled case relevant for some problem case? What if the set of decisions leads to contradictory conclusions?

These and other questions are addressed by defining the concepts used for dealing with case-based reasoning by case comparison, such as 'case' and

'precedent'. Moreover, the conclusions are specified that follow by case comparison from a given set of cases.

The outline of the chapter is as follows. First an extended example is presented informally, involving a number of dismissal cases (Section 3.1). A method is introduced to represent cases and the argumentation that is possible in them (Sections 3.1, 3.2 and 3.3). Besides this, the conclusions are defined that follow without resorting to precedents (Sections 3.4 and 3.5). Then the method of reasoning by case comparison is exposed, and the outcomes of case comparison are formally defined (Section 3.6). After that a criterion is formulated for when settled cases are relevant as precedents (Section 3.7), and the conclusions following by case comparison are defined (Section 3.9). It is further shown how a number of reasoning patterns in case comparison can be accommodated in the proposed framework, such as analogising and distinguishing (Section 3.9). The chapter is closed with a number of points not addressed by the theory, and a discussion of the theory's capability to deal with these points (Section 3.10).

3.1 Example with analysis

In this section an extended example is presented first. After that the example is systematically analysed using graphical representations of argumentation, called dialectical arguments.

3.1.1 An extended example

Consider the following dismissal case.

Al is a bank employee. One day he causes considerable damage to company property by dropping an expensive computer, despite a warning in advance to be careful. Al is then dismissed for having caused considerable damage to company property, despite a warning in advance.

He challenges the dismissal in court, arguing that it can be voided. He supports this conclusion by pointing out that he is highly esteemed as a colleague. He also states that he has substantial interests in keeping his job, and he states that this is supported by pointing out that he has a family to maintain. He further supports the latter statement by pointing out that his children are planning to go to university, and that his house is mortgaged.

The employer states that Al has caused considerable damage to company property, and he states that this attacks the conclusion that the dismissal can be voided. He further supports this statement by pointing out that a warning was given in advance. Besides the employer says that Al's wife has a good income, attacking Al's statement that by pointing out that he has a family to maintain, he supports the conclusion that he has substantial interests in keeping his job. The

employer also says that Al once deceived him with a forged diploma, and he states that this attacks the conclusion that the dismissal can be voided. He further supports the latter statement by pointing out that the diploma was very relevant for the job.

Besides arguing directly for the conclusion that his dismissal can be voided, Al also cites a similar dismissal case, in which the decision was taken that the dismissal could be voided:

Bob was a steel welder. One day he caused considerable damage to company property by dropping a pipe, despite a warning in advance to be careful. Bob was then dismissed for having caused considerable damage to company property, despite a warning in advance.

He challenged the dismissal in court, arguing that it could be voided. He supported this conclusion by pointing out that he was highly esteemed as a colleague. He also stated that he had substantial interests in keeping his job, supporting this by pointing out that he had a family to maintain. He further supported the latter statement by pointing out that his children were planning to go to university.

The employer attacked the conclusion that the dismissal could be voided, by pointing out that Bob had a criminal record. He also stated that Bob had caused considerable damage to company property, and he stated that this attacked the conclusion that the dismissal could be voided. He further supported this statement by pointing out that a warning was given in advance. The employer also said that Bob once deceived him with a forged diploma, and he stated that this attacked the conclusion that the dismissal can be voided. He further supported the latter statement by pointing out that the diploma was very relevant for the job. Besides, the employer said that Bob's wife had a good income, attacking Bob's statement that by pointing out that he had a family to maintain, he supported the conclusion that he had substantial interests in keeping his job.

The judge decided that the dismissal could be voided.

There may be decided cases with an opposite conclusion. An example is the following case:

Chris worked for a company as a salesman. One day he caused considerable damage to company property by cracking up a company car. Chris was then dismissed for having caused considerable damage to company property.

He challenged the dismissal in court, arguing that it could be voided. He supported this conclusion by pointing out that he was highly esteemed as a colleague and that the dismissal had not affected

the working atmosphere. He also stated that he had substantial interests in keeping his job, and he stated that this was supported by pointing out that he had a family to maintain. He further supported the latter statement by pointing out that his children were planning to go to university, and that his house was mortgaged.

The employer attacked the conclusion that the dismissal could be voided, by stating that Chris had caused considerable damage to company property. The employer further said that Chris once deceived him with a forged diploma, and he stated that this attacked the conclusion that the dismissal could be voided. He further supported the latter statement by pointing out that the diploma was very relevant for the job.

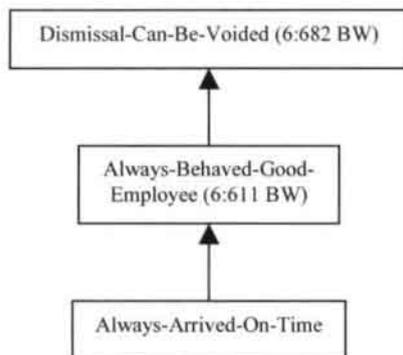
The judge decided that the dismissal could not be voided.

Given these two decided cases, how should the issue be resolved whether Al's dismissal can be voided? Let us analyse the cases systematically to answer this question.

3.1.2 *Dialectical arguments*

To arrive at a systematic analysis of cases, it is convenient to have a graphical representation of the argumentation in the cases. In particular, it is handy to represent graphically that statements support or attack conclusions.

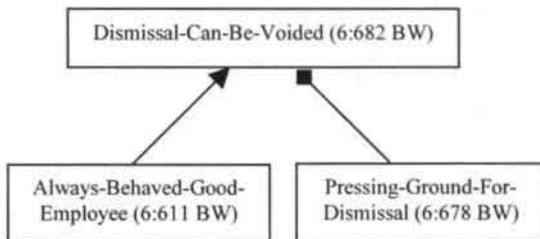
To this end tree-like structures are introduced, called *dialectical arguments*; cf. Verheij's (2000; 2001) naïve dialectical arguments and Loui's (1997) and Loui and Norman's (1995, p. 164) records of disputation. Dialectical arguments consist of statements that support or attack other statements. These support and attack relations between statements are represented by arrows. A simplified example for the domain of Dutch dismissal law is in the following figure.



At the top of this figure one finds a rectangle with the legal conclusion that the dismissal can be voided, to which a judge can decide on the basis of article 6:682 of the Dutch Civil Code (art. 6:682 BW). The rectangle in the middle

contains the statement that the dismissed person has always behaved like a good employee, a general obligation for employees that is codified in article 6:611 of the Dutch Civil Code (art. 6:611 BW). An arrow upwards indicates that the conclusion that the dismissal can be voided is supported by the statement that the dismissed person has always behaved like a good employee. This statement is in turn supported by the statement that the employee always arrived on time for work.

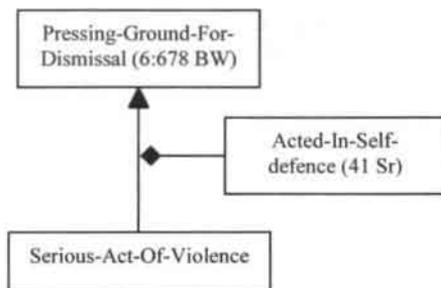
Conclusions cannot only be supported, but they can be attacked as well. Attacks are represented by arrows ending in a solid square. In the following figure one finds an example of this, again from the domain of Dutch dismissal law.



Here the conclusion that the dismissal can be voided is attacked by the statement that there is a pressing ground for dismissal according to article 6:678 of the Dutch Civil Code (art. 6:678 BW).

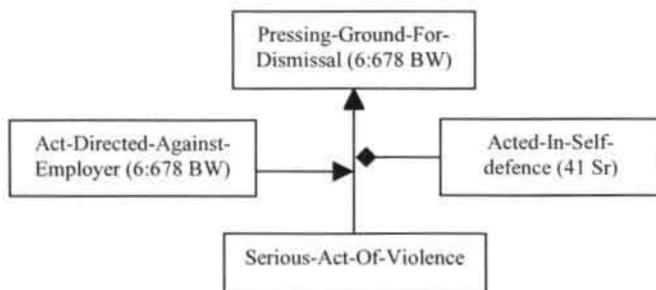
It is a key feature of the present model that it can also be supported and attacked that a statement supports or attacks a conclusion. Attacking that a statement supports or attacks a conclusion is akin to what is often called an undercutting exception in the literature (Hage 1997, p. 166; Verheij 1996, pp. 200-201; Pollock 1995, pp. 41 and 86; Pollock 1987, p. 485). Supporting that a statement supports or attacks a conclusion is comparable to Toulmin's backing of a warrant (1958, pp. 98f.).

As said, in the present model it can itself be supported and attacked that a statement supports or attacks a conclusion. A step forward to deal with this is to treat it as a statement itself that the conclusion is supported or attacked (cf. Verheij's DEFLOG, 2000, pp. 5f.; see also Verheij 1999a, pp. 45f.). Accordingly, one can represent by an arrow pointing at another arrow that it is supported or attacked that a statement supports or attacks a conclusion. This gives rise to a kind of entanglement of dialectical arguments (Roth 2001b, pp. 31-33). An example is in the following figure.



At the top of this figure one finds the conclusion that there is a pressing ground for dismissal according to article 6:678 of the Dutch Civil Code. The conclusion is supported by the statement that the employee committed a serious act of violence. However, it is attacked that having committed a serious act of violence supports that there is a pressing ground for dismissal. The attacking statement is that the employee acted in self-defence, which is a general ground of justification according to article 41 of the Dutch Penal Code (art. 41 Sr). As a result, on the sole ground that the employee committed a serious act of violence, the conclusion does not follow that there is a pressing ground for dismissal. Note that this does not mean that the opposite conclusion would follow here, viz. that there is no pressing ground for dismissal. There may be other reasons for the conclusion that there is a pressing ground, such as that the employee insulted a superior.

It can also be supported that a statement supports or attacks a conclusion. An example is in the following figure.



Here the statement is made that the violent act was directed against the employer. This statement supports that having committed a serious act of violence supports that there is a pressing ground for dismissal, in accordance with article 6:678 of the Dutch Civil Code (art. 6:678 BW).

Dialectical arguments will be used throughout this book as a convenient graphical representation of the argumentation in cases. However, the question which conclusions follow is not answered by interpreting dialectical arguments regarding the status of the statements involved (cf. Verheij's dialectical interpretations, 2000). This question is formally dealt with in Sections 3.5 and

3.8, in a way that does not involve any explicit reference to dialectical arguments.

3.1.3 *An analysis of the example*

To arrive at a systematic analysis of the example, for each case the complete argumentation can be represented graphically by a dialectical argument. Then the dialectical arguments of the cases can be compared to see for each settled case whether it is relevant for the problem case.

But which statements must be included in these dialectical arguments? For instance, must the employee's profession be included? That would make all three cases incomparable, since the employee's profession is different in each case. Likewise, if considerable damage was done to company property, does the precise cause of the damage matter? Does one have to include that Al dropped an expensive computer, for instance, or that Bob dropped a pipe? How must choices like these be made?

It is one of the main points of this book that in the law such choices are to some extent contingent, that is, they can differ from one type of case to another (see Section 3.6.2 on the comparison basis). In particular, it can depend on the domain of law at hand which statements are relevant for the purpose of case comparison. It may be, for instance, that for dismissal cases it is relevant whether a warning was given before a person caused damage, while this may not be relevant in the domain of tort.

If reasoning by case comparison is to be a method that can be equally applied to all cases of some type, then a uniform choice must be made here. Presently the employee's profession will not be included in the dialectical arguments for the three dismissal cases, for instance, while the statement that the employee caused considerable damage will.

The result for Al's case is in the following figure.

Note that the main conclusion is followed by a plus sign. This is done to indicate that in Bob's case the conclusion was actually drawn that the dismissal could be voided.

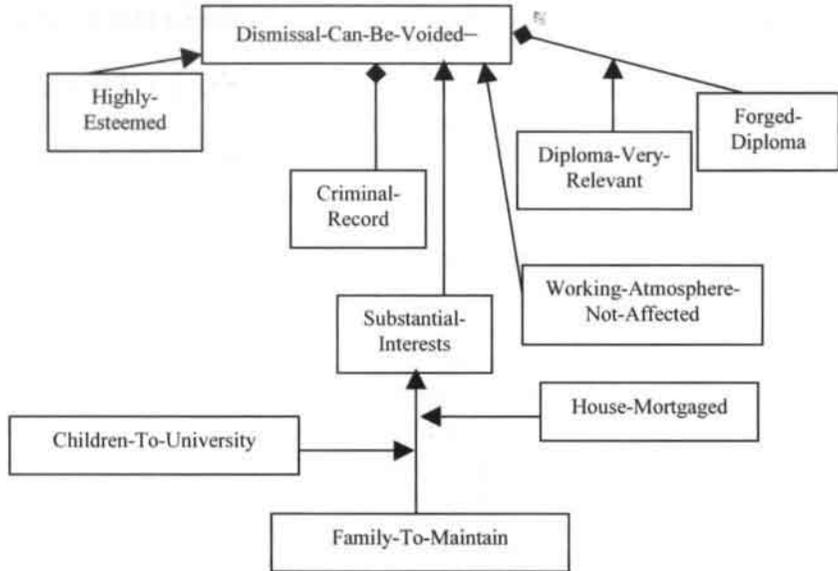
The two dialectical arguments clearly show the relevant differences and similarities between both cases. Among the similarities one has, for instance, that in both cases the employee has a family to maintain. Another similarity is that in both cases a diploma was forged, and that this diploma was very relevant for the job.

A difference is that Bob had a criminal record while this cannot be stated about Al. This attacks the conclusion that the dismissal can be voided in Bob's case, while in Al's case this attack is not possible. Another difference is that in Al's case the employee's house was mortgaged, while this could not be stated in Bob's case. In Al's case this supports that having a family to maintain supports having substantial interests in keeping the job. The statement that the employee has substantial interests in keeping the job supports the conclusion that the dismissal can be voided. As will be made formally precise in Subsection 3.6.5, the statement that Al's house is mortgaged indirectly "contributes to" the conclusion that his dismissal can be voided. No statement of a mortgaged house could be made in Bob's case.

Each of the differences make Al's case provide more support for the conclusion that the dismissal can be voided. If in Bob's case the judge decided that the dismissal could be voided, then apparently, by a kind of argument *a fortiori*, the same should hold for Al's dismissal.

The argumentation in Chris's case can be represented by a dialectical argument as well, as in the following figure.

Chris's dismissal case



Note that the conclusion that the dismissal can be voided is followed by a minus sign, to indicate that the opposite conclusion was drawn.

Just like Bob's case, Chris's case shares some similarities with Al's. For example, in both cases the employee said that he was highly esteemed as a colleague, supporting the conclusion that the dismissal could be voided. Another similarity is that in both cases a diploma was forged, and that this diploma was very relevant for the job.

A difference is that Chris supported the disputed conclusion that the dismissal could be voided by pointing out that the working atmosphere had not been affected. This supports the conclusion that the dismissal can be voided in Chris's case, while it cannot be stated in Al's case. Another difference is that Al's employer did not merely state that having caused considerable damage attacks the disputed conclusion, but he further supported this by pointing out that Al had been given a warning in advance. As will be made formally precise later on (Subsection 3.6.5), this indirectly "detracts from" the conclusion that the dismissal can be voided. A third difference is that Al's employer pointed out that Al's wife had a good income, attacking the statement that having a family to maintain supports the conclusion that the dismissal can be voided. This also indirectly detracts from the conclusion that the dismissal can be voided, while the statement could not be made in Chris's case.

All these differences make that Chris's case provides more support than Al's for the conclusion that the dismissal can be voided. If in Chris's case the judge decided that the dismissal could not be voided, then apparently the same should hold for Al's dismissal.

There is a contradiction between the conclusions to which both decided cases lead. If one follows Bob's case then Al's dismissal can be voided, but it is not if one follows Chris's case. What causes this contradiction? To answer this question one can compare both decided cases with each other. One then finds a number of differences, each of which makes Chris's case provide more support than Bob's for the conclusion that the dismissal can be voided. While in both cases it was stated that having a family to maintain supports having substantial interests in keeping one's job, in Bob's case this statement was attacked by pointing out that his wife had a good income.

However, in spite of all this the judge in Bob's case concluded that the dismissal could be voided, while in Chris's case the opposite conclusion was drawn. In a sense these two decisions contradict each other, and no conclusion at all can therefore be drawn from them.

The foregoing illustrates a number of points. First, it is shown that argumentation involves making statement that support or attack conclusions. An employee can support that he has substantial interests in keeping his job, for instance, by pointing out that he has a family to maintain. Moreover, it can itself be supported or attacked that a statement supports or attacks another statement. By pointing out that the employee's wife has a good income, for instance, an employer can attack that having a family to maintain supports having substantial interests in keeping one's job.

The second point illustrated by the foregoing is that one must decide which statements are relevant for the purpose of case comparison, and which are not. If in one case an expensive computer was dropped and in another a pipe, for instance, then in both cases considerable damage was done. If it is considered relevant what was dropped, then the cases are different. If it is only considered relevant that considerable damage was done, however, then this cannot make the cases differ.

The third point is that one can follow a settled case with a certain conclusion if the case at hand provides at least as much support for the conclusion as the settled case. If it can indeed be shown that there is equal or more support for a conclusion, then by a variant of reasoning a fortiori the conclusion usually holds again. Roughly, this makes all decided cases relevant that provide equal or less support for the disputed conclusion than the case at hand.

The fourth point illustrated by foregoing is that it is possible that precedents lead to contradictory conclusions. In that situation the decided cases cannot be used to derive an unambiguous conclusion.

3.2 Support and attack

In this section some convenient terminology is introduced to deal with the dialectical mechanisms of *support* and *attack*. What can be expressed by a descriptive sentence of some language is called a *state of affairs* (Hage 1997, p.

131-134; Verheij 1996, p. 21, 23). For instance, the English sentence 'Al is a bank employee' expresses the state of affairs that Al is a bank employee. If a state of affairs obtains in a case, then it is called a *fact* in that case.

Dialectical argumentation involves making *statements*, and to this end sentences are used. In this way sentences have two functions. They express states of affairs, and they are used in statements.

The simplest examples of statements involve case facts, such as that Al caused considerable damage. Such statements are called *elementary*, and to contrast them with the more complex ones introduced next they make use of atomic sentences. By convention, variable atomic sentences are small letters written in italics, e.g., *a*, *b* and *c*. The intended meaning of a variable atomic sentence will often be made explicit by means of an abbreviated sentence, in the same way as in HYPO (Ashley 1990), CATO (Aleven 1997) and Prakken and Sartor's (1998) model. An example is *a*: Considerable-Damage, which says that the employee caused considerable damage.

Statements can *support* or *attack* other statements, and it can itself be a statement that such a relation of support or attack holds.⁷ Statements of this kind are called *compound*. An example of a compound statement of support is that being highly esteemed as a colleague supports the statement that the dismissal can be voided. An example of a statement of attack is that having caused considerable damage attacks that the dismissal can be voided. A statement that is supported or attacked by another statement is called a *conclusion*, and the supporting or attacking statement is called a *reason* for or against the conclusion, respectively. An example of a conclusion is that the dismissal can be voided, an example of a reason is that the employee is highly esteemed as a colleague.

It is quite common that there are both reasons for and reasons against a conclusion in one case. Case-based reasoning is a technique to resolve such conflicts. The present model therefore allows for such conflicting reasons. No *ex falso* rule is authorised to the effect that anything would follow from them. Instead a conflict merely blocks the derivation of a conclusion (cf. Hage 1997, pp. 124 and 140), unless the conflict can be resolved by following a precedent. If the employee is highly esteemed as a colleague and has caused considerable damage, for example, then these could be conflicting reasons for and against the conclusion that the dismissal can be voided. In the absence of any precedent that can resolve the conflict between these reasons, no conclusion can be derived from them.

Support and attack are expressed by means of a special connective \nearrow . The conditional sentence $a \nearrow b$, for instance, informally expresses that the statement that *a* supports the statement that *b*, or '*a* supports *b*' for short. Instead of '*a* supports *b*' one can also say 'if *a* then *b*', provided that it is kept in mind that

⁷ The possibility of withdrawing statements is not considered (cf. Verheij 2000c, p. 216).

the support relation is not intended as a standard material implication (see below for more on this; cf. also Verheij 2003, p. 323 on 'primitive implication'). An example of support is that being highly esteemed as a colleague (*a*) supports the statement that the dismissal can be voided (*b*).

To express attack the connective \nearrow is combined with negation, denoted \neg . Thus the sentence $d \nearrow \neg b$, for instance, has the informal reading that '*d* attacks *b*' or 'if *d* then not *b*'. An example is that having a criminal record (*d*) attacks the statement that one's dismissal can be voided (*b*). If a sentence is used in a statement, then by convention the statement using the negation of the sentence is said to be the *opposite* statement. Attacking a statement therefore comes down to supporting its opposite. The opposite of the statement that the dismissal can be voided, for instance, is the statement that the dismissal cannot be voided. Moreover, by supporting that the dismissal cannot be voided one attacks that the dismissal can be voided.

It is a key element of the present model that support and attack are not restricted to statements using atomic sentences in simple small letters, but are also possible for compound statements of support or attack.⁸ To express the support or attack of such statements, a nested notation is used that involves brackets. The sentence $a \nearrow (b \nearrow c)$, for instance, then informally says that *a* supports that *b* supports *c*. An example of this mechanism is that having children planning to go to university (*a*) supports that having a family to maintain (*b*) supports that there are substantial interests in keeping the job (*c*). Likewise the sentence $d \nearrow \neg (b \nearrow c)$, for instance, expresses that *d* attacks that *b* supports *c*. An example of this mechanism is that having a wife with a good income (*d*) attacks that having a family to maintain (*b*) supports that there are substantial interests in keeping the job (*c*).

The connective \nearrow for support can be used for deriving conclusions by a variant of modus ponens (from *p* and if *p* then *q*, conclude *q*). In Sections 3.5 and 3.8 below it will be made formally precise which derivations with this connective are allowed. At this point it suffices to say that unlike the standard material implication, the relation of support allows for non-monotonic modus ponens inferences, in the sense that a conclusion may have to be withdrawn if it is attacked. If *a* and $a \nearrow c$ are both derivable, for instance, then *c* is normally derivable by modus ponens. However, if next to these sentences *b* and $b \nearrow \neg c$ are derivable, then *c* ceases to be derivable by modus ponens.

In the present model such conflicts can be resolved by assigning the same conclusions to them as in comparable cases that have been settled by a judge. In particular, by following settled cases the model can deal with situations where there are both reasons supporting a conclusion, and reasons attacking it. The

⁸ *Attacking* a statement of support or attack is similar to what is often called an 'undercutting exception'. See Hage 1997 p. 166; Verheij 1996, pp. 200-201; Pollock 1995 pp. 41 and 86; Pollock 1987 p. 485. *Supporting* a support or attack statement is comparable to Toulmin's 'backing' of a 'warrant' (1958, pp. 98f.).

model does not resolve such conflicts by explicit weighing knowledge (Hage 1997) or priority information (Prakken and Sartor 1998).

3.3 Representing cases

Cases are (finite) sets of states of affairs that can be expressed by sentences of a language of support and attack, called the *case representation language*. The language is abbreviated as CRL and defined concisely as follows. The convention is adopted that small Greek letters in italics are metavariables for sentences of this language, for example α . A set of atomic sentences (see the previous section) is presupposed, all of which are by definition sentences of CRL. Furthermore, if α and β are sentences, then $(\alpha \wedge \beta)$ and $\neg\alpha$ are sentences too. The formal definition of the case representation language is as follows.

Definition 1 (Case representation language, CRL)

Given a set of atomic sentences, the case representation language, abbreviated CRL, is defined as the smallest set of sentences such that conditions 1. and 2. hold:

1. if α is an atomic sentence, then α is a sentence, and
2. if α and β are sentences, then $(\alpha \wedge \beta)$ and $\neg\alpha$ are sentences.

As usual, outer brackets are omitted (cf. Gamut I, p. 37).

A doubly negated sentence is treated as equivalent to the sentence itself. The sentence $\neg\neg a$, for instance, is equivalent to the sentence a . In accordance with this, the convention is adopted that double negations ($\neg\neg$) vanish everywhere. Thus $\neg\neg c$, for instance, becomes c , and $a \wedge \neg\neg c$ becomes $a \wedge c$. Note that as a result, the opposite of the opposite of a statement is treated as being equal to the statement itself. Accordingly, attacking a conclusion's opposite is the same as supporting the conclusion.

Cases are represented by (finite) sets of sentences of the case representation language CRL.⁹ Such a set of sentences is called a *case representation*. Case representations are denoted in a dedicated way, namely as (strings of) characters in bold type, starting with an uppercase. Examples are **Case** and **Precedent**. To prevent confusion each case is assumed to have one unique denotation, and different cases have different denotations.

Using these denotations one can express states of affairs about cases. Suppose, for instance, that the fact expressed by the sentence a obtains in the case with case representation **Case**. Then this formally means that the sentence a is in the case representation **Case**: $a \in \mathbf{Case}$.

Conclusions of decided cases are treated as statements of facts obtaining in these cases. If conclusion c was drawn in a case with case representation

⁹ Here one encounters the problem of pragmatics, that is, the problem which sentences are to be used to represent cases (cf. Gamut I, 1991, pp. 195f.). This problem is beyond the scope of the present work.

PastCase, for instance, then this formally means that the sentence c is in case representation **PastCase**: $c \in \mathbf{PastCase}$. Note that as a result of this practice, case representations do not distinguish between sentences expressing case facts and sentences used in conclusions. In other words, all sentences in a case representation are treated on a par.

Dedicated variables will stand for arbitrary case representations. By convention these variables are written as strings of characters in bold italics and starting with an uppercase, for instance ***Case***. The expression $a \in \mathbf{Case}$, for instance, then says that the arbitrary case representation ***Case*** has the property that it includes the sentence a .

A final remark is that cases will not always be distinguished strictly from their case representations. Instead of saying that a case with case representation **Precedent** is a precedent, for instance, it will also be said briefly that **Precedent** is a precedent.

3.4 The background

While most that is relevant for case comparison is likely to be included in case representations, it is not at all obvious that cases contain all the relevant information. There are also states of affairs in the world which are general in the sense that they obtain as facts in all cases, and which for that very reason may have been left implicit. In this connection one may think of, for example, facts about the valid rules of statutory law, common-sense rules and rules of conduct. An example could be the common sense rule that having a family to maintain supports the conclusion that there are substantial interests in keeping a job. Besides this, facts generally agreed upon can play a role in case comparison as well. Among these there are, for instance, conventional facts (e.g., Mommers 2002, pp. 66 and 137; Peczenik and Hage 2000). It is a conventional fact, for example, that the Constitution of the United States is authoritative, or that 'A' is the first letter of the Latin alphabet.

Much of this kind of information would take the form of statements of support or attack, which in the present formalism make use of sentences involving the connective \nearrow . Facts generally agreed upon and conventional facts can sometimes be expressed by sentences that do not have a conditional structure, however. An example could be the conventional fact that there are 24 hours in a day.

The existence of relevant information outside case representations means that a background of information has to be postulated. Formally this background is represented by a set of sentences of the case representation language CRL introduced above, and it will be abbreviated BG. The background is comparable to the 'common sense' knowledge in Prakken and Sartor's model (1998, p. 261; see also Section 4.3). Note that there is a difference with what they call 'Background Information', since this also includes cases in addition to common sense knowledge.

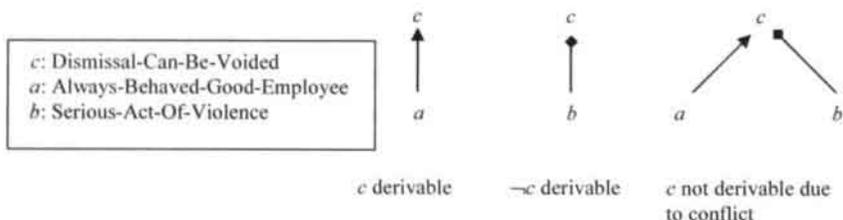
As remarked above, facts given as background are often left implicit in the case representations, while they obtain in all cases and can be relevant in case comparison. The role of the background can therefore be captured by extending the case representations, such that all facts given as background obtain in the extended cases. Formally this is done by adding all sentences of the set BG to each case representation.

3.5 Reasoning with support and attack

This section deals with the question which conclusions are derivable in a case, given the background and using the relations of support and attack between statements. In other words, the derivations discussed here are possible without resorting to other cases as precedents. This makes them a special case of the general consequence relation defined later on (Definition 15 in Section 3.8), in the sense that the latter will also allow for derivations on the basis of a comparison with settled cases.

The notion of derivability presented here is sceptical, in the sense that no conflicts of reasons at all may occur in the derivation of a conclusion. More specifically, a conclusion can only be derivable if there is a dialectical argument for it such that none of the statements involved can be attacked. The idea behind this sceptical choice is that in the present model precedents are the only means by which conflicts of reasons can be resolved.

The premises from which conclusions are derived are the statements of case facts and background. In the following figures a number of example derivations are shown. By convention, the premises in these figures are the statements at which no arrow is pointing. To keep the dialectical arguments relatively compact, the elementary statements in them will use variable atomic sentences instead of rectangles with text (as in Section 3.1). The intended meaning of the variables will be included in a legend. See the following figure.

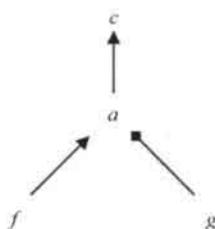


In this figure there are three simple dialectical arguments with at their top the conclusion c . On the left c is supported by a , while it is not attacked. Accordingly, conclusion c is derivable in this situation, and its opposite $\neg c$ is not derivable. In the dialectical argument in the middle, conclusion c is only attacked and not supported. Accordingly, the conclusion $\neg c$ is derivable there, while its opposite c is not.

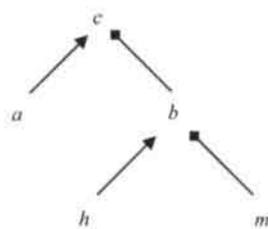
In the third dialectical argument, however, conclusion c is both supported and attacked. In other words, there is a conflict of reasons with regard to c . In the absence of a precedent with which this conflict can be positively or negatively resolved, neither conclusion c nor its opposite $\neg c$ are derivable here. Note that this is in accordance with the convention that attacking a conclusion is the same as supporting the opposite, and that double negations vanish. By virtue of these conventions the opposite conclusion $\neg c$ is supported by b and attacked by a , so that there is a conflict of reasons regarding the opposite conclusion too.

The next figure shows two somewhat more complicated situations.

c : Dismissal-Can-Be-Voided
a : Always-Behaved-Good-Employee
f : Always-Arrived-On-Time
g : Once-Insulted-Superior
b : Serious-Act-Of-Violence
h : Punched-In-Face
m : No-Serious-Injuries



a and c not derivable due to conflict



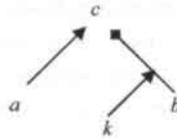
b and c not derivable due to conflict

On the left the intermediate a is supported and attacked. In the absence of a precedent that can resolve the conflict regarding a , the conclusion c is not supported by a derivable reason. As a result, the conclusion c at the top is not derivable here either.

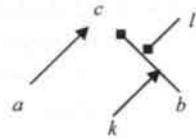
On the right the intermediate b is not derivable because it is both supported and attacked. According to the sceptical notion of derivability adopted here, this means that c is not derivable either. The intuitive reason for this is that the conflict regarding b may be resolved positively with a precedent. The intermediate b would then be derivable and conclusion c would actually be attacked by it. Note that a less sceptical choice in this situation would be to say that b is not derivable and therefore does not attack c . According to this choice, the conclusion c would be reinstated by the attack of b , and the conclusion would accordingly follow.

Conflicts can also occur with regard to statements of support or attack, as in the following figure. On the left it is derivable that b attacks c (expressed by $b \nearrow \neg c$), because this is supported by k .

<p> <i>c</i>: Dismissal-Can-Be-Voided <i>a</i>: Always-Behaved-Good-Employee <i>b</i>: Serious-Act-Of-Violence <i>k</i>: Act-Directed-Against-Superior <i>l</i>: Agitated-Atmosphere </p>



c not derivable due to conflict



b attacks *c*, and *c*, not derivable due to conflict

On the right it is not derivable that *b* attacks *c*, because $b \not\vdash \neg c$ is attacked by *l*. In other words, there is a conflict regarding the conclusion that *b* attacks *c*. According to the present sceptical notion of derivability, this means that *c* is not derivable either. The intuitive reason for this is that a precedent could positively resolve the conflict regarding the conclusion that *b* attacks *c*. As a result, conclusion *c* would then actually be attacked by *b*. Note that a less sceptical choice would be to say that the attack by *b* is blocked due to *l*, and that accordingly the conclusion *c* follows.

More generally, only *conflict-free* derivations are presently possible, because without resorting to precedents no conflicts of reasons can be resolved. Moreover, in the absence of conflict resolving precedents each attack must be taken seriously, even if the attack is itself attacked. From now on conclusions following from such conflict-free derivations will be called *CF-derivable* for short. Informally, a conclusion is CF-derivable if there is a dialectical argument for it starting from the case facts and the background, such that none of the statements involved can be attacked.

More formally, the CF-derivable conclusions in a case with case representation *Case*, are defined as follows. Recall that the background BG can also be involved in the derivation of conclusions, since facts given as background are typically left implicit in case representations. Accordingly, the premises can be represented by the union $Case \cup BG$ of two sets. The first is the case representation *Case*, and the second is the background BG.

As above, metavariables for sentences of this language are small Greek letters in italics, so that γ can stand for a derivable conclusion. The consequence relation itself is written \vdash , with CF (for Conflict-Free) as a subscript. Formally

$$Case \cup BG \vdash_{CF} \gamma$$

then says that γ is CF-derivable from *Case* and BG.

The formal definition of this consequence relation involves the *modus ponens closure* of the set of sentences $Case \cup BG$. Informally, this closure is formed by making all possible modus ponens inferences (meaning that if β and $\beta \not\vdash \alpha$ are in the closure, then α is too). Formally the modus ponens closure is defined as follows for any set of sentences *S*.

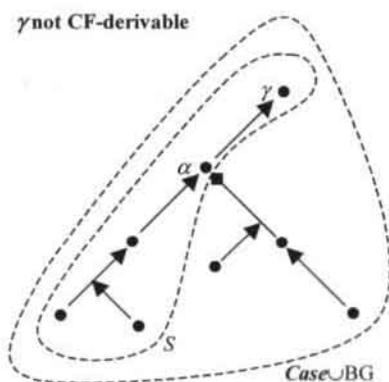
Definition 2 (Modus ponens closure)

Let S be a set of sentences. Then the *modus ponens closure* of S , denoted $\text{Th}(S)$, is the smallest set Th for which 1. and 2. hold:

1. S is a subset of Th , and
2. if β is in Th and $\beta \rightarrow \alpha$ is in Th , then α is in Th .

If a conclusion γ is in the modus ponens closure $\text{Th}(S)$ of some set S , then it will also be said that γ follows by modus ponens from S .

If a conclusion γ is to be CF-derivable from *Case* and BG, then there must be a chain of modus ponens inferences leading to it from the premises $\text{Case} \cup \text{BG}$, that is, case facts and background. Formally this means that there must be a subset S of $\text{Case} \cup \text{BG}$, such that γ follows by modus ponens from S . The following figure is an illustration of this.



The large area enclosed by a dashed line is meant to represent the set of premises $\text{Case} \cup \text{BG}$, and everything that follows by modus ponens from this set. The smaller area enclosed by a dashed line stands for a subset S of $\text{Case} \cup \text{BG}$ from which conclusion γ follows by modus ponens. The chain of modus ponens inferences leading to γ is made visible by a part of the dialectical argument for γ that involves support relations only. If this modus ponens chain were the complete dialectical argument for γ , then it would be CF-derivable.

However, one can attack an intermediate conclusion α in the modus ponens chain leading to γ , because $\neg\alpha$ follows from $\text{Case} \cup \text{BG}$ by modus ponens. In the figure the modus ponens chain leading to $\neg\alpha$ is indicated by another part of the dialectical argument for γ which only involves support relations next to the attack of α . As a result of this attack, at this point the conclusion is not CF-derivable.

Note, however, that if such an intermediate conclusion can be attacked, then it may be possible to find *another* modus ponens chain for conclusion γ which does not involve this intermediate conclusion. If such a chain indeed

exists, then the conclusion to γ may still be CF-derivable (provided that there are no other attacks).

In sum, if a conclusion γ is to be CF-derivable from the case facts and the background that are given as premises, then two conditions must hold. First, there must be a chain of modus ponens inferences leading to γ from the premise set $Case \cup BG$. In other words, there must be a subset S of $Case \cup BG$, such that γ follows by modus ponens from S . The second condition is that none of the statements involved in the derivation can be attacked. Formally this means that if α follows by modus ponens from S , then its opposite $\neg\alpha$ must not follow from $Case \cup BG$ by modus ponens. The formal definition is as follows.

Definition 3 (CF-derivable conclusions)

A conclusion γ is CF-derivable from $Case$ and BG , briefly CF-derivable, if and only if there is a subset

$S \subseteq Case \cup BG$ for which conditions 1. and 2. hold:

1. γ follows by modus ponens from S , and
2. for all α following by modus ponens from S : $\neg\alpha$ does not follow by modus ponens from $Case \cup BG$.

Note that this definition makes the derivation of opposite conclusions like c and $\neg c$ impossible. Note further that in such situations of conflicting premises, the model does not authorise an ex falso rule to the effect that everything would be derivable.

Observe further that even if a conclusion c is among the premises ($c \in Case \cup BG$), it is not derivable if it can be attacked (i.e. $\neg c$ follows by modus ponens from $Case \cup BG$) or if its opposite is given as a premise ($\neg c \in Case \cup BG$).

3.6 Case comparison

In this section it will be explained how case comparison is treated formally in the proposed model. Briefly, the comparison of cases will involve a comparison of the dialectical arguments that represent the argumentation in these cases. In particular, the dialectical arguments can involve statements supporting or attacking that another statement supports or attacks a conclusion.

An important feature of the present account of case comparison is, that it is acknowledged explicitly that in the law it depends on a contingent choice which case features are relevant in comparing cases. This choice of relevant case features is captured formally by what is called a comparison basis, and the contingency of the choice is made explicit by stating the comparison outcomes relative to this comparison basis.

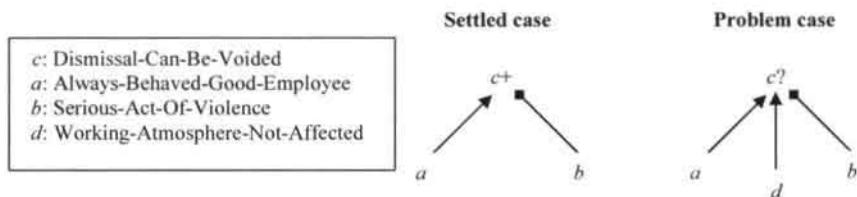
3.6.1 Comparing cases

It is the purpose of case comparison to determine whether a settled case can be followed in a problem case. Intuitively there are two situations in which one can

certainly follow a settled case where a conclusion was drawn. First, it may be that in the problem case there is as much support for the disputed conclusion as in the settled case. Second, it can happen that there is more support for the conclusion in the problem case, so that by a kind of reasoning a fortiori the conclusion should hold again.

In the present model, the support for a conclusion is determined by the dialectical argument with the conclusion at its top. In this connection the present model will speak of *dialectical support* for a conclusion. Case comparison, then, comes down to comparing dialectical arguments regarding the dialectical support for their conclusion.

In the following a number of examples of increasing complexity is given. As before, the elementary statements will use variable atomic sentences, and the intended meaning of the variables will be included in a legend. See the following figure.



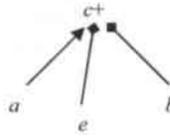
In this figure there is a settled case on the left where the conclusion (*c*) was drawn that a person's dismissal could be voided, as indicated by the plus sign. On the right there is a problem case where this conclusion is an issue, as indicated by the question mark. The plus sign and the question mark represent additional information on the *cases* involved, viz. that some conclusion was drawn or that it is an issue, respectively. They do not stand for evaluations of the dialectical arguments.

In both cases the conclusion (*c*) that the dismissal can be voided is supported by the statement (*a*) that the person has always behaved like a good employee. The conclusion *c* is attacked by the statement (*b*) that the employee committed a serious act of violence. In the problem case the conclusion *c* is also supported by the statement (*d*) that the working atmosphere has not been affected by the dismissal. This supporting statement cannot be made in the settled case. As a result, there is more dialectical support for *c* in the problem case, so that it should follow there as well.

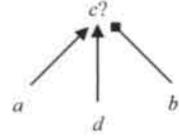
Another example is in the following figure.

c: Dismissal-Can-Be-Voided
a: Always-Behaved-Good-Employee
b: Serious-Act-Of-Violence
d: Working-Atmosphere-Not-Affected
e: Criminal-Record

Settled case



Problem case

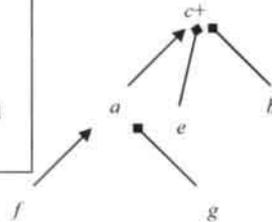


In the settled case the conclusion *c* is attacked by the statement (*e*) that the employee has a criminal record. This attacking statement cannot be made in the problem case. Together with the difference *d* already discussed, this means that there is more dialectical support for *c* in the problem case. As a result, the conclusion (*c*) that the dismissal can be voided can hold there as well.

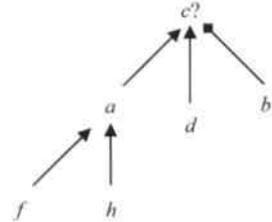
Dialectical arguments can have a more complex structure. An example is in the following figure.

c: Dismissal-Can-Be-Voided
a: Always-Behaved-Good-Employee
f: Always-Arrived-On-Time
g: Once-Insulted-Superior
h: Always-Dressed-Properly
b: Serious-Act-Of-Violence
d: Working-Atmosphere-Not-Affected
e: Criminal-Record

Settled case



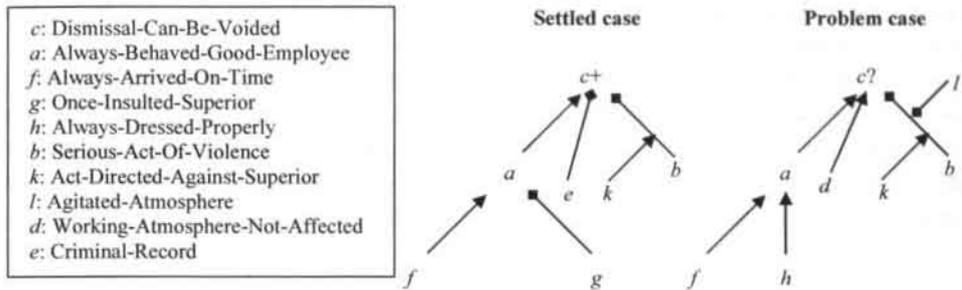
Problem case



In both the settled case and the problem case, the statement (*a*) that the person has always behaved like a good employee is supported by the statement (*f*) that the person always arrived on time. In the problem case the statement (*a*) that the person has always behaved like a good employee is also supported by the statement (*h*) that the employee was always dressed properly. Moreover, in the settled case the statement *a* is attacked by the statement (*g*) that the employee once insulted a superior.

As a result, there is more dialectical support for the statement *a* in the problem case. In accordance with this, one can say that the differences *g* and *h* tend to make the problem case provide more dialectical support than the settled case for the conclusion (*c*) that the dismissal can be voided. Together with the differences *d* and *e* discussed above, this means that there is more dialectical support for conclusion *c* in the problem case than in the settled case. As a result, the conclusion can follow in the problem case as well. Note that for concluding to the outcome that there is more dialectical support for *c* in the problem case, it does not matter how the conflict with regard to the intermediate *a* is to be resolved. In other words, for comparing the cases regarding their dialectical support for *c*, it makes no difference whether or not the intermediate conclusion *a* was actually drawn in the settled case.

It can itself be supported or attacked that one statement supports or attacks another. An example is in the following figure.



The difference with the previous situation is that in the problem case it is subject of argumentation whether the conclusion (*c*) that the dismissal can be voided is attacked by the statement (*b*) that the employee committed a serious act of violence. That *b* attacks *c* is supported in both cases by the statement (*k*) that the violent act was directed against a superior. Moreover, the attack by *b* is attacked in the problem case by the statement (*l*) that the violent act took place in an agitated atmosphere. The statement of an agitated atmosphere cannot be made in the settled case. As a consequence, in the problem case there is less dialectical support for the statement that *b* attacks *c*. In accordance with this, one can say that the difference *l* tends to make the problem case provide more dialectical support than the settled case for the conclusion (*c*) that the dismissal can be voided. Together with the other differences discussed before, this means that the problem case provides more dialectical support for conclusion *c* than the settled case. As a result, the conclusion can follow in the problem case as well. Note that for concluding to this result it is not necessary to resolve the conflict with regard to the attack by *b*. In other words, for the purpose of case comparison one does not have to decide whether *b* actually attacks *c* in the problem case.

These examples illustrate that cases can be compared with respect to their dialectical support for the disputed conclusion, and that this comparison involves the dialectical arguments with the conclusion at their top. In other words, only that part of cases is compared that counts for the dialectical support for the disputed conclusion, while other case features can be left out of consideration. This makes case comparison relative to the conclusion for which a settled case is cited. In this respect the present model is comparable to Prakken and Sartor's (1998), who speak in this connection of the 'citation of portions of precedents' (p. 268, original emphasis) and 'local' comparisons (p. 269; see Section 4.3). Cf. also Branting's precedent constituents (2000, p. 66; see Section 4.5).

There are two situations in which the comparison of dialectical arguments has the result that a settled case can be followed in a problem case. The first

situation is that the dialectical argument for the conclusion is the same in the problem case as in the settled case. In that situation the dialectical argument of the problem case involves the same statements as that of the settled case. This holds not only for the elementary statements involved, but also for the compound statements of support or attack.

The second situation where a settled case can be followed is that there are differences in the dialectical arguments, but the differences are such that the problem case provides more dialectical support for the disputed conclusion than the settled case. By a kind of reasoning *a fortiori*, in that situation the conclusion should hold in the problem case.

3.6.2 *The comparison basis*

When comparing dialectical arguments one finds different types of statement. The conclusion at the top of these dialectical arguments tends to be the statement of an abstract legal state of affairs, such as that there is liability for damage. At the bottom of the arguments one normally finds statements of non-legal and concrete facts, such as that John smashed a window by throwing a brick.

Not all these states of affairs are relevant for the purpose of case comparison. The mere fact that John smashed a window, for example, can be irrelevant for the purpose of comparing his case with another case where liability for damage is an issue. At this point a distinction is therefore postulated between relevant and irrelevant states of affairs. From this point onwards all statements corresponding to relevant states of affairs will be called *factors*, to stress their role in case comparison.¹⁰ Accordingly, statements corresponding to irrelevant states of affairs will be called *non-factors*. It can be a non-factor, for instance, that Al dropped an expensive computer, but a factor may be the statement that the employee caused considerable damage.

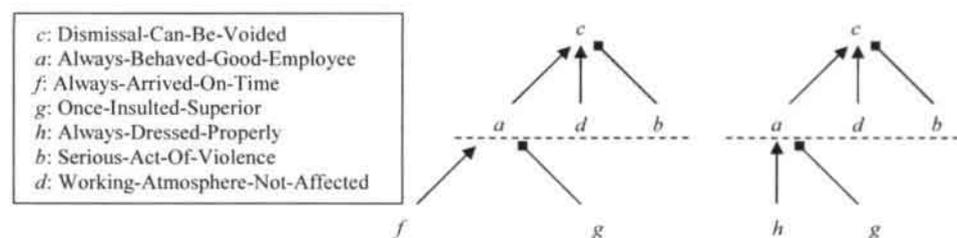
In general it is not a fixed matter but disputable which statements are factors and which are not. In other words, in the law it depends on a contingent choice which factors are taken into account when comparing cases. For a start this choice will depend, of course, on the legal domain under consideration, such as trade secret law or dismissal law. However, even within one domain one can in principle choose different factors for the purpose of case comparison. A possible criterion to guide the choice could be that factors correspond to legal states of affairs, while other statements do not. As discussed further on, however, there are situations in which this criterion does not work.

Some factors have the intuitive property that what supports them can be ignored for the purpose of case comparison, and that for this purpose all that

¹⁰ Cf. the 'factors' in HYPO and CATO, which are generalised, legally relevant sets of case facts (cf. Ashley 1990, pp. 37-38; Alevén 1997, p. 19). A difference is that it is presently acknowledged explicitly that the choice of factors is contingent, that is, it is disputable which statements will count as factors.

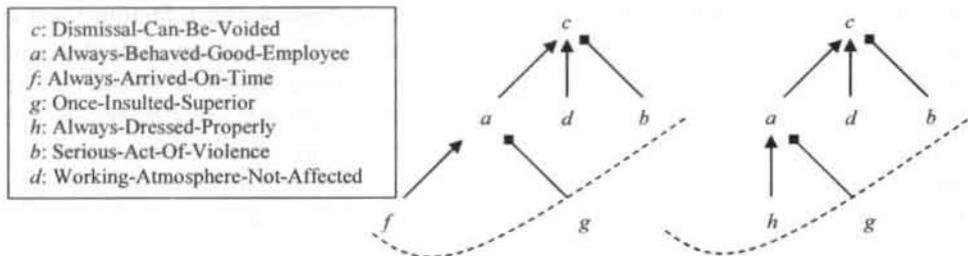
matters is whether or not they can be derived. By definition these factors form a set called the *comparison basis*, and the factors in this set are called *basic* relative to the comparison basis.¹¹ Suppose, for instance, that it is a basic factor (*a*) that a person has always behaved like a good employee, and suppose that this is supported by the statement (*f*) that the person has always arrived on time. Then for comparing cases the factor (*a*) is relevant that the person has always behaved like a good employee, while the supporting statement (*f*) can be ignored that the person has always arrived on time.

In the figure below the comparison basis is visualised by drawing a line of division through the dialectical arguments. Statements above the line correspond to relevant states of affairs that have to be considered for the purpose of case comparison. States of affairs below the line are irrelevant and can be ignored in comparing cases. If the comparison basis is chosen as in the figure, it can be ignored that on the left the factor (*a*) that the person has always behaved like a good employee is supported by a different statement than on the right. On the left the supporting statement is (*f*) that the employee has always arrived on time. On the right the supporting statement is (*h*) that the employee was always dressed properly. As said, this difference can be ignored here, so the two cases are comparable regarding their dialectical support for the conclusion (*c*) that the dismissal can be voided.



The following figure shows a different choice for the comparison basis in the same situation. According to the alternative choice the supporting factors *f* and *h* are relevant as well. In other words, it makes a difference whether the employee was always dressed properly, or always arrived on time. As a result, the two cases are not comparable regarding the dialectical support for the conclusion (*c*) that the dismissal can be voided.

¹¹ Cf. The 'base-level factors' in CATO (Aleven 1997 pp. 23 and 47).



These figures illustrate that the comparison outcomes depend on the particular division made between factors and non-factors. In other words, the conclusions derived by case comparison depend on the comparison basis. Moreover, the comparison basis is contingent, in the sense that it is disputable which statements will count as basic factors. It can depend on the issue at stake, for instance, which factors are basic and which are not. Suppose, for example, that the issue is whether the dismissal can be voided, and that a basic factor relative to this issue is that the employee has caused considerable damage to company property. This means that it is not relevant whether the damage was caused by dropping an expensive computer or by cracking up a car. If the issue is whether the employee is liable to compensate for the damage, however, then the exact cause of the damage may be very relevant. Accordingly, it is then not a basic factor that the employee caused considerable damage.

Once a particular comparison basis has been chosen, it determines which conclusions can be derived. If case comparison is to be a method that can be equally applied to all cases, however, this gives rise to a problem. The problem is that different comparison bases may be used, resulting in different comparison outcomes from which different conclusions can be derived. From a modelling perspective this is problematic, in particular for defining the conclusions that follow by case comparison.

The present model deals with this problem in two ways. First, the comparison basis is assumed to be fixed in advance and to be uniform, in the sense that it is the same for all comparisons made. Second, the outcomes of case comparison are stated relative to this fixed comparison basis, while the same is done with the consequence relation.

Formally the comparison basis is represented by a set of sentences of the case representation language CRL defined above, and it is referred to by the abbreviation CB. For example, let the sentence *a* abbreviate the sentence 'The employee caused considerable damage due to carelessness'. Then the sentence *a* may be used in a basic factor. Next to factors that take the form of elementary statements, the comparison basis can contain statements for support or attack. For example, it can be a basic factor (*a* / *c*) that having substantial interests in keeping one's job (*a*) supports the conclusion (*c*) that the dismissal can be voided. At the same time it need not be a basic factor, however, that (*a*) one has substantial interests in keeping one's job.

It is also possible that a statement (say f) is a factor while its opposite ($\neg f$) is not. This can happen if for case comparison it is only relevant whether the statement is involved in a case, or not (see Subsection 3.6.5 for the factor statements involved). It is then not necessarily also relevant whether or not the opposite statement can be made. As an example, consider the statement (f) that the employee committed a serious act of violence (art. 7:678 of the Dutch Civil Code), which is relevant for the issue whether a dismissal can be voided. Accordingly, the statement that a serious act of violence was committed will usually be treated as a factor that is relevant in comparing dismissal cases. Cases in which no serious act of violence was committed can then be dealt with, however, as ones that do not involve the statement (f) that such an act was committed. In other words, to deal with dismissal cases one does not necessarily have to treat the opposite statement ($\neg f$) as a relevant factor.

Although the comparison basis is contingent, at first sight there appear to be two distinctions that can guide the choice of the basic factors. First, it seems that concrete states of affairs are usually not involved in case comparison, so that they have to be restated at some higher level of abstraction. The fact that Al dropped an expensive computer, for instance, can be irrelevant in comparing his dismissal case with other cases, while it may very well be relevant that considerable damage was done to company property.

The second distinction is supported by the observation that only legal states of affairs normally seem to be involved in case comparison, so that case facts first need to be legally qualified. If an employee has insulted his superior, then this may not play a role in comparing cases. However, if the offence gives rise to a pressing ground for dismissal (art. 7:677 Dutch Civil Code), then the presence of such a pressing ground may be very relevant. In sum, only abstract and legal states of affairs seem to matter in case comparison.

Neither of these distinctions is a fully satisfactory guide in choosing a comparison basis, however, because there are situations where highly concrete and non-legal states of affairs can become relevant in case comparison. As an example, suppose that an employee is dismissed and that the issue is whether the dismissal can be voided. Suppose, moreover, that the employee's case is compared with another dismissal case. Then it may very well be relevant that the employee has a family to maintain, or that his children are planning to go to university. Such statements are very concrete and non-legal, but can still be relevant as factors in case comparison.

3.6.3 *The entangled factor hierarchy*

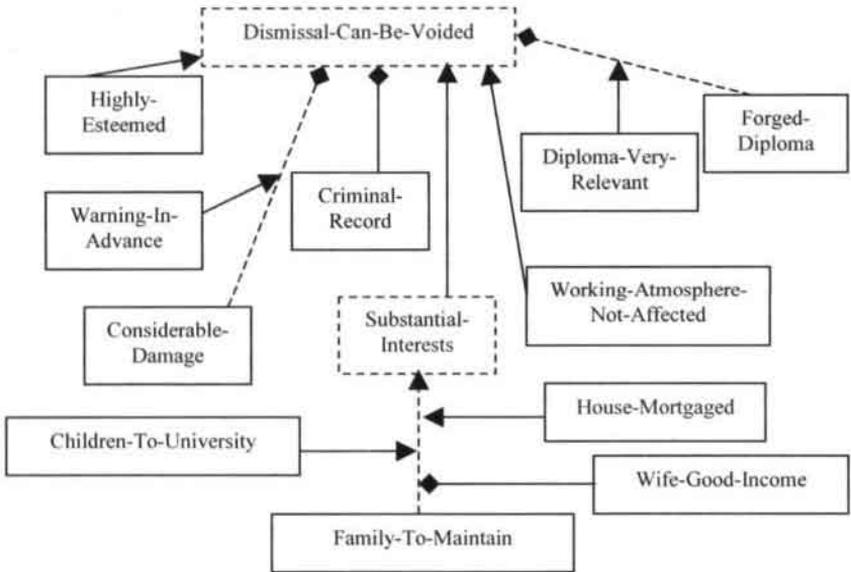
In this section all factors that follow by modus ponens from the whole comparison basis are considered. Together these factors make up a representation of the argumentation pertaining to some domain of law. An example of such a representation is depicted in the figure below.

To distinguish between basic and other factors, the following conventions are used in this figure. The elementary basic factors are placed within a solid

box, while other elementary statements in dashed boxes. An elementary basic factor is the statement, for instance, that the working atmosphere has not been affected by the dismissal. An example of an elementary non-basic factor is that the employee has substantial interests in keeping his job.

Basic compound support or attack statements are indicated by a solid arrow, while non-basic support or attack statements are indicated with dashed arrows. For instance, a basic compound statement is that being highly esteemed as a colleague supports that one's dismissal can be voided. An example of a non-basic compound factor is the statement that having a family to maintain supports having substantial interests in keeping one's job.

Entangled factor hierarchy (part)



The representation thus obtained is similar to the so-called 'Factor Hierarchy' employed by CATO (Aleven 1997, pp. 44-45) as a case-independent knowledge structure. In both structures there is a notion of support and attack of conclusions, and in both the support relations are represented by arrows. However, in contrast with CATO the present model also allows for supporting and attacking statements of support or attack. In the representation this corresponds to arrows that point at other arrows. This makes the representation 'entangled' in comparison to CATO's Factor Hierarchy. For these reasons the present structure is called an *entangled factor hierarchy* (cf. Roth 2001b, pp. 31-33).

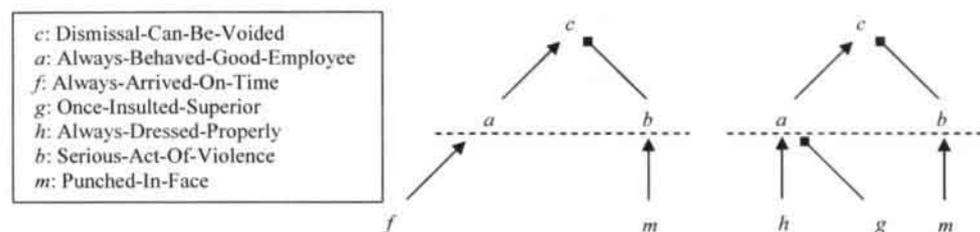
The example in the figure above merely represents an illustration of what the entangled factor hierarchy may look like for the domain of Dutch dismissal law. Although no claim is made that it is a realistic representation of actual legal

practice in the domain, most of the factors were found in reports of decided dismissal cases.

3.6.4 Deriving factors from non-factors

Cases are compared in terms of factors only, because non-factors are by definition not relevant to this end. The question which factors can be involved in case comparison is addressed in the next subsection. In this subsection this question is answered for basic factors.

It is a special feature of the present model that the basic factors involved in case comparison do not have to be given in advance, but can be derived from the non-factors of the case representation and the background. Briefly, the basic factors that can be involved are those that are CF-derivable according to Definition 3 (Section 3.5). See the following figure.



As in Section 3.6.2, the division between factors and non-factors is indicated by a dashed line. The statement (*a*) that the dismissed person has always behaved like a good employee is a basic factor. Another basic factor is the statement (*b*) that the employee committed a serious act of violence. In both cases the latter basic factor is supported by the statement (*m*) that the employee punched a person in the face. As a result, basic factor *b* is CF-derivable in both cases (see Definition 3, Section 3.5).

On the left the statement (*f*) that the person always arrived on time supports the basic factor (*a*) that he has always behaved like a good employee. On the right this basic factor is supported by the statement (*h*) that the employee was always dressed properly. However, on the right it is also attacked that (*a*) the person has always behaved like a good employee, namely by the statement (*g*) that the employee once insulted a superior. As a result, on the right the basic factor *a* is not CF-derivable.

For case comparison one must first determine the basic factors that can be involved. As said above this is done by taking the basic factors that are derivable in a conflict-free way in each case, in accordance with the definition of CF-derivability above (Definition 3, Section 3.5). Recall from that definition that briefly stated, conclusions are CF-derivable as long as this does not lead to conflicts due to attack.

The CF-derivable basic factors are said to *apply* to a case (relative to CB and BG).¹² It is convenient to have a shorthand notation for the set of applying basic factors, and to this end the abbreviation ApplBF is used, with CB and BG added as subscripts. The formal definition is as follows.

Definition 4 (Applying basic factors in a case)

The set of *applying basic factors in Case* relative to CB and BG, denoted $\text{ApplBF}_{\text{CB, BG}}(\text{Case})$, is defined as the set of basic factors α relative to CB, such that α is CF-derivable from *Case* and BG.

According to this definition, the applying basic factors are derived from non-factors in a conflict-free way. In principle, however, other definitions could be given too. For instance, one could consider all basic factors that are *at all* supported in the case, thus allowing for possible conflicts among the basic factors that apply. If it is both supported and attacked that (a) the dismissed person has always behaved like a good employee, for instance, then one could treat this statement and its opposite both as an applying basic factor.

As said the choice is presently made that the applying basic factors are derived in a conflict-free way, however. There is a combination of two intuitive reasons for this choice. First, within the present model the only means of resolving conflicts is by case comparison. Second, by definition case comparison is done in terms of factors, so that it can only deal with conflicts among factors. As a result, within the present model there is no way of resolving conflicts among non-factors, and this is why the applying basic factors are derived in a conflict-free way.

3.6.5 Which factor statements are relevant?

This section deals with the question which factors are relevant when cases are compared regarding their dialectical support for a conclusion. To this end the following points are formally addressed. First, which factor statements can be made in a case, given the applying basic factors? Second, which of these statements are involved in the dialectical argument for a conclusion, and how do they affect the dialectical support for the conclusion? As will be seen, the formal treatment of these points will lead to a notion of factors' relevance for case comparison that is context-dependent, that is, dependent on the cases involved.

The question which factor statements can be made is captured formally by the notion of *grounded factors*. These are defined as those statements that follow by modus ponens from the set of applying basic factors. An example of a grounded factor may be the statement (k) that the employee has substantial

¹² Cf. the *applicable* factors in HYPO, which are inferred if their *prerequisites* are satisfied. See Ashley 1990, pp. 37-38.

interests in keeping his job, assuming that it is supported by some basic factor (l) that the employee has a family to maintain.

It is possible – and common – that a factor and its opposite are both grounded in a case, namely if the factor is supported and attacked. It may be, for instance, that the factor (a) that a person has always behaved like a good employee, is supported by the basic factor (f) that the person has always arrived on time. At the same time the factor (a) concerning the person's behaviour may be attacked by the basic factor (g) that the employee once insulted a superior.

The definition of grounded factors involves the set $\text{ApplBF}_{\text{CB, BG}}(\text{Case})$ of basic factors that apply to *Case* (relative to CB and BG). The grounded factors are captured formally by taking the modus ponens closure of this set of applying basic factors, denoted $\text{Th}(\text{ApplBF}_{\text{CB, BG}}(\text{Case}))$. This is the smallest set containing all applying basic factors, such that if β and $\beta \rightarrow \alpha$ are in the closure, then α is too (cf. Definition 2). By definition, factors using sentences of the set thus formed are called *grounded* in *Case* relative to CB and BG. Furthermore, instead of saying that γ is a grounded factor it will be said for short that γ is grounded. The formal definition is as follows.

Definition 5 (Factors grounded in a case)

γ is a *grounded factor in Case* relative to CB and BG (is *grounded in Case* relative to CB and BG) if and only if $\gamma \in \text{Th}(\text{ApplBF}_{\text{CB, BG}}(\text{Case}))$.

Intuitively the grounded factors are those statements that can be made in a case, given the applying basic factors. If a certain conclusion is at stake in case comparison, then the next question is which of the grounded statements are involved in the dialectical argument for the conclusion.

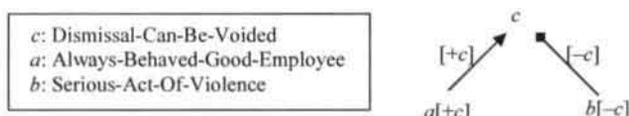
A key intuitive property of factors involved in a dialectical argument is, that they affect the dialectical support for the argument's conclusion. They can have two effects. First, they can "contribute to" the conclusion, and such factors will be treated formally as *relevant for* the conclusion. Second, they can "detract from" the conclusion, and these factors will be treated formally as relevant for the *opposite* conclusion. The simplest example of a factor that is relevant for a conclusion is a reason for the conclusion, and the simplest example of a factor that is relevant for a conclusion's opposite is a reason against the conclusion. As will be seen, however, in the present model there are many more ways of contributing to or detracting from a conclusion.

One could choose to treat as relevant only the directly supporting or attacking factors that are derivable, and for which it is derivable that they support or attack the conclusion (cf. Hage 1997, pp. 185-187). The present approach will be more general, however, in the sense that one can also treat other factors as relevant. In particular, case comparison can presently also involve intermediate conclusions, or factors that support or attack that another factor supports or attacks a conclusion.

In the following it is illustrated with examples how the relevance relation will be formally defined. Each of the statements in the examples is assumed to be grounded, that is, it follows by modus ponens from the applying basic factors.

A trivial situation in which a grounded factor is involved in a dialectical argument is, when it appears as the argument's conclusion. Apart from this trivial situation, the simplest way in which a factor can be involved in a dialectical argument is by being a reason for or against the argument's conclusion. More precisely, if a factor is grounded in combination with the statement that it supports or attacks a conclusion, then both are involved in the dialectical argument for this conclusion. A reason for a conclusion is relevant for the conclusion, and the same holds for the statement that the reason supports the conclusion. Likewise, a reason against a conclusion is relevant for the conclusion's opposite, and the same holds for the corresponding statement of attack. See the example of the following figure.

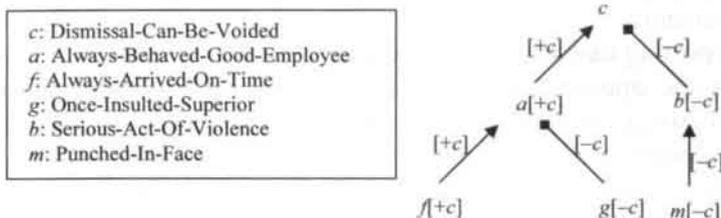
Relevance for c and $\neg c$



In this situation one has the factor (a) that a dismissed person has always behaved like a good employee, which supports the conclusion (c) that the dismissal can be voided. The conclusion is attacked by the factor (b) that the employee committed a serious act of violence. In the figure the statements relevant for conclusion c are indicated with a plus sign followed by the conclusion, placed between square brackets (i.e. $[+c]$). The statements relevant for $\neg c$ are indicated in a similar way with a minus sign (i.e. $[-c]$). As is thus shown, both a and a/c are relevant for c , and b and $b/\neg c$ are relevant for $\neg c$. This illustrates that relevance is not restricted to elementary factor statements using atomic sentences, but is also attributed to compound statements of support and attack.

Statements can also be involved in the dialectical argument in a less direct way than by being reasons for or against a conclusion, for they can also support or attack such reasons. A statement that supports a reason for a conclusion is relevant for the conclusion, and a statement that supports a reason against the conclusion is relevant for the opposite conclusion. Similarly, a statement that attacks a reason for a conclusion is also relevant for the conclusion's opposite. An example is in the following figure.

Relevance for c and $\neg c$

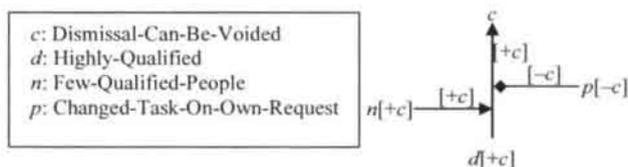


In the figure one finds the statement (f) that the employee has always arrived on time, which supports the statement (a) that the dismissed person has always behaved like a good employee. The latter statement is attacked by the statement (g) that the person once insulted a superior. In addition to this, the statement (b) that the employee committed a serious act of violence, is supported by the statement (m) that the employee punched a person in the face.

In this situation f and $f \wedge a$ are relevant for a and g and $g \wedge \neg a$ are relevant for $\neg a$. As indicated in the figure, one has in addition that both f and $f \wedge a$ are relevant for the main conclusion c , and that both g and $g \wedge \neg a$ are relevant for $\neg c$. Likewise, both m and $m \wedge b$ are relevant for b and for $\neg c$, as indicated in the figure. This example illustrates that relevance can be “passed on” by intermediate conclusions.

Other factors involved in a dialectical argument are factors supporting or attacking that a statement is a reason for or against a conclusion. By supporting that a factor is a reason for a conclusion, for instance, a statement is relevant for the conclusion too. Similarly, if a statement attacks that a factor is a reason for a conclusion, then it is relevant for the opposite conclusion. An example is in the following figure.

Relevance for c and $\neg c$



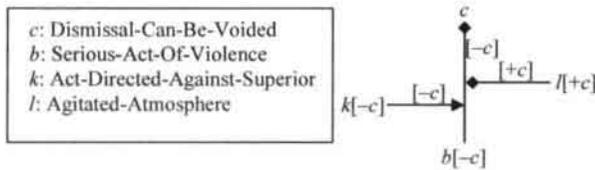
Here one has the factor (n) that there are only few people qualified for the job, which supports that being highly qualified for the job (d) supports the conclusion (c) that the dismissal can be voided. In addition one finds the factor (p) that the employee was given a new task on his own request, which attacks that being highly qualified (d) supports that the dismissal can be voided (c).

In this example n and $n \wedge (d \wedge c)$ are relevant for $d \wedge c$, and p and $p \wedge \neg (d \wedge c)$ are relevant for $\neg (d \wedge c)$. As indicated in the figure, one has in addition that both

n and $n \neq (d \neq c)$ are relevant for c , and that both p and $p \neq \neg(d \neq c)$ are relevant for $\neg c$. This example illustrates that relevance can also be passed on by a relevant support statement.

By supporting that a factor is a reason against a conclusion, a statement is relevant for the opposite conclusion. Similarly, if a statement attacks that a factor is a reason against a conclusion, then it is relevant for the opposite of the conclusion's opposite. By virtue of the convention that double negations vanish, equivalently it is then relevant for the conclusion itself. See the example in the following figure.

Relevance for c and $\neg c$



Here one has the factor (k) that a serious act of violence was directed against a superior, which supports that having committed the act (b) attacks the conclusion (c) that the dismissal can be voided. In addition one finds the factor (l) that the serious act of violence was committed in an agitated atmosphere, which attacks that having committed the act (b) attacks that the dismissal can be voided (c).

In this situation k and $k \neq (b \neq \neg c)$ are relevant for $b \neq \neg c$, and l and $l \neq \neg(b \neq \neg c)$ are relevant for $\neg(b \neq \neg c)$. As indicated in the figure, one also has that both k and $k \neq (b \neq \neg c)$ are relevant for $\neg c$, and that both l and $l \neq \neg(b \neq \neg c)$ are relevant for c . This example shows that relevance can also be passed on by a relevant statement of attack.

The examples illustrate the relevance relation between factors. The formal definition of this relation is as follows.

Definition 6 (Relevance relation)

The *relation of relevance in Case* relative to CB and BG is the smallest relation such that 1., 2. and 3. hold:

1. if α is grounded in *Case* relative to CB and BG, and $\alpha' \beta$ is grounded in *Case* relative to CB and BG, then α and $\alpha' \beta$ are *relevant for β in Case* relative to CB and BG;

and

2. if α is relevant for β in *Case* relative to CB and BG, and β is relevant for γ in *Case* relative to CB and BG, then α is *relevant for γ in Case* relative to CB and BG;

and

3. if α is relevant for $\neg\beta$ in *Case* relative to CB and BG, and β is relevant for γ in *Case* relative to CB and BG, then α is *relevant for $\neg\gamma$ in Case* relative to CB and BG.

As was done in the examples above, abbreviations will be used in dialectical arguments to indicate that a factor is relevant for a conclusion or its opposite. In this way [+c] then abbreviates, for instance, that a factor is relevant for conclusion c . Similarly, [-c] says that a factor is relevant for $\neg c$.

Clause 1 of this definition informally says that if a factor statement can be made in combination with the statement that it supports some conclusion, then both are relevant for the conclusion. This makes attacking statements relevant for the opposite conclusion, since attack is defined as support for a conclusion's opposite.

Clause 2 informally says that if a statement is relevant for a second and the second is relevant for a third, then the first is also relevant for the third (transitivity). Intuitively, this clause holds that one can contribute to a conclusion by contributing to a factor that contributes to the conclusion.

Clause 3 informally says that if a statement is relevant for the opposite of a second statement, and the second is relevant for a third, then the first is relevant for the opposite of the third. Intuitively, this clause holds that one can detract from a conclusion by detracting from a factor that contributes to the conclusion.

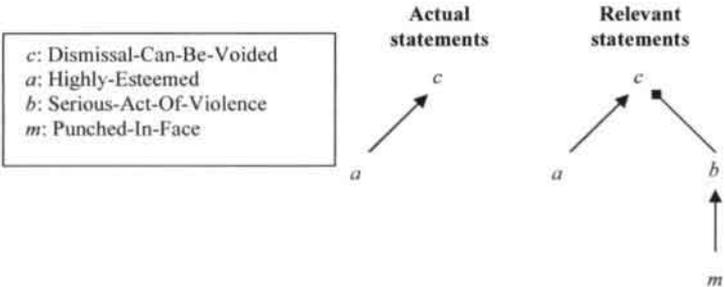
Note that the relation of relevance is defined as the smallest relation satisfying these three clauses, to make sure that no relevances can appear out of the blue. As a result, all relevant factor statements are by definition also grounded in the case under consideration, which intuitively means that they can all be made in the case, given the applying basic factors.

Note further that relevance is a context-sensitive notion, since it is defined as a relation holding within cases. More specifically, it can depend on other factors whether or not a factor must be taken into account when comparing a case to other cases. This context-dependence of relevance is in correspondence with legal practice. As an example, suppose that having caused considerable damage attacks that one's dismissal can be voided, provided that a warning was given in advance. Suppose, moreover, that considerable damage was done in

some dismissal case. Then if a warning was given in advance, the statement that considerable damage was done is relevant for the conclusion that the dismissal cannot be voided. If no warning was given in advance, however, then the statement of considerable damage is not relevant.

For modelling case comparison, precisely those factors are taken into account that count as relevant according to the formal definition above. However, at this point it must be remarked that with regard to the factor statements made, the model can differ from the actual argumentation by parties in real cases. Two types of difference can occur here. First, in real cases not all factor statements are always made that are relevant according to the definition just given (cf. Loui and Norman 1995, p. 164 on arguments not advanced). Second, in real cases statements are sometimes made that are not relevant, for instance because they are not grounded.

An example of the first situation is in the following figure. The basic factors are those at which no support or attack arrow is pointing.

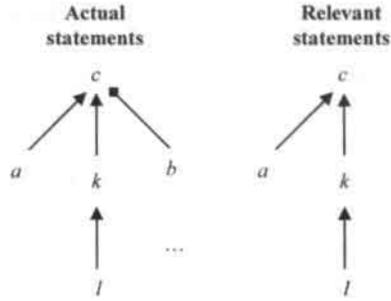


In this figure the dialectical argument on the left represents the actual situation, while on the right the version is represented that contains all relevant factor statements. In this version the conclusion (*c*) that the dismissal can be voided is attacked by the factor (*b*) that the employee committed a serious act of violence. This factor is in turn supported by the basic factor (*m*) that the employee punched a person in the face.

The figure also shows, though, that in actual argumentation this attack does not take place. An intuitive reason for this could be that the opportunity of attacking is not seized by a party, for instance because this attack is believed to be unsuccessful anyway.

The second type of difference between actual and relevant argumentation is that a statement is made that is not relevant, for instance because it is not grounded. An example is in the following figure.

<p>c: Dismissal-Can-Be-Voided a: Highly-Esteemed k: Substantial-Interests l: Family-To-Maintain b: Serious-Act-Of-Violence</p>
--



As the figure shows, the conclusion (*c*) that the dismissal can be voided is attacked by the factor (*b*) that the employee committed a serious act of violence. The figure also shows, though, that this attacking factor is not relevant because in the case at hand it is not grounded. An intuitive reason why the attacking factor statement (*b*) is still made could be, that it is believed to be grounded in the case. Another intuitive reason could be that the attacking factor (*b*) is known to be not grounded, but that a party speculates on its acceptance by others.

3.6.6 Comparison outcomes

In this section the outcomes of case comparison are formally defined. One outcome is that two cases provide equal dialectical support for a conclusion, another is that one case provides at least as much dialectical support for a conclusion as another case. A third outcome is that in one case there is strictly more dialectical support for a conclusion than in another.

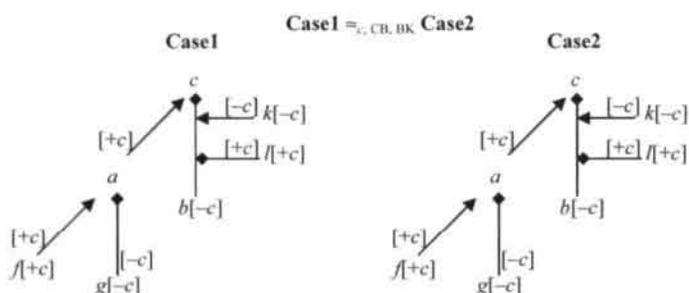
A convenient set of sentence types is introduced first to express these comparison outcomes. For each comparison outcome there is a special type of sentence, as in the following list:

1. **Case1** $\approx_{\gamma, CB, BG}$ **Case2**
This expresses that **Case1** and **Case2** provide *equal dialectical support* for conclusion γ relative to comparison basis CB and background BG.
2. **Case1** $\geq_{\gamma, CB, BG}$ **Case2**
This expresses that **Case1** provides *at least as much dialectical support* for conclusion γ as **Case2** relative to comparison basis CB and background BG.
3. **Case1** $>_{\gamma, CB, BG}$ **Case2**
This expresses that **Case1** provides *strictly more dialectical support* for conclusion γ than **Case2** relative to comparison basis CB and background BG.

Comparison outcomes concern the dialectical support for a conclusion, and they relate this dialectical support in one case to that in another case. The dialectical support for a conclusion in a case is determined by the dialectical argument that can be constructed for the conclusion. As discussed earlier (Subsection 3.6.5), dialectical arguments used in case comparison involve grounded statements only, that is, statements that follow by modus ponens from the applying basic factors. In other words, the dialectical arguments involved in case comparison are constructed from the applying basic factors and by using modus ponens. This means that the dialectical support for a dialectical argument's conclusion is in the end determined by the applying basic factors with which the dialectical argument is constructed.

The idea behind the present formalisation of case comparison is, that for each case and for each conclusion the basic factors can be divided into two sets, namely those relevant for the conclusion and those relevant for the opposite conclusion. (In exceptional situations they can be both, for instance if one and the same basic factor is both a reason for and against the conclusion.) The sets of relevant basic factors are then involved in the definition of comparison outcomes.

One outcome of case comparison is that two cases provide equal dialectical support for a conclusion. This happens precisely then when the dialectical arguments for the conclusion are identical.¹³ This dialectical argument is constructed by applying modus ponens to the basic factors that are relevant for the conclusion or its opposite. One can therefore also say that there is equal dialectical support for a conclusion in two cases, if and only if in the first case the same basic factors are relevant for the conclusion as in the second, while in addition the same basic factors are relevant for the conclusion's opposite. A formal example is in the following figure.

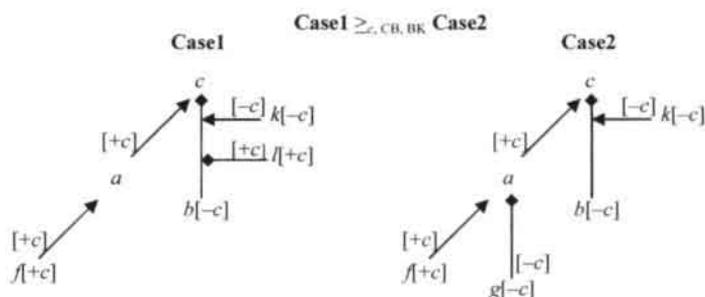


¹³ The requirement that the dialectical arguments are identical could be relaxed by taking into account that sentences can be equivalent, for instance the sentences $p \wedge q$ and $q \wedge p$. This could be done, for instance, by treating the sentences of CRL as equivalence classes of sentences of a richer language. This does not provide more insight, however.

In this figure two cases are compared regarding the dialectical support for conclusion c . The relevant basic factors are those at which no support or attack arrow is pointing, and this holds for both elementary (e.g., f) and compound ($f \wedge a$) basic factors. Basic factors relevant for the conclusion are indicated by $[+c]$, and basic factors relevant for $\neg c$ are indicated with $[-c]$.

In the situation of the figure the dialectical arguments are identical, so that there is equal dialectical support for c in both cases. Accordingly, in both cases the same basic factors are relevant for the conclusion, and the same are relevant for the opposite conclusion. Among the basic factors that are relevant for c one has, for instance, the factors f and $a \wedge c$. Among the basic factors relevant for $\neg c$ one has, for instance, the factors g and b .

Another outcome of case comparison is that one case provides at least as much dialectical support for a conclusion as another case. This is captured by the condition that the same or more basic factors are relevant for the conclusion in the first case than in the second. Moreover, in the first case the same or fewer basic factors are relevant for the opposite conclusion than in the second case. See the example of the following figure.



In **Case1** there are more basic factors relevant for c than in **Case2**, since **Case1** has the basic factors l and $l \wedge \neg(b \wedge \neg c)$ that are relevant for c and that are lacking in **Case2**. At the same time **Case1** has fewer basic factors relevant for $\neg c$, since **Case2** in addition has the basic factors g and $g \wedge \neg a$ that are relevant for $\neg c$. As a result, one can say that **Case1** provides more dialectical support for c than **Case2**.

To define comparison outcomes in general, it is convenient to have a shorthand notation for the sets of basic factors relevant for a conclusion and its opposite, respectively. For the basic factors that are relevant for a conclusion the abbreviation ProBF is introduced, while for the basic factors relevant for a conclusion's opposite the abbreviation ConBF will be used. Moreover, the conclusion involved will be added as a subscript, while the same is done with the abbreviations CB and BG for the comparison basis and the background, respectively. Thus $\text{ProBF}_{\gamma, CB, BG}(\text{Case})$ denotes the set of basic factors relative to CB, which are relevant for γ in *Case* relative to CB and BG. Likewise,

$\text{ConBF}_{\gamma \text{ CB, BG}}(\text{Case})$ denotes the set of basic factors relative to CB, which are relevant for $\neg\gamma$ in *Case* relative to CB and BG.

With these abbreviations, the following definition formally states necessary and sufficient conditions for the comparison outcome that there is equal dialectical support in two cases. In words, these conditions hold that in the first case the same basic factors are relevant for the conclusion as in the second case, and the same are relevant for the conclusion's opposite.

Definition 7 (Equal dialectical support)

For any case *Case*, let $\text{ProBF}_{\gamma \text{ CB, BG}}(\text{Case})$ denote the set of basic factors relative to CB, which are relevant for γ in *Case* relative to CB and BG. Likewise, let $\text{ConBF}_{\gamma \text{ CB, BG}}(\text{Case})$ denote the set of basic factors relative to CB, which are relevant for $\neg\gamma$ in *Case* relative to CB and BG.

Then *Case1* and *Case2* provide *equal dialectical support* for γ relative to CB and BG, denoted $\text{Case1} \approx_{\gamma \text{ CB, BG}} \text{Case2}$, if and only if 1. and 2. hold:

1. $\text{ProBF}_{\gamma \text{ CB, BG}}(\text{Case1}) = \text{ProBF}_{\gamma \text{ CB, BG}}(\text{Case2})$, and
2. $\text{ConBF}_{\gamma \text{ CB, BG}}(\text{Case1}) = \text{ConBF}_{\gamma \text{ CB, BG}}(\text{Case2})$.

The next definition formally states the necessary and sufficient conditions under which one case provides at least as much dialectical support as another case. In words, these conditions say that the set of basic factors relevant for the conclusion in the first case is a (strict or non-strict) superset of the corresponding set of basic factors in the second case. At the same time the set of basic factors relevant for the opposite conclusion in the first case is a (strict or non-strict) subset of the corresponding set of basic factors in the second case.

Definition 8 (At least as much dialectical support)

For any case *Case*, let $\text{ProBF}_{\gamma \text{ CB, BG}}(\text{Case})$ denote the set of basic factors relative to CB, which are relevant for γ in *Case* relative to CB and BG. Likewise, let $\text{ConBF}_{\gamma \text{ CB, BG}}(\text{Case})$ denote the set of basic factors relative to CB, which are relevant for $\neg\gamma$ in *Case* relative to CB and BG.

Then *Case1* provides *at least as much dialectical support* for γ as *Case2* relative to CB and BG, denoted $\text{Case1} \succeq_{\gamma \text{ CB, BG}} \text{Case2}$, if and only if 1. and 2. hold:

1. $\text{ProBF}_{\gamma \text{ CB, BG}}(\text{Case1}) \supseteq \text{ProBF}_{\gamma \text{ CB, BG}}(\text{Case2})$, and
2. $\text{ConBF}_{\gamma \text{ CB, BG}}(\text{Case1}) \subseteq \text{ConBF}_{\gamma \text{ CB, BG}}(\text{Case2})$.

The third comparison outcome distinguished here is that one case provides strictly more dialectical support for a conclusion than another case. This outcome follows if at least one of the set inclusions in Definition 8 is strict. This happens precisely then if there is at least as much dialectical support in the first case as in the second, while the cases are not equal regarding their dialectical support. As a result, the outcome of strictly more dialectical support can be defined in terms of the previous two, as follows.

Definition 9 (Strictly more dialectical support)

Case1 provides *strictly more dialectical support for γ* than *Case2* relative to CB and BG, denoted $\mathbf{Case1} >_{\chi \text{ CB, BG}} \mathbf{Case2}$, if and only if $\mathbf{Case1} \geq_{\chi \text{ CB, BG}} \mathbf{Case2}$ and not $\mathbf{Case1} \approx_{\chi \text{ CB, BG}} \mathbf{Case2}$.

Thus defined, the outcomes of case comparison have some properties that are useful to understand the conditions under which a body of settled cases yields a conclusion in a problem case (see the next section, Section 3.7). One of these properties is *transitivity*. Transitivity informally means that if, for instance, one case provides at least as much dialectical support as a second, and the second as much as a third, then the first case also provides at least as much dialectical support as the third. Formally these properties can be stated as follows.

Transitivity

If $\mathbf{Case1} \approx_{\chi \text{ CB, BG}} \mathbf{Case2}$ and $\mathbf{Case2} \approx_{\chi \text{ CB, BG}} \mathbf{Case3}$, then $\mathbf{Case1} \approx_{\chi \text{ CB, BG}} \mathbf{Case3}$.

If $\mathbf{Case1} \geq_{\chi \text{ CB, BG}} \mathbf{Case2}$ and $\mathbf{Case2} \geq_{\chi \text{ CB, BG}} \mathbf{Case3}$, then $\mathbf{Case1} \geq_{\chi \text{ CB, BG}} \mathbf{Case3}$.

If $\mathbf{Case1} >_{\chi \text{ CB, BG}} \mathbf{Case2}$ and $\mathbf{Case2} >_{\chi \text{ CB, BG}} \mathbf{Case3}$, then $\mathbf{Case1} >_{\chi \text{ CB, BG}} \mathbf{Case3}$.

Transitivity follows trivially from the transitivity of the conditions of identity ($=$) and set inclusion (\supseteq) that must hold between the sets of relevant basic factors involved in Definition 7 and 8, respectively.

Two other properties are *invariance under negation* and *inversion under negation*. Informally, invariance under negation means that if two cases provide equal dialectical support for some conclusion, then the two also provide equal dialectical support for the opposite conclusion. Furthermore, the property of inversion under negation means that if one case provides at least as much dialectical support for some conclusion as another, then the second case provides at least as much dialectical support for the opposite conclusion as the first case. Formally these properties can be stated as follows.

Invariance and inversion under negation

$\mathbf{Case1} \approx_{\chi \text{ CB, BG}} \mathbf{Case2}$ if and only if $\mathbf{Case1} \approx_{\neg \chi \text{ CB, BG}} \mathbf{Case2}$.

$\mathbf{Case1} \geq_{\chi \text{ CB, BG}} \mathbf{Case2}$ if and only if $\mathbf{Case2} \geq_{\neg \chi \text{ CB, BG}} \mathbf{Case1}$.

$\mathbf{Case1} >_{\chi \text{ CB, BG}} \mathbf{Case2}$ if and only if $\mathbf{Case2} >_{\neg \chi \text{ CB, BG}} \mathbf{Case1}$.

These properties follow almost trivially from the definitions of comparison outcomes, and the observation that $\text{ProBF}_{\neg \chi \text{ CB, BG}}(\mathbf{Case}) = \text{ConBF}_{\chi \text{ CB, BG}}(\mathbf{Case})$ and $\text{ConBF}_{\neg \chi \text{ CB, BG}}(\mathbf{Case}) = \text{ProBF}_{\chi \text{ CB, BG}}(\mathbf{Case})$, the latter of which holds thanks to the convention that double negations vanish (Section 3.3).

3.6.7 Summary

Recapitulating the main points of Section 3.6:

- Case comparison involves comparing dialectical arguments regarding their dialectical support for a conclusion.
- For concluding to a comparison outcome regarding the relative dialectical support for a conclusion, it is not necessary to decide whether the conclusion or intermediate conclusions are actually derivable.
- It can itself be supported and attacked that one statement supports or attacks another, so that statements of support and attack must also be taken into account when comparing cases.
- A comparison basis is postulated as the division between what is relevant and irrelevant for the purpose of case comparison, and the comparison outcomes are stated relative to this comparison basis.
- For the purpose of comparing cases, the applying basic factors are derived from non-factors in combination with the background.
- Dialectical arguments represent an modelled version of the argumentation, including all relevant factor statements that can be made.

3.7 Which settled cases are relevant?

In this section the comparison outcomes are used to select the settled cases that are relevant to resolve a given problem case. To this end the problem cases and settled cases are formally defined. Roughly, problem cases are cases where some conclusion is disputed, while settled cases have the property that the conclusion was drawn and motivated by a judge. Moreover, settled cases are relevant for problem cases that provide at least as much dialectical support for their conclusion. Such relevant settled cases are called precedents.

A case base is presupposed as set of cases from which the precedents are selected. Informally, this set contains all cases that can be involved in reasoning by case comparison, and this includes not only settled cases, but problem cases as well. More specifically, it depends on the conclusion at stake whether a case is a settled case or a problem case.

Once the precedents have been selected, the question is whether they can be followed. In particular, what if the precedents contradict each other? To answer this question the settled cases are ordered regarding their dialectical support, and the cases and their ordering are depicted graphically as dots on a line.

3.7.1 Problem cases, settled cases and precedents

Informally a *problem case* for a conclusion is a case where the conclusion is an issue, in the sense that neither the conclusion nor its opposite was drawn. This is captured formally by the demand that the sentence used in the conclusion is not in the case representation, while the same holds for the negated sentence.

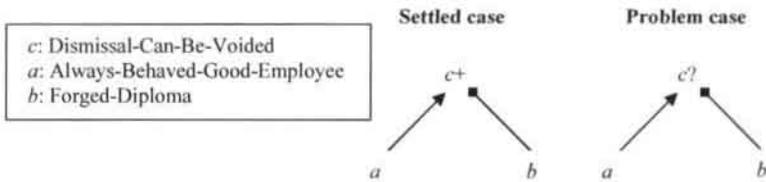
Definition 10 (Problem cases)

Case is a *problem case* for γ if and only if $\gamma \notin \mathbf{Case}$ and $\neg\gamma \notin \mathbf{Case}$.

Note that a problem case is a case according to this definition, thanks to the convention that strings of characters in bold type stand for case representations, that is, sets of sentences of the language CRL (Section 3.3).

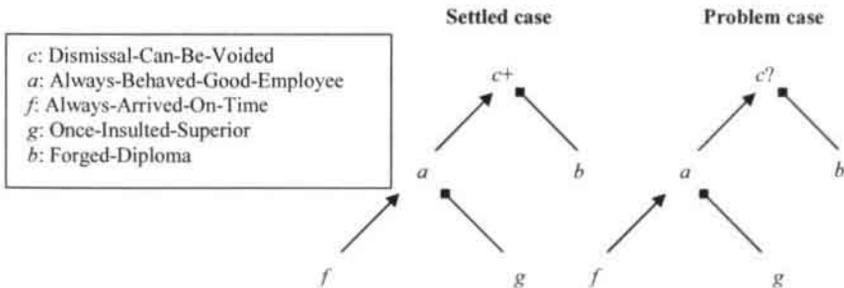
Informally, a *settled case* is a case in which a judge took a decision that carries judicial authority. Deciding cases typically involves resolving conflicts, and the resolution of these conflicts can help decide problem cases with comparable conflicts.

In the following a number of examples is given. In each example the basic factors are those at which no support or attack arrow is pointing.



In the settled case in this figure conclusion c is supported by a and attacked by b . Conclusion c was drawn in the settled case, as indicated by the plus sign. The settled case can be used to resolve the conflict in the problem case, and conclusion c can be derived there.

There can be more than one conflict in a settled case, as in the following figure.



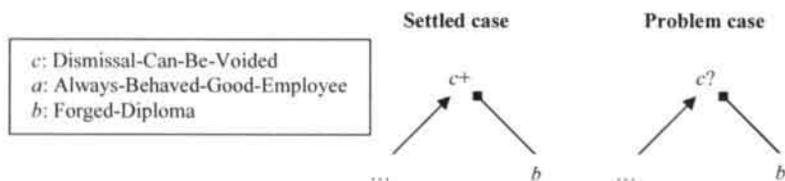
In this figure there is not only a conflict with respect to c , but also one with respect to a . While the former conflict was resolved by drawing c as a conclusion, no explicit decision was taken with respect to the latter conflict concerning a . Accordingly, only c is followed by an plus sign.

Since it is left implicit whether and how the conflict with respect to a is resolved, the settled case cannot be used to resolve this conflict in other cases. As the figure shows, however, the settled case can still be useful for resolving problem cases where c is an issue. What matters is that c was drawn in the settled case, and that the problem case provides equal dialectical support for this conclusion. The outcome that there is equal dialectical support holds irrespective of the issue whether or not the intermediate conclusion a was drawn in the settled case.

If a case has been decided by drawing a conclusion, then this is captured formally by the property that the state of affairs corresponding to the conclusion obtains in the case. If conclusion c was drawn in a case with case representation **Case**, for instance, then this is expressed by the sentence $c \in \mathbf{Case}$.

Settled cases typically involve conflicts that are resolved. Therefore a settled case where a conclusion was drawn must not involve the opposite conclusion at the same time. Thus if $c \in \mathbf{Case}$ then one must not have $\neg c \in \mathbf{Case}$ at the same time, because otherwise **Case** is not settled for c .

A complication is that it is possible that a conclusion was drawn without being supported at all, as in the following figure.



Such cases are not accepted as settled cases, however, because intuitively their decisions are not sufficiently motivated, and are therefore presently not assumed to carry authority. Formally this is captured by the requirement that if a case is to be a settled case for some conclusion, then the conclusion must be a grounded factor in the case. Recall that a factor is grounded in a case if it is supported by (conclusions supported by) the basic factors that apply to the case, whereby the applying basic factors are derivable from case facts and the background.

In sum, settled cases have the property that a grounded conclusion was drawn, such that the corresponding state of affairs obtains. Moreover, the opposite conclusion was not drawn, and its corresponding state of affairs does not obtain. The formal definition is

Definition 11 (Settled cases)

Case is a *settled case for γ* relative to CB and BG if and only if $\gamma \in \mathbf{Case}$ and $\neg\gamma \notin \mathbf{Case}$, and γ is grounded in **Case** relative to CB and BG.

Note that a settled case is a case according to this definition, thanks to the convention that case representations are denoted by strings of characters in bold type and starting with an uppercase.

Precedents are settled cases that can help resolve conflicts in a problem case. Informally, a precedent is a settled case for which it holds that the problem case provides at least as much dialectical support for the conclusion. The formal definition is as follows.

Definition 12 (Precedents)

A case **Precedent** is a *precedent for γ with respect to ProblemCase* relative to CB and BG if and only if

Precedent is a settled case for γ relative to CB and BG, and
ProblemCase is a problem case for γ , and
ProblemCase $\geq_{\chi_{CB, BG}}$ **Precedent**.

By convention a precedent for an *opposite* conclusion (like $\neg c$) is also said to be a precedent *against* the conclusion (c).

Note that precedents are defined with respect to a certain problem case. As a result, it is possible that a settled case is a precedent for one problem case, while it is not for another. This is an entirely appropriate choice for the case comparison method, where the selection of precedents is based on the mutual comparison of cases. The present definition of precedents does not correspond, though, to the view on precedents as general sources of law that are in principle applicable to all cases. This view would be more appropriate to the rule extraction method discussed in the previous chapter.

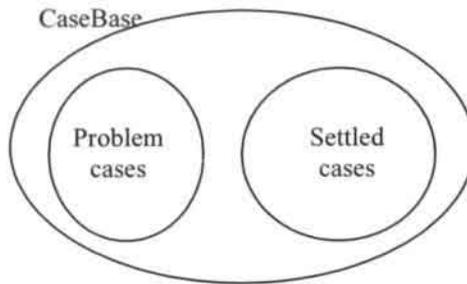
3.7.2 The case base

The definitions of the previous subsection state the conditions under which a case is a precedent at all. In case-based reasoning the precedents are always selected from a given pool of cases (cf. HYPO's Case Knowledge Base, Ashley 1990, p. 25; CATO's case database, Aleven 1997, p. 253; cf. also Kolodner's Case Library, 1993, Part II).

At this point a *case base*, denoted CaseBase, is therefore postulated. Informally the case base is the (finite) set of cases from which precedents must be selected to derive conclusions by case comparison. Formally, CaseBase is a (finite) set of case representations **Case**, where each case representation is a (finite) set of sentences of the case representation language CRL (Section 3.3).

The case at hand can, but need not, itself be included in the case base (cf. Ashley 1990, p. 131), because the case base only restricts the set of possible

precedents. Moreover, not all cases in the case base are necessarily settled cases, but problem cases can be included as well, for instance cases that have not been decided on all issues. More precisely, it depends on the conclusion at stake which of the cases in CaseBase are problem cases, and which are settled ones. See the following figure.

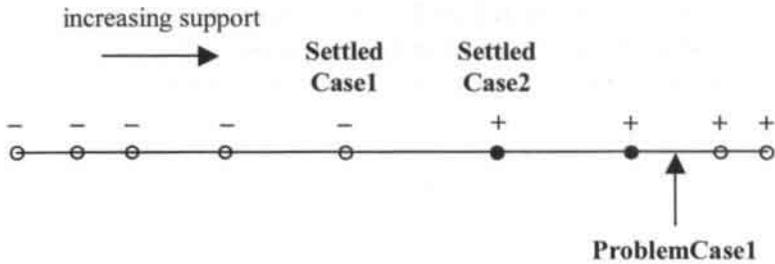


The large oval stands for the case base. The smaller oval inside and on the left represents the set of problem cases for a conclusion, for instance the conclusion that an employee's dismissal can be voided. On the right one finds a small oval representing the settled cases for the conclusion.

Given a case at hand, for each conclusion only precedents are selected that are among the settled cases in the case base. As will be explained at more length in Section 3.9, it therefore depends on the case base which conclusions are derivable by case comparison. As a result, these conclusions will be formally defined relative to the set CaseBase (Definitions 14 and 15).

3.7.3 *When do precedents conflict?*

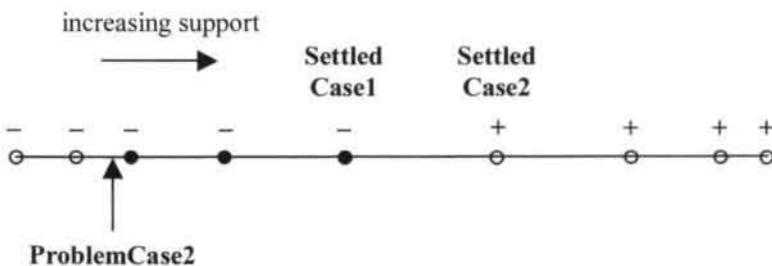
To clarify the conditions under which settled cases can help resolve problem cases, one can order the settled cases in the case base by their dialectical support for the disputed conclusion. The ordering by dialectical support is transitive (Subsection 3.6.8) and can therefore be visualised by placing the cases on a line. Note that not all cases in the case base are necessarily comparable regarding their dialectical support for the conclusion. In particular, not all precedents are always mutually comparable. Accordingly, each line can only represent one subset of mutually comparable cases in the case base, and it may not be possible to place all precedents on one line. However, despite this limitation it is enlightening to visualise the ordering by dialectical support, as in the following figure.



In this figure the small circles represent settled cases. The conclusion of each settled case is indicated by a plus or a minus sign. The plus signs stand for a conclusion (say c), while the minus signs stand for the opposite conclusion ($\neg c$). From left to right there is increasing dialectical support for the conclusion (c). As the plus and minus signs indicate, cases are decided negatively on the left, while on the right they are decided positively. Going from left to right, **SettledCase1** is the negatively decided case that provides the highest dialectical support. **SettledCase2** is the positively decided case with the least dialectical support.

ProblemCase1 is a problem case which has two positively decided cases on its left, represented by solid small circles. These are precedents for their conclusion (c), because **ProblemCase1** provides more dialectical support for the conclusion (c). Moreover, **ProblemCase1** is positioned such that all negatively decided cases are to the left of it. Accordingly, it provides more dialectical support for the conclusion (c) than the negatively decided cases. As a consequence of this and the property of inversion under negation, **ProblemCase1** provides *less* dialectical support for the *opposite* conclusion ($\neg c$). As a result, with respect to **ProblemCase1**, none of the negatively decided cases is a precedent for its negative conclusion ($\neg c$). The only two precedents with respect to **ProblemCase1** are the ones represented by solid circles.

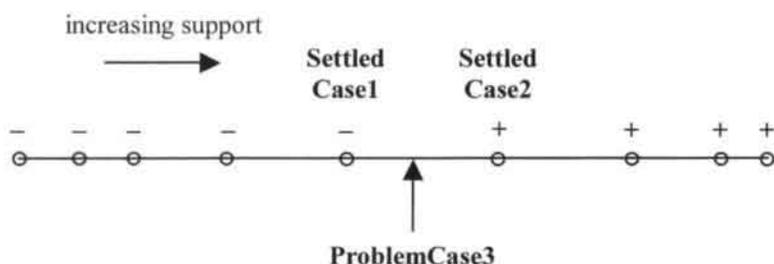
In the following figure, **ProblemCase2** is a problem case which has three negatively decided cases on its right, represented by solid small circles.



ProblemCase2 provides less dialectical support for the positive conclusion (c) than these negatively decided cases. In other words, it provides more dialectical support for the opposite conclusion ($\neg c$). As a result, the negatively decided

cases to the right of **ProblemCase2** are precedents for their negative conclusion ($\neg c$). Moreover, **ProblemCase2** is positioned such that all positively decided cases are to the right of it. Accordingly, it provides less dialectical support for the conclusion (c) than the positively decided cases. As a result, none of the positively decided cases is a precedent for its conclusion (c).

In the next figure, **ProblemCase3** is a problem case without positively decided cases on its left, and without negatively decided cases on its right.

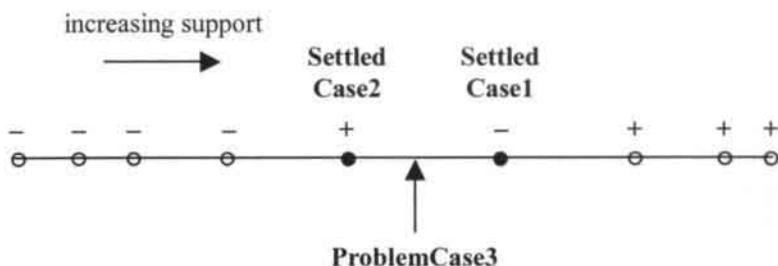


Accordingly, **ProblemCase3** provides less dialectical support for the positive conclusion (c) than all positively decided cases. It also provides less dialectical support for the opposite conclusion ($\neg c$) than all negatively decided cases. In other words, with respect to **ProblemCase3** none of the settled cases is a precedent. Therefore **ProblemCase3** cannot be decided by following precedents.

One might expect that all the positively decided cases always provide more dialectical support for the conclusion than all the negatively decided ones. If a case provides more dialectical support for a conclusion than another case that was positively decided, then one would expect the former case to be decided positively too. Likewise, if a case provides less dialectical support than one which was negatively decided, then one would expect the former case to be negatively decided. These expected decisions would be in accordance with the *stare decisis* principle as it is presently formalised, viz. as a kind of reasoning *a fortiori*.

In exceptional cases it is possible, however, that a judge decides a case without taking recourse to the *stare decisis* principle. A reason for this could be, for instance, that a judge wants to take into account a change of view on the law. Such a change can bring about, for example, that in dismissal cases it becomes much more important than before whether an employee has caused considerable damage due to carelessness. Another reason for a departure from *stare decisis* could be, that this is desirable in the light of the changing values and policies underlying the law. From the perspective of women's emancipation it could be desirable, for instance, that a female employee's dismissal can be voided, irrespective of decisions taken in comparable cases with male employees.

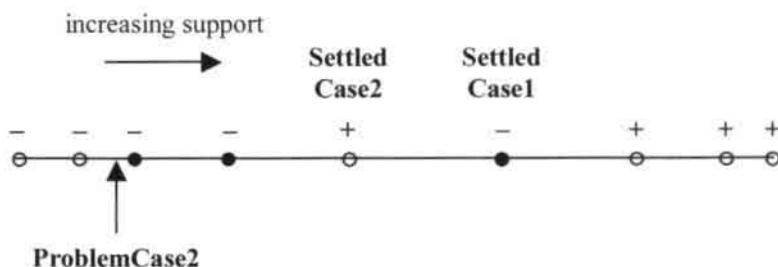
Whatever the reasons to depart from stare decisis, if it is done the set of precedents can become contradictory, and then it can become problematic to draw conclusions in problem cases. An example is in the following figure.



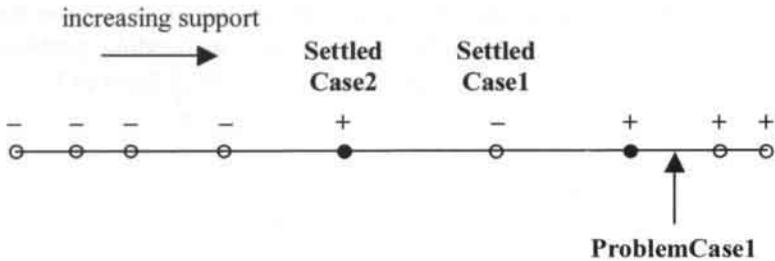
The main difference with the previous figure is that **SettledCase1** is now to the right of **SettledCase2**. Accordingly, there is more dialectical support for the conclusion (c) in **SettledCase1** than in **SettledCase2**. As a result, the two decisions are contradictory, in the sense that the case with more support was decided negatively, while in the case with less support a positive conclusion was drawn. In other words, for whatever reason, at least one of these cases must have been decided against the principle of stare decisis (unless the decisions were delivered simultaneously, or the later decision was made by a higher court).

Now consider **ProblemCase3**, which is positioned between these two settled cases. Obviously, **SettledCase1** is a precedent for the opposite conclusion ($\neg c$) with respect to **ProblemCase3**, while **SettledCase2** is a precedent for the conclusion itself (c). As a result, the precedents cannot both be followed here.

If **ProblemCase2** and **ProblemCase1** are considered, however, then both can still be resolved by following precedents. **ProblemCase2**, for instance, has negatively decided precedents only, and should accordingly be decided negatively. See the following figure.



Likewise, **ProblemCase1** has positively decided precedents only, and should therefore be decided positively. See the following figure.



For **ProblemCase1** and **ProblemCase2** it does not matter that **SettledCase1** and **SettledCase2** have been decided in a contradictory way, that is, against the principle of stare decisis. Both problem cases can be resolved here by following precedents.

This example shows that if precedents have been decided in a contradictory way and against the principle of stare decisis, then sometimes no conclusion follows from them in a problem case. It also shows, however, that even then there are situations in which an unambiguous conclusion follows. To derive such an unambiguous conclusion, it suffices to have or more precedents for the conclusion, and no precedents against.

3.7.4 Summary

Recapitulating the main points of Section 3.7:

- A settled case is relevant as a precedent for its conclusion in a problem case, if and only if the problem case provides at least as much dialectical support for the conclusion (Definition 12).
- As a result, settled cases are precedents relative to the comparison basis, because the outcomes of case comparison are.
- A case base is presupposed as a set of cases from which the precedents must be selected to derive conclusions.
- It is possible that relative to the comparison basis, settled cases have been decided in a contradictory way and against the principle of stare decisis as it is presently formalised.
- If that happens there can, but need not, be precedents for opposing conclusions, in which situation no conclusion follows in the problem case.

3.8 Reasoning by case comparison and reasoning with support and attack

This section formally defines the conclusions that follow in a case from the statements of case facts and background, given the comparison basis and the case base. The derivations discussed here are a generalisation of those according to the definition of CF-derivability above (Definition 3), in the sense that more conclusions can presently be derived thanks to the resolution of conflicts by a comparison with settled cases. Moreover, as before the derivation of conclusions can presently involve entangled statements, that is, it can be supported or attacked that a statement is a reason for or against a conclusion. For simplicity the examples below will not include these mechanisms of entanglement, however.

A special category among the derivable conclusions defined here are those that are directly derivable by following a precedent in the case base. Such conclusions will be called *derivable by case comparison*. Taking the conclusions derivable by case comparison as a starting point, the relations of *support and attack* allow for the derivation of additional conclusions.

In these additional derivations the conflict resolving potential of settled cases is exploited in two ways. First, it can be that an intermediate conclusion is derivable by case comparison. In that situation the intermediate conclusion can be used directly as a premise in the derivation of another conclusion. As a result, in the derivation of the latter conclusion the conflicts in connection with the intermediate conclusion do not play a role any more. The second way in which settled cases can allow for additional derivations is by blocking attacks. More specifically, when an attack involves an intermediate conclusion, then the attack is blocked by a precedent against this intermediate conclusion.

The derivations defined here depend on the outcomes of case comparison and the case base, since together these determine which settled cases can be followed as precedents. Moreover, the outcomes of case comparison in turn depend on the comparison basis and the background. As a result, the derivable conclusions will be defined relative to the comparison basis, the background and the case base.

More formally, the derivable conclusions defined here will be called CBR-*derivable* (where CBR stands for Case-Based Reasoning), to distinguish them from the CF-derivable ones introduced earlier. Formally the consequence relation is denoted \vdash , to which the three sets CB, BG and CaseBase are added as subscripts. At the premise side the consequence relation involves the union $Case \cup BG$ of a case representation *Case* with the background BG. Conclusions make use of sentences of the case representation language CRL, and one can accordingly let γ stand for the sentence used in a conclusion. Formally

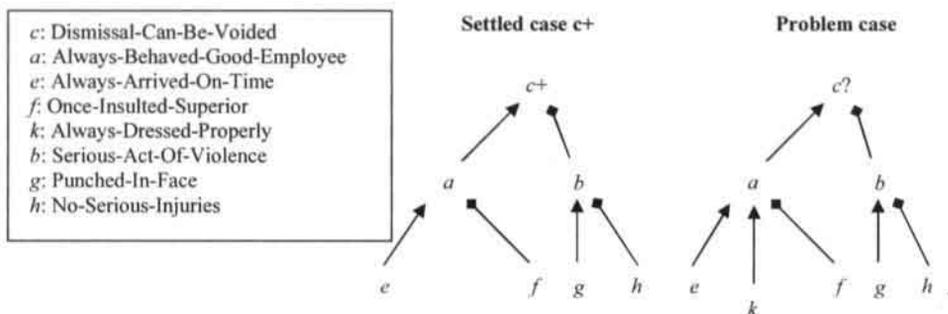
$$Case \cup BG \vdash_{CB, BG, CaseBase} \gamma$$

then says that relative to CB, BG and CaseBase, γ is CBR-derivable from *Case* and BG.

In the following, four examples are given of derivations by case comparison. In each example there is the same problem case with the same disputed conclusion, but the set of relevant settled cases is different. As before the basic factors involved are those at which no support or attack arrow is pointing.

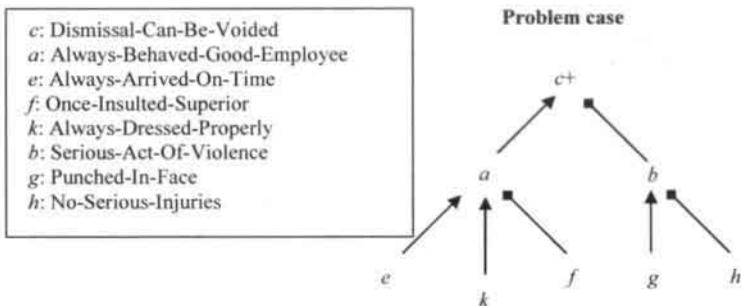
Example 1: A precedent for the conclusion

Consider an example derivation of a conclusion from a precedent, as in the following figure.



In this figure there is a problem case on the right where conclusion c is an issue, as indicated by the question mark. On the left there is a settled case where conclusion c was drawn, as indicated by the plus sign. At this point the set CaseBase contains this settled case for the conclusion only, and no other settled cases for the conclusion or its opposite. There is more dialectical support for conclusion c in the problem case than in the settled case, because of the additional relevant factor k in the problem case. As a result, the settled case can be followed as a precedent and conclusion c can be derived, provided that there are no precedents against c .

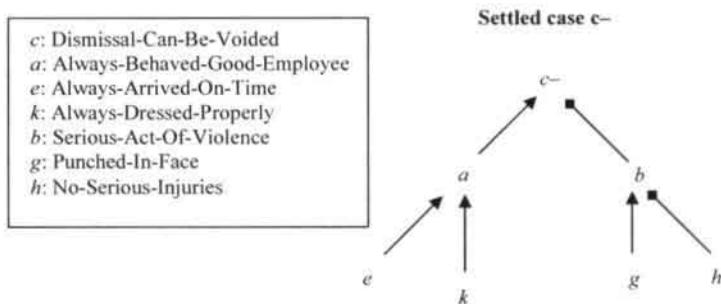
That conclusion c can be derived from a precedent can be indicated by labelling it with a plus sign in the dialectical argument of the problem case, as in the following figure.



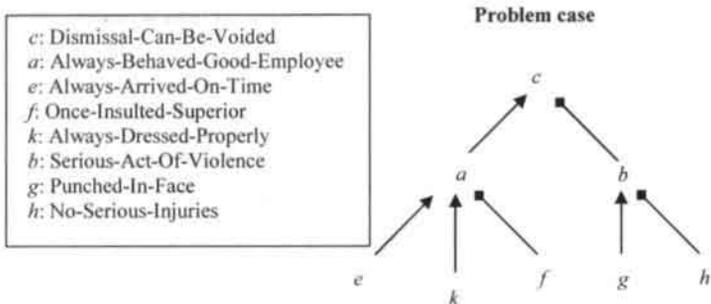
While conclusion c can apparently be derived on the basis of the settled case, one can leave implicit how the conflicts regarding a and b are to be resolved. It does not matter, for instance, whether or not the intermediate a actually followed in the settled case. All that matters is that conclusion c was drawn in the problem case, and that there is more dialectical support for c in the problem case. Accordingly, these intermediate conclusions were not labelled in the dialectical argument in the figure above.

Example 2: Conflicting precedents

There is one exception to the principle that precedents are always followed on the basis of the authority of their decisions, namely if these decisions involve opposing conclusions. Recall from the previous section that this happens if there is a precedent for and against a conclusion. To illustrate this with an example, suppose that next to **Settled Case c^+** one also has the case **Settled Case c^-** in CaseBase:



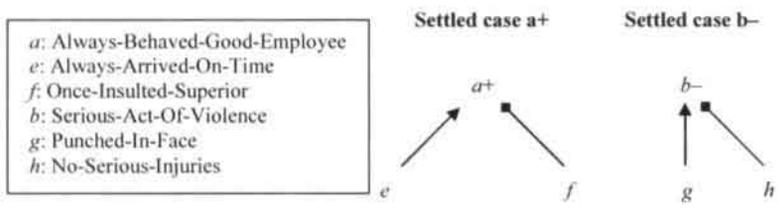
For reasons explained in the previous section, the choice is made here that the conclusion c is not derivable by case comparison. Briefly, these reasons are that the two conflicting precedents with opposing conclusions have been decided in a contradictory way, and that such precedents are of no use to derive conclusions. To indicate this the conclusion c is *not* labelled, as in the following figure.



Note that the conclusion *c* would not be labelled either if there were no precedents for or against it at all. This is in accordance with the fact that in the latter situation no conclusions follow by case comparison either.

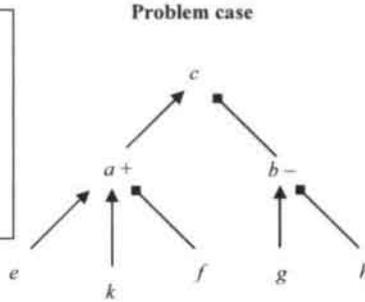
Example 3: Precedents for or against intermediate conclusions

Suppose that neither conclusion *c* nor its opposite $\neg c$ can be derived by case comparison. It could be, for instance, that no precedents for or against the conclusion can be found, or that there are conflicting precedents. Can the conclusion still be derived then with the help of other cases? Intuitively this is possible if the intermediate supporting *a* can be derived, while in addition there is no attack by the intermediate *b* because its opposite $\neg b$ follows. Now suppose that the following two settled cases are in CaseBase.



In the settled case on the left conclusion *a* was drawn, as indicated by the plus sign. On the right conclusion $\neg b$ was drawn, as indicated by the minus sign. Obviously, the problem case provides more dialectical support for *a* than the settled case on the left, because of the additional relevant *k*. The problem case provides as much dialectical support for $\neg b$ as the settled case on the right, because with regard to this conclusion the same factors are relevant in the former case as in the latter. As a result, both *a* and $\neg b$ can be derived in the problem case (provided that there are no precedents against these conclusions). Again this can be indicated by labelling these conclusions in the dialectical argument, as in the following figure.

c: Dismissal-Can-Be-Voided
 a: Always-Behaved-Good-Employee
 e: Always-Arrived-On-Time
 f: Once-Insulted-Superior
 k: Always-Dressed-Properly
 b: Serious-Act-Of-Violence
 g: Punched-In-Face
 h: No-Serious-Injuries



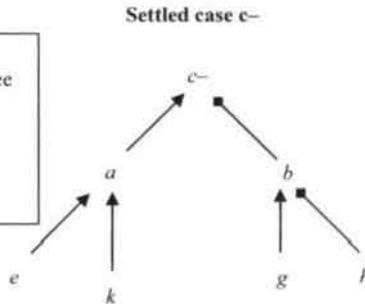
In this figure a is labelled + because it can be derived from precedent **Settled Case a+**, while b is labelled – because its opposite $\neg b$ can be derived from precedent **Settled Case b–**.

In this situation the choice is made to let c be derivable, namely on the basis of the support provided by a . The intuitive reason for this choice is that the conflict regarding the supporting a is resolved positively, while that regarding the attacking b is resolved negatively.

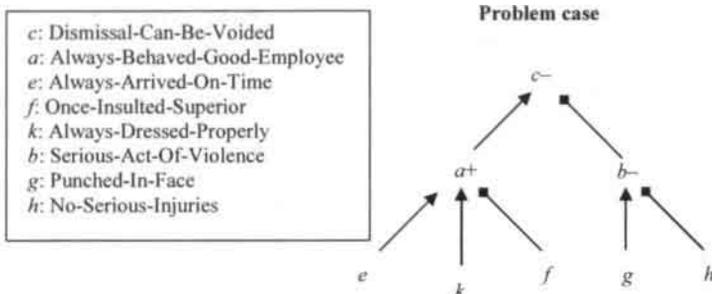
Example 4: Indirectly conflicting precedents

Suppose that in addition to the precedents **Settled Case a+** and **Settled Case b–** of the previous example, there is a precedent **Settled Case c–** against c in CaseBase. Suppose moreover that there is no precedent for c in CaseBase. See the following figure.

c: Dismissal-Can-Be-Voided
 a: Always-Behaved-Good-Employee
 e: Always-Arrived-On-Time
 k: Always-Dressed-Properly
 b: Serious-Act-Of-Violence
 g: Punched-In-Face
 h: No-Serious-Injuries



In this case the opposite conclusion $\neg c$ was drawn, as indicated by the minus sign. It is easily seen that the problem case provides more dialectical support for the conclusion $\neg c$, because of the additional relevant f . As a result, this settled case is a precedent for $\neg c$, so that apparently this conclusion is derivable in the problem case. This can be indicated by labelling the conclusion c with a minus sign, as follows.



At first sight this seems to be a somewhat strange situation. On the basis of the two precedents in CaseBase for *a* and against *b*, respectively, one would expect *c* to be derivable rather than its opposite, just as in the previous example. On the basis of the precedent against *c*, however, the opposite conclusion $\neg c$ apparently follows. It appears that the three precedents give conclusions that are *indirectly conflicting*, that is, conflicting in the light of the support relation with *a* and the attack relation with *b*.

In a situation with such indirectly conflicting precedents, one could choose to let no conclusion follow by case comparison. In effect this choice would come down to saying that judges' decisions are useless to derive conclusions, if they are not in accordance with inferences with support or attack that are possible on the basis of other decisions.

Here another approach is taken, however, in which the authority of all decisions is equally acknowledged, and in which reasoning by case comparison has primacy over reasoning with support and attack. More specifically, in this approach all precedents can make their conclusions follow immediately (if there are no precedents against), even if these conclusions are indirectly conflicting. Accordingly, no *ex falso* rule is authorised to the effect that anything would follow in a situation of such indirectly conflicting conclusions. Instead the conclusions following by case comparison are assumed as given, and support and attack relations can come into play only if reasoning by case comparison has failed to produce a conclusion.

In the situation above, for instance, the conclusions *a*, $\neg b$ and $\neg c$ are all derivable by case comparison. They cannot be questioned for being indirectly conflicting in the light of the support relation between *a* and *c*, and the attack relation between *b* and *c*. As in Example 3, these support and attack relations could still come into play, though, if the conflict regarding *c* is not resolved by case comparison.

These examples illustrate which conclusions follow by case comparison, in combination with the support and attack relations between statements. The conclusions following by case comparison are indicated by *labelling* them by a plus or a minus sign. In the following figure a schematic dialectical argument is drawn.

Definition 14 (Conclusions derivable by case comparison)

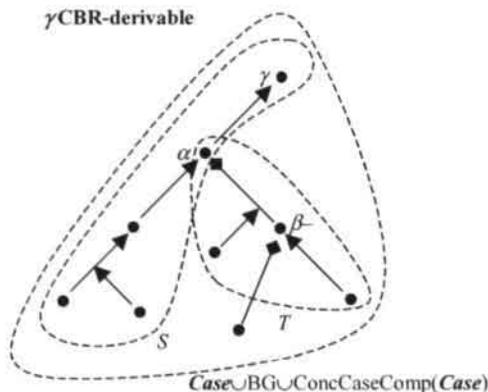
The set of conclusions of *Case derivable by case comparison* relative to CB, BG and CaseBase, denoted $\text{ConcCaseComp}_{\text{CB, BG, CaseBase}}(\text{Case})$, is defined as the set of conclusions γ labelled + in *Case* relative to CB, BG and CaseBase.

Note that this definition also includes the opposites of all conclusions that are labelled -, since these opposites are all labelled +.

Given the set of conclusions derivable by case comparison, additional derivations with support or attack can be made in the present model. There are two ways in which these additional derivations become possible. First, it can be that an intermediate conclusion is derivable by case comparison, such as the intermediate conclusion a in Example 3. The intermediate conclusion can then be used directly as a premise in the derivation of another conclusion, while conflicts in connection with the intermediate conclusion do not play a role any longer. Once conclusion a has been derived by case comparison in Example 3, for instance, it no longer matters that it is attacked by the factor f .

The second way in which settled cases allow for additional derivations is by blocking attacks, such as that by the factor b in Example 3. Specifically, when an attack involves an intermediate conclusion, then the attack is blocked by a precedent against this intermediate conclusion. In Example 3 the attack by the intermediate b was blocked because there was a precedent against it. As a result, conclusion c could be derived on the basis of the supporting factor a .

To capture all derivable conclusions, the union of the three sets *Case*, BG and $\text{ConcCaseComp}(\text{Case})$ is formed first. This set is then used as a basis for derivations with support or attack relations between statements. A situation where a conclusion is derivable is sketched schematically in the following figure.



The large area enclosed by a dashed line represents the set $Case \cup BG \cup ConcCaseComp(Case)$. This set includes all conclusions established by case-based inferences, and is now used as a basis for additional derivations with support or attack. A smaller area contains a chain of modus ponens inferences, leading to conclusion γ from a subset S of $Case \cup BG \cup ConcCaseComp(Case)$. If this modus ponens chain were the only part of the dialectical argument for γ then it would be derivable.

However, the modus ponens chain from S to γ involves an intermediate conclusion α that can be attacked. The attack is possible because there is a subset T of $Case \cup BG \cup ConcCaseComp(Case)$, from which $\neg\alpha$ follows by modus ponens. This is visualised by another part of the dialectical argument for γ which is in a third area enclosed by a dashed line. In principle such attacks could cause the conclusion γ to be not CBR-derivable.

However, the modus ponens chain leading to $\neg\alpha$ involves an intermediate conclusion β against which there is a precedent. In other words, the conflict regarding the intermediate conclusion β is resolved negatively, as becomes apparent by the fact that it is labelled $-$. Intuitively this means that the attack of α on the basis of T fails. In accordance with this, the choice is made here to let conclusion γ be CBR-derivable (provided that there are no other attacks that do not fail).

In sum, if a conclusion γ is to be CBR-derivable in the present model, then for a start there must be a subset S of $Case \cup BG \cup ConcCaseComp(Case)$ from which γ follows by modus ponens. In addition to this, all attacks must fail that are aimed at any intermediate conclusion α in the modus ponens chain leading to γ . More formally, suppose that there is a subset T from which $\neg\alpha$ follows by modus ponens. Then there must be a β following from T by modus ponens that is labelled $-$, and whose opposite $\neg\beta$ accordingly follows by case comparison. The formal definition is as follows.

Definition 15 (CBR-derivable conclusions relative to CB, BG and CaseBase)

A conclusion γ is CBR-derivable from $Case$ and BG relative to CB, BG and CaseBase, if and only if there is a subset $S \subseteq Case \cup BG \cup ConcCaseComp_{CB, BG, CaseBase}(Case)$ for which 1. and 2. hold:

1. γ follows by modus ponens from S ;
- and
2. for all subsets $T \subseteq Case \cup BG \cup ConcCaseComp_{CB, BG, CaseBase}(Case)$ for which there is an α such that α follows by modus ponens from S , while $\neg\alpha$ follows by modus ponens from T :
 - there is a β that follows by modus ponens from T and is labelled $-$ in $Case$ relative to CB, BG and CaseBase.

Observe first that according to this definition, all conclusions that are derivable by case comparison are also CBR-derivable. Each of those conclusions trivially follows by modus ponens from a subset S that contains only the sentence used

in the conclusion. Moreover, all attacks at such a conclusion fail, because the opposite conclusion is labelled –.

Observe further that the earlier definition of CF-derivability above (Definition 3) is a special case of this definition, because it deals with the situation that no precedents at all are taken into consideration as a means to resolve conflicts. As a result, in this situation no conclusions can follow by case comparison. Formally this means that no conclusions are labelled +, and that the set of conclusions derivable by case comparison is empty. This also means, however, that attacks cannot fail thanks to a precedent, because there are no conclusions labelled – either. It is relatively straightforward to see that under these conditions the definition of CBR-derivability reduces to that of CF-derivability.

Note finally that the definition of CBR-derivability also captures derivations involving entangled statements, that is, derivations in which it is supported or attacked that a statement is a reason for or against a conclusion. This is because the definition treats compound statements involving support or attack in the same way as elementary statements using atomic sentences. More specifically, compound support and attack statements can themselves act as (intermediate) conclusions, and they can as such be derived according to the definition.

The definition of CBR-derivability captures all conclusions that follow by case comparison, in combination with the support and attack relations between statements. The definition concludes the presentation of the present model's formalism proper. In the next section some applications of the model are presented.

3.9 Applications of the theory

This section presents two applications of the formal theory of case comparison that was exposed in the foregoing. First, two common reasoning patterns in case comparison are formalised, viz. analogising and distinguishing (Ashley 1990, pp. 25f.; Alevén 1997, pp. 58f.). Second, it is shown how a number of arguments on the importance of distinctions and similarities can be accommodated, viz. downplaying and emphasising (cf. Alevén 1997, pp. 62f.).

The account of analogising and distinguishing is an explication of earlier work on argument moves in case comparison (Roth 2001a; 2001b). While this earlier work merely indicated the types of move that could be made in case comparison, in the present account it is explained why certain moves count as analogising or distinguishing. Briefly, this is done by defining the factors that can be involved in analogising and distinguishing moves, in terms of relevance for the conclusion or its opposite.

3.9.1 Analogising and distinguishing

This subsection aims at a formal account of two reasoning patterns in case comparison, viz. analogising and distinguishing. Briefly, analogising comes down to pointing out significant similarities between cases, while distinguishing

involves pointing out significant differences. A significant similarity is a shared relevant factor, and a significant distinction is a relevant factor that is not shared, and that tends to make the settled case provide more dialectical support for the conclusion than the problem case. In other words, a significant distinction is a factor that prevents the settled case from becoming a precedent with respect to the problem case.

Suppose, for instance, that in a problem dismissal case the issue is whether the dismissal can be voided. Suppose, moreover, that in the problem case this conclusion can be supported by stating that the employee is highly esteemed as a colleague. Furthermore, suppose that a settled dismissal case is cited where the dismissal could be voided, and where the employee was also highly esteemed as a colleague. Then one can analogise the two cases by pointing out the high esteem for the employer as a significant similarity. Suppose in addition that in the settled case one can support the conclusion that the dismissal can be voided, by stating that the working atmosphere has not been affected by the dismissal. If this statement on working atmosphere cannot be made in the problem case, then it can be pointed out as a significant distinction.

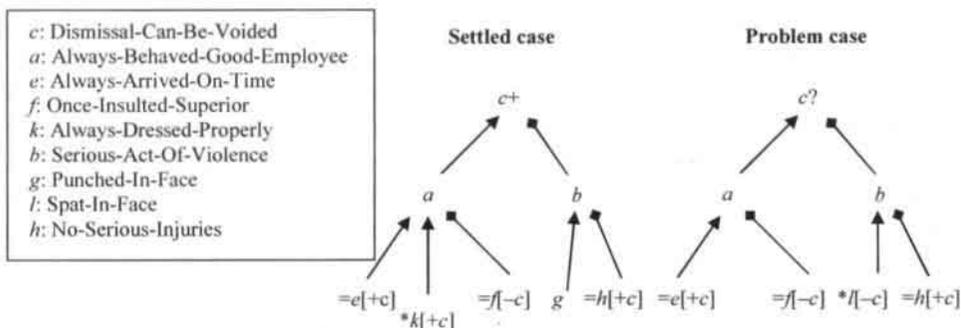
In the following a general definition is given of the factors that count as significant similarities and distinctions, respectively. The definition yields an explanation of why the factors count as significant similarities or distinctions, and this explanation is done in terms of the relevance relation defined earlier (Definition 6). In contrast with HYPO and CATO where only two opposing conclusions can be at issue, the significant similarities and distinctions are presently defined relative to any conclusion for which one likes to cite settled cases. If the issue is whether the dismissal can be voided, for instance, then the significant similarities and differences are different from when the issue is whether the employee has substantial interests in keeping his job.

The present account only involves applying basic factors in the analogising and distinguishing moves, leaving the generalisation to non-basic factors for future research. Another simplifying assumption is that the applying basic factors of each case are not relevant for a conclusion and the conclusion's opposite at the same time. This could happen, for instance, if a basic factor both supports and attacks a conclusion. A concrete example of this is that having substantial interests in keeping one's job would not only support the conclusion that the dismissal can be voided, but also attack this conclusion.

The basic factors involved in such unusual situations complicate the definition of significant similarities and distinctions. The intuitive reason for this is that similarities and distinctions should have only one effect on the dialectical support for a conclusion. A significant distinction that applies to the settled case and not to the problem case, for instance, should intuitively only strengthen the settled case for the conclusion (Aleven 1997, p. 60). Having substantial interests in keeping one's job, for example, normally does make one's case stronger for the conclusion that the dismissal can be voided. Thanks to the assumption that

no basic factor can be relevant for opposing conclusions, however, the unusual situations sketched above are ruled out in advance.

Consider the situation of the following figure.



The basic factors are those at which no support or attack arrow is pointing. Note that the two cases are incomparable regarding the dialectical support for the disputed conclusion c , so that the settled case cannot be followed as a precedent. As said the central question presently is which distinguishing and analogising moves can be made, however, rather than whether the disputed conclusion follows. Note further that following a convention adopted in CATO (Aleven 1997, p. 62), the significant similarities and distinctions involved in the analogising and distinguishing moves are marked with = and *, respectively. In addition it is indicated for each distinction and similarity whether it is relevant for conclusion c ($[+c]$), or for its opposite ($[-c]$).

Why, then, are e , f and h significant similarities? Note that the two cases are incomparable, so that one cannot say that e , f and h are significant similarities because they actually make both cases provide equal dialectical support for the conclusion.

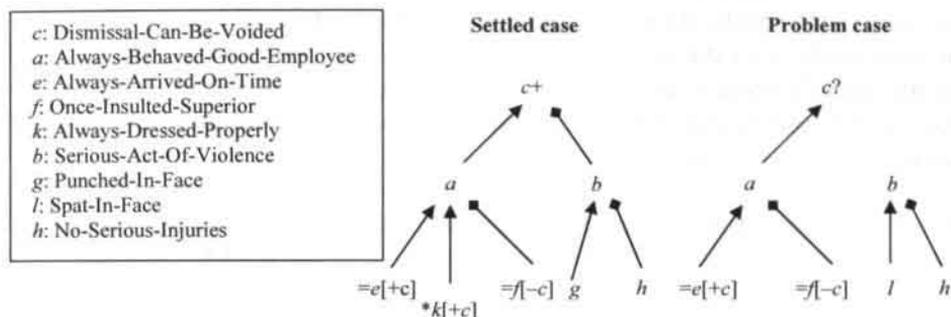
What one can say, though, is that in accordance with Definition 6, the shared factors e and h are relevant for the disputed conclusion c in each case. In other words, one can say that e and h are “shared strengths” for the conclusion (Ashley 1990, p. 66). Likewise, the factor f is relevant for the conclusion’s opposite $-c$. In other words, f is a “shared weakness” (Ashley 1990, p. 66).

In a similar way one can account for significant distinctions such as k and l . Intuitively, a distinction is a factor which is not shared and which tends to make the settled case provide more dialectical support for the conclusion than the problem case (cf. Ashley 1990, pp. 15 and 66). Again the problem is, however, that both cases are incomparable, so that one cannot say that k and l are distinctions because they actually make the problem case provide less dialectical support for the conclusion.

What one can say, though, is that the factor k that applies to the settled case is relevant for the disputed conclusion c . In other words, factor k is a strength of the settled case that is not shared by the problem case. Similarly, the factor l that

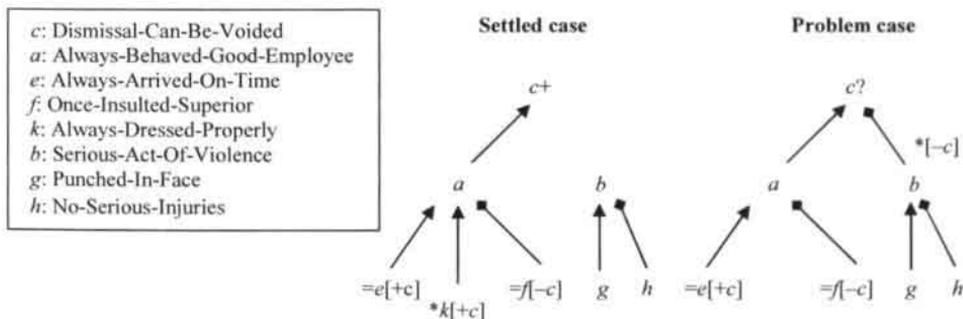
applies to the problem case is relevant for the opposite conclusion $\neg c$. In other words, factor l is a weakness of the problem case that is not shared with the settled case.

Next, consider the following situation.



The difference with the situation encountered before is that b cannot attack c in the problem case here because the corresponding attack statement does not apply. (The reason why it does not apply is not represented in the figure; there may be an attack by a non-factor, for instance.) As a result, in the problem case the factors l and h are not relevant for conclusion c or its opposite $\neg c$. Accordingly, one cannot distinguish the cases by pointing out l , and one cannot analyse them with h . In the figure this becomes evident in the fact that neither of these two is marked with $=$ or $*$.

As an example of distinguishing with a compound basic factor, consider the following situation.



The difference with the previous example is that the basic factor that b attacks c now applies to the problem case instead of the settled case. As a result, in the problem case the basic factor that b attacks c is relevant for the opposite conclusion $\neg c$. Accordingly, one can distinguish the cases by pointing out that b attacks c in the problem case, and this compound statement is therefore marked $*$ in the figure. Note that the factor g applies to both cases but is only relevant in

the problem case, namely for the opposite conclusion $\neg c$. Presently the choice is made that such shared factors cannot be used to distinguish cases, and accordingly factor g is not marked * in the problem case. The intuitive reason for this choice is that a shared factor is not a difference between cases, and can therefore not be used to distinguish them either. For the purpose of distinguishing cases it can still be exploited, though, that a shared factor is only relevant in one case and not in the other. However, to this end one must use other basic factors, namely those that cause the shared basic factor to be relevant in one case and not in the other. An example of such a basic factor is the statement that b attacks c in the situation of the figure.

Summing up the observations made in the examples, the following points can be made.

- Significant similarities are basic factors that apply both to the settled case and to the problem case.
- A significant similarity is relevant for the conclusion in the settled case as well as in the problem case, or relevant for the conclusion's opposite in the settled case as well as in the problem case.
- Significant distinctions are factors that apply to the settled case and not to the problem case, or vice versa.
- If a significant distinction applies to the settled case and not to the problem case, then in the settled case the distinction is relevant for the conclusion.
- If a significant distinction applies to the problem case and not to the settled case, then in the problem case the distinction is relevant for the conclusion's opposite.

These points straightforwardly lead to the following formal definition.

Definition 16 (Significant similarities and distinctions)

Let *SettledCase* be a settled case for γ relative to CB and BG, and let *ProblemCase* be a problem case for γ .

1. Similarities:

A factor σ is a *significant similarity for γ between *SettledCase* and *ProblemCase* relative to CB and BG*, if and only if a. or b. hold

a. σ is a basic factor relative to CB that applies to *SettledCase* and to *ProblemCase* relative to CB and BG, and is relevant for γ in *SettledCase* and in *ProblemCase* relative to CB and BG;

or

b. σ is a basic factor relative to CB that applies to *SettledCase* and to *ProblemCase* relative to CB and BG, and is relevant for $\neg\gamma$ in *SettledCase* and in *ProblemCase* relative to CB and BG.

2. Distinctions:

A factor δ is a *significant distinction for γ between *SettledCase* and *ProblemCase* relative to CB and BG*, if and only if a. or b. hold

a. δ is a basic factor relative to CB that applies to *SettledCase* relative to CB and BG and does not apply to *ProblemCase* relative to CB and BG, and is relevant for γ in *SettledCase* relative to CB and BG;

or

a. δ is a basic factor relative to CB that applies to *ProblemCase* relative to CB and BG and does not apply to *SettledCase* relative to CB and BG, and is relevant for $\neg\gamma$ in *ProblemCase* relative to CB and BG.

With this definition analogising comes down to pointing out a significant similarity, while distinguishing involves a significant distinction.

What the definition shows is that the model proposed in this chapter can deal with the well-known analogising and distinguishing argument moves in case comparison. More specifically, it is explained why certain basic factors count as significant similarities or distinctions, and this explanation is done in terms of the relevance relation of Definition 6. In this way it is also shown that analogising and distinguishing can be defined relative to the conclusion at stake, because different conclusions yield different sets of significant similarities and distinctions.

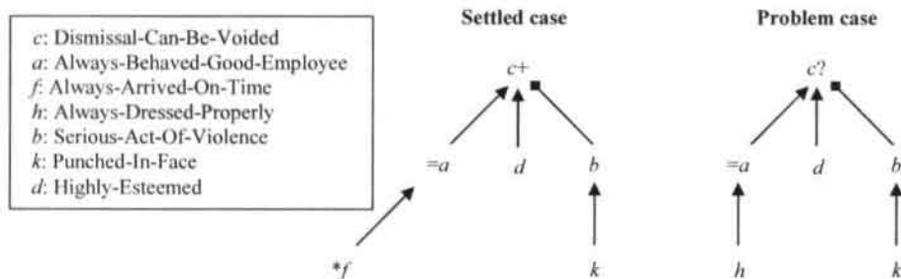
The set of significant similarities could also be used to define a measure of on pointness (cf. Ashley 1990, pp. 128f.). Briefly, then, a settled case would be more on point to the extent that it shares more significant similarities with the problem case.

3.9.2 Downplaying and emphasising

An interesting application of the definition of distinctions and similarities is an account of argument moves on the importance of distinctions and similarities, viz. downplaying and emphasising (cf. Alven 1997, pp. 62f.; see also Subsection 4.2.2). Intuitively, downplaying is showing that a distinction or similarity is not very important, while emphasising is showing that it is.

In the following it will be illustrated with examples how downplaying and emphasising move types can be accommodated within the present formalism. Briefly, this is done by making a local shift of the comparison basis to a more abstract or a more concrete level, and by comparing the situation before and after the shift. Each downplaying or emphasising argument type thus accommodated will be described informally, yielding a descriptive account of what downplaying and emphasising is within the present formalism. No heuristics are formulated to find a 'strategically best' downplaying or distinguishing move in a concrete context, however (cf. CATO's Heuristic Criteria, Alven 1997, pp. 68f.; see also Subsection 4.2.3). Moreover, no claim is made that the present downplaying and emphasising move types make up an exhaustive list, leaving a more systematic approach for future research.

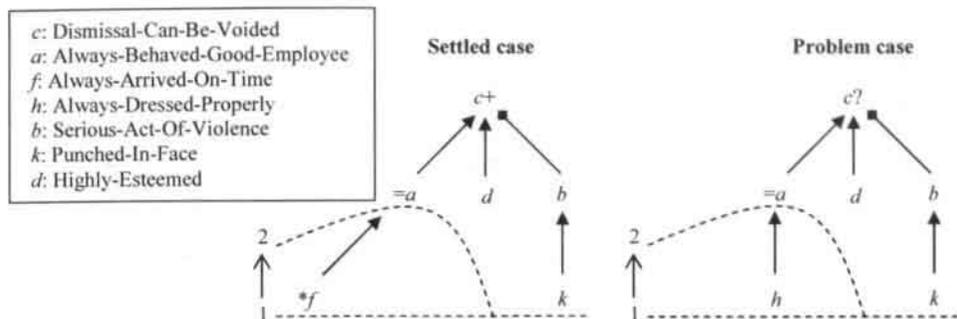
Suppose that a settled case is cited for the conclusion (*c*) that the dismissal can be voided, as in the following figure.



Suppose, moreover, that the settled case is distinguished from the problem case by pointing out the significant distinction (*f*) that the employee always arrived on time, as indicated by the *. How can this distinction be downplayed or emphasised?

Downplaying with an abstract interpretation that is a significant similarity

This way of downplaying is done by replacing a distinction with an abstract interpretation of it. Then the cases are compared in terms of this abstract interpretation. The abstract interpretation is a significant similarity. See the following figure.



In this figure there are two comparison bases. One is indicated with a straight dashed line at level 1, a second is represented by a curved dashed line at level 2.

The figure shows that if the comparison basis at level 1 is chosen, then the factor (f) that the dismissed employee always arrived on time is a significant distinction (marked *). The distinction applies to the settled case and not to the problem case, and it is relevant for the more abstract factor (a) that the person always behaved like a good employee.

The comparison can also be done in terms of this more abstract factor, which comes down to choosing another comparison basis. This second comparison basis differs from the first in that the original factor f is replaced with its abstract interpretation a . In the figure this second comparison basis is indicated by the curved dashed line at level 2.

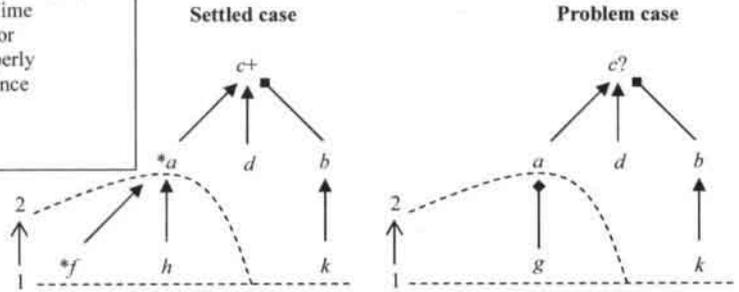
Relative to this second comparison basis the factor f is no significant distinction because it is not a factor. Its abstract interpretation a is a basic factor that applies to both cases, because it is supported and not attacked. In the settled case the factor a is supported by the statement f , and in the problem case the factor a is supported by the statement (h) that the person was always dressed properly. As a result, relative to the second comparison basis the factor a is a significant similarity between the cases, and is therefore marked =.

In general, downplaying with an abstract interpretation comes down to comparing two situations. Relative to one comparison basis there is a significant distinction (e.g., f). In another comparison basis this distinction is replaced by a factor (e.g., a) for which the original distinction is relevant: an abstract interpretation. Relative to the second comparison basis, this abstract interpretation (a) is a significant similarity.

Emphasising with an abstract interpretation that is a significant distinction

Emphasising a distinction also involves an abstract interpretation of the distinction. The difference is, though, that here the abstract interpretation is a significant distinction and not a significant similarity. See the following figure.

c: Dismissal-Can-Be-Voided
a: Always-Behaved-Good-Employee
f: Always-Arrived-On-Time
g: Once-Insulted-Superior
h: Always-Dressed-Properly
b: Serious-Act-Of-Violence
k: Punched-In-Face
d: Highly-Esteemed



Again the focus is on the distinction *f*, which is why the other distinctions *h* and *g* are not marked. The original distinction *f* is replaced with its abstract interpretation *a*, and again this yields another comparison basis at level 2.

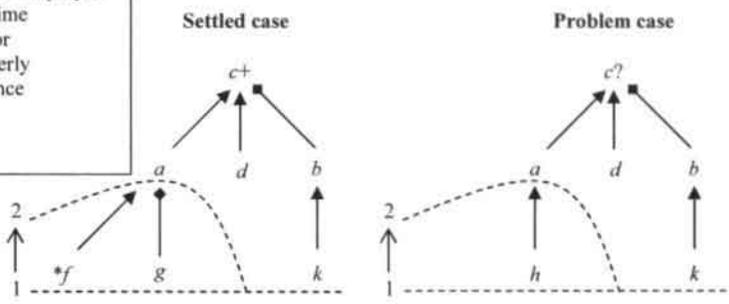
This abstract interpretation *a* is a basic factor that applies to the settled case but not to the problem case, because it is attacked there by the statement (*g*) that the employee once insulted a superior. As a result, relative to the second comparison basis the factor *a* is a significant distinction between the cases, and it is therefore marked *.

In general, emphasising with an abstract interpretation comes down to comparing two situations. Relative to one comparison basis there is a significant distinction (e.g., *f*). In another comparison basis this distinction is replaced by a factor (e.g., *a*) for which the original distinction is relevant: an abstract interpretation. Relative to the second comparison basis, this abstract interpretation (*a*) is a significant distinction as well.

Downplaying with an abstract interpretation that is not significant

This way of downplaying is also done by replacing a distinction with an abstract interpretation of it. Then the cases are compared in terms of this abstract interpretation. The difference with the first way of downplaying is, though, that here the abstract interpretation is neither a significant similarity, nor a significant distinction. See the following figure.

- c: Dismissal-Can-Be-Voided
- a: Always-Behaved-Good-Employee
- f: Always-Arrived-On-Time
- g: Once-Insulted-Superior
- h: Always-Dressed-Properly
- b: Serious-Act-Of-Violence
- k: Punched-In-Face
- d: Highly-Esteemed



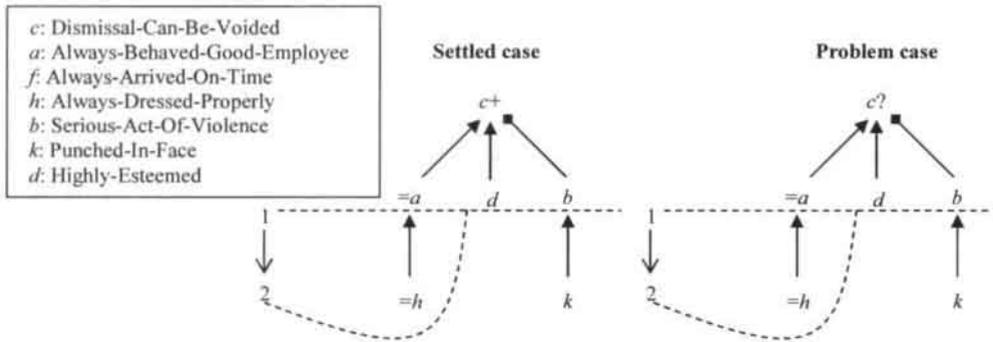
Again the distinction *f* is to be downplayed by shifting the comparison basis to level 2. Relative to the second comparison basis, the abstract interpretation *a* is a basic factor that does not apply to the settled case, because it is attacked there by the statement (*g*) that the employee once insulted a superior. In the problem case the abstract interpretation *a* does apply because it is supported by the statement (*h*) that the employee was always dressed properly. However, the factor *a* is not a significant distinction because in the problem case it is relevant for the conclusion (*c*) that the dismissal can be voided. As a result, relative to the second comparison basis the factor *a* is neither a significant similarity, nor a significant distinction.

In general, downplaying with an opposite interpretation comes down to comparing two situations. Relative to one comparison basis there is a significant distinction (e.g., *f*). In another comparison basis this distinction is replaced by a factor (e.g., *a*) that is relevant for the conclusion (e.g., *c*) and for which the original distinction (*f*) is relevant: an abstract interpretation. Relative to the second comparison basis, this abstract interpretation (*a*) is neither a significant similarity, nor a significant distinction.

In a similar way one can downplay or emphasise significant similarities, as in the following examples.

Emphasising with a specification that is a significant similarity

This way of emphasising is done by replacing a similarity with a specification of it. Then the cases are compared in terms of this specification. The specification then is a significant similarity as well. See the following figure.



Here the focus is on the similarity *a*, which is why the other similarities are not marked.

In the figure there are two comparison bases. One is indicated with a straight dashed line at level 1, a second is represented by a curved dashed line at level 2.

The figure shows that if the comparison basis at level 1 is chosen, then the factor (*a*) that the dismissed person always behaved like a good employee is a significant similarity (marked =). The similarity applies to both cases, where it is supported by the more concrete statement (*h*) that the person was always dressed properly.

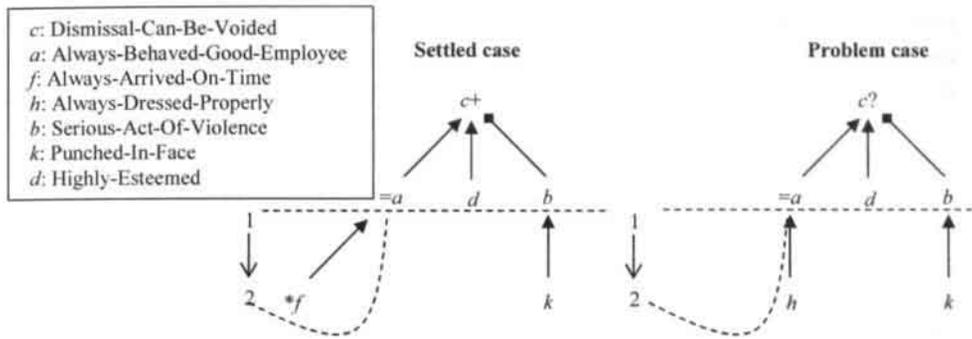
The comparison can also be done in terms of this more concrete factor, which comes down to choosing another comparison basis. This second comparison basis differs from the first, in that the original factor *a* is replaced with its specification *h*. In the figure this second comparison basis is indicated by the curved dashed line at level 2.

Relative to this second comparison basis the factor *a* is no significant similarity. Its specification *h* is a basic factor that applies to both cases. As a result, relative to the second comparison basis the factor *h* is a significant similarity between the cases, and is therefore marked =.

In general, downplaying with a specification comes down to comparing two situations. Relative to one comparison basis there is a significant similarity (e.g., *a*). In another comparison basis this similarity is replaced by a factor (e.g., *h*) that is relevant for the similarity: a specification. Relative to the second comparison basis, the specification (*a*) is a significant similarity as well.

Downplaying with a specification that is a significant distinction

This way of emphasising is done by replacing a similarity with a specification of it. Then the cases are compared in terms of this specification. The specification is a significant distinction. See the following figure.



Again the focus is on the similarity *a*, and again other similarities are not marked.

The figure shows that if the comparison basis at level 1 is chosen, then the factor (*a*) that the dismissed person always behaved like a good employee is a significant similarity (marked =). The similarity applies to both cases. In the problem case the similarity is supported by the more concrete statement (*h*) that the person was always dressed properly. In the settled case it is supported by the more concrete statement (*f*) that the person always arrived on time.

The comparison can also be done in terms of the latter more concrete factor, which comes down to choosing another comparison basis. This second comparison basis differs from the first, in that the original factor *a* is replaced with its specification *f*. In the figure this second comparison basis is indicated by the curved dashed line at level 2.

Relative to this second comparison basis the factor *a* is no significant similarity. Its specification *f* is a basic factor that applies to the settled case and not to the problem case. As a result, relative to the second comparison basis the factor *f* is a significant distinction between the cases, and is therefore marked *. (Note that this distinction can, of course, in turn be downplayed by shifting the comparison basis back to level 1.)

In general, downplaying with a specification comes down to comparing two situations. Relative to one comparison basis there is a significant similarity (e.g., *a*). In another comparison basis this similarity is replaced by a factor (e.g., *f*) that is relevant for the similarity: a specification. Relative to the second comparison basis, the specification (*f*) then is a significant distinction.

These examples show that downplaying and emphasising can be accommodated within the present model, namely as a comparison between two situations with different comparison bases. Distinctions can be downplayed or emphasised by choosing the comparison basis at a more abstract level. Downplaying or emphasising similarities is done by choosing a more concrete comparison basis.

In all of these examples there was only one suitable abstraction or specification. However, in general there will be more than one way of downplaying or emphasising within the present model, because there can be more

than one suitable abstraction or specification. The reason for this is that within the entangled factor hierarchies involved there can be different specifications relevant for a factor, and there can be different abstractions for which a factor is relevant.

As a result, a choice often has to be made between the different abstractions or specifications that are suitable for downplaying or emphasising. However, the present account given is merely descriptive, and no heuristics are proposed to select the abstractions or specifications that serve the end of downplaying or emphasising in a 'strategically best' way (cf. CATO's heuristic criteria to select Focal Abstractions, Alevén 1997, pp. 68f.).

3.10 Discussion

This chapter started with an extended example, after which a formalisation was presented of the method of reasoning by case comparison. A number of comparison outcomes were defined, yielding a criterion for when settled cases are relevant as precedents. After that the conclusions following by case comparison were defined. A number of reasoning patterns typically found in case comparison were also accounted for, such as analogising and distinguishing. This section briefly addresses some points regarding legal case-based reasoning that have not been addressed, and it discusses the possibilities to deal with these points within the proposed theory.

3.10.1 Reasoning with hypothetical cases

In case-based reasoning one often deals with hypothetical rather than real cases. It is possible, for instance, that a judge explores the consequences of a decision for hypothetical future cases, thereby trying to avoid consequences that are undesirable in the light of goals or policies underlying the law. If the dismissal of a woman employee cannot be voided, for example, then from a perspective of women's emancipation this could have unwanted effects for comparable later cases. This kind of reasoning is called consequentialist or forward-looking case-based reasoning (MacCormick 1997, p. 129 f.; cf. also Alexy 1978, p. 247).

Another possibility is that a judge considers hypothetical variations of a problem case and tries to attach conclusions to those variations (Ashley 1990, pp. 84-86 and 148-155; Wiarda, pp. 107f.; Nieuwenhuis 1976, pp. 507f.). Then by comparison with the actual problem case one can attempt to arrive at a reasonable decision for that case. If an employee is dismissed for having caused considerable damage despite a warning in advance, for instance, then the judge may consider a hypothetical variant where no warning in advance was given. Then the judge may try to decide whether in this hypothetical variant the dismissal could be voided.

In principle both ways of reasoning with hypothetical cases could be accommodated in the proposed model as well, namely by including them in the case base. Then the comparison mechanism could be applied to the hypothetical cases in the same way as it is applied to real ones. More specifically, the

comparison of cases would yield hypothetical problem cases and hypothetical precedents. By comparing a hypothetical problem case to cases in the case base, for instance, the comparison mechanism could tell which conclusions would follow in the hypothetical problem case. Likewise, one could also tell which conclusions would come to hold on the basis of hypothetical precedents included in the case base.

3.10.2 *Fact types in cases*

The case representation language only has two connectives, namely one for negation and one for the support relation between statements. As a result, phenomena like conjunction and disjunction ('and' and 'or') cannot be modelled in the present approach. It is not possible, for instance, to combine several supporting statements into one conjunctive reason. The reason for keeping the language abstract is that in the present model the focus is on the role of support and attack in case comparison and in the derivation of conclusions, and that a richer language would distract from these mechanisms. To accommodate other desired phenomena one could extend the case representation language with extra features, and subsequently generalise the definitions of derivable conclusions and of comparison outcomes.

The case representation language does not have the means to express weighing or priority information. It is not possible for instance, to make the statement that the reasons for a conclusion outweigh those against. Accordingly, it is not possible to resolve such conflicts on the basis of explicit weighing knowledge. Instead the only way to resolve conflicts in the present model is by resorting to precedents.

Comparison outcomes cannot be expressed in the case representation language either. In reality it may be, however, that the conclusion is drawn in a case that a certain comparison outcome holds. It is conceivable, for instance, that a judge decides that two cases are comparable regarding the issue whether a dismissal can be voided.

An obvious possible way to accommodate situations like this is by including sentences expressing comparison outcomes in case representations. For example, suppose that in **Case3** a judge concludes that relative to CB and BG, **Case1** and **Case2** provide equal dialectical support for conclusion c . Then this could be formalised by including the sentence $\text{Case1} \approx_{c, \text{CB}, \text{BG}} \text{Case2}$ in the case representation **Case3**: $(\text{Case1} \approx_{c, \text{CB}, \text{BG}} \text{Case2}) \in \text{Case3}$.

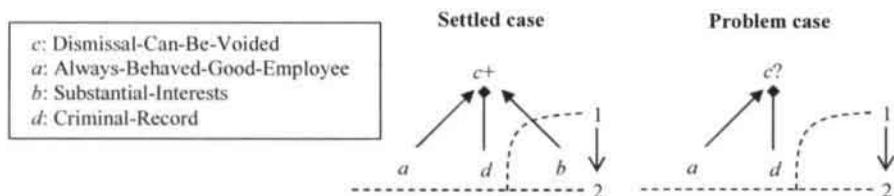
3.10.3 *Overruling precedents*

Another pattern of reasoning with cases that is sometimes distinguished is overruling the decision made in a settled case (Cross 1977, pp. 129-131). The authority to overrule decisions made by judges is usually attributed to courts higher up in the judicial hierarchy. From the perspective of the proposed model, overruling in effect comes down to saying that although the settled case is a precedent its decision was wrong, so that it does not have to be followed. This

kind of argument by judicial authority is not modelled in the approach proposed in this chapter. To accommodate it in the present model one would have to assume an ordering on the cases, to reflect the hierarchical relations between the courts involved. Note that such an ordering would also be useful in a situation of conflicting precedents, because it could be used to break ties by preferring precedents decided by higher courts to those decided by lower courts.

Another form of authority that is sometimes attributed to courts is the power to introduce distinctions that were originally not made (Raz 1979, pp. 185f.). Such distinctions tend to be factors which were not taken into consideration when the prior decision was made, but which are in retrospect seen as relevant for that decision. Suppose, for instance, that an employee is dismissed for having insulted a superior. Suppose, moreover, that the judge decided that the dismissal could be voided, without taking into consideration that the employee had substantial interests in keeping his job. Then if a court is confronted with this settled case, it could declare that the employee's substantial interests were in fact very relevant for the conclusion that the dismissal could be voided.

From the perspective of the model proposed in this chapter, this kind of reasoning would come down to enforcing, with judicial authority, a change in the comparison basis. For instance, if a court declares that having substantial interests in keeping a job supports that one's dismissal can be voided, then this would amount to including this as a factor in the comparison basis. This kind of reasoning is not modelled in the present approach, but this could be done straightforwardly, in a way similar to that in which the downplaying and emphasising moves were introduced above (Section 3.9). The only difference is that in these moves the comparison basis was changed by replacing a factor with another one, while the change is presently brought about by adding a factor to the comparison basis. See the following figure.



This figure shows a situation in which the comparison basis is at first in correspondence with the dashed line at level 1. Relative to this comparison basis the two cases provide equal dialectical support for the conclusion (*c*) that the dismissal can be voided. Accordingly, the settled case can be followed then. Suppose, however, that the court has the authority to declare that the statement (*b*) is relevant that the employee has substantial interests in keeping his job. Then this comes down to changing the comparison basis by adding this statement as a relevant factor, in accordance with the straight line at level 2. As

a result of this addition of the extra relevant statement to the comparison basis, the problem case provides less support for the conclusion (c) that the dismissal can be voided. Accordingly, the settled case cannot be followed then.

This example shows that the present model can account for changes in the comparison basis, which result from decisions by judges. The example does not answer the question, however, on what grounds such changes are enforced and, more generally, how the comparison basis must be chosen in the first place. This question is briefly addressed in the next subsection.

3.10.4 *How to choose the comparison basis?*

In the proposed model a comparison basis was introduced explicitly as a division between the statements that are relevant in case comparison, and those that are not. The explicit introduction of a comparison basis can be seen as a contribution to the on-going discussion in legal theory on the 'problem of importance' (Burton 1985, pp. 31f. and Part II), that is, the problem which case features are relevant in case comparison.

The problem of importance can be addressed in different ways from a legal theoretical stance. When following a method proposed by Goodhart (1930, pp. 161f.; 1959, pp. 117f.), for instance, one would have to isolate from each opinion the 'material facts', that is, the facts that the deciding court deemed relevant for its decision. This method has been criticised with some vigour by Stone (1959, pp. 597f.; 1964, pp. 271f.), however, on the grounds that different reasonable selections of material facts are often possible, and that facts selected as material can usually be stated at varying levels of generality. Further pursuing Stone's argument in this respect, one can conclude that it depends on later decisions which case features are relevant for the purpose of case comparison. A similar view is held by Raz (1979, pp. 183f.), who stresses the importance of distinguishing as a judicial instrument to change rules of case law.

Which grounds could underlie such a change of the comparison basis? One typical ground for a change is that a certain comparison basis would give comparison outcomes and conclusions that are undesirable in the light of the values, goals or policies underlying the law. A reason to include the statement that the employee has committed an act of violence could be, for instance, that the law generally seeks to prevent violent conduct.

This suggests that the comparison basis is rooted in the goals or values upheld by the legal system. More specifically, one can see the comparison basis as a means to promote the goals or values of the legal system as much as possible. In other words, if one particular comparison basis is adopted rather than another, then this can be explained in terms of the goals that both promote or detract from. This is comparable to the way in which according to legal coherentism, the theory of positive law should be embedded in a more general theory involving legal values, goals or principles (see for coherentism in the law Hage 2001; Hage *et al.* 2000, pp. 370f.; see Dancy 1985, Chapters 8 and 9 for

coherentism in epistemology). The present model could be extended in line with legal coherentism, namely by choosing the comparison basis such that it promotes the values underlying the system of law as well as possible.

Chapter 4: Related research

In this chapter some of the most relevant related research on case-based reasoning will be discussed and compared to the model proposed in the previous chapter. Three questions will be addressed in particular. First, is it explained which conclusions follow from settled cases, or is there more emphasis on the reasoning patterns along which these conclusions follow? In the proposed model both aspects are captured, but as will be seen the related research tends to address only one of them. Second, is adherence to decisions treated as a kind of rule extraction, or is the case comparison method chosen to this end? In the proposed model the case comparison method is formalised, but in related research both case comparison and rule extraction occur as methods of adhering to decisions. Third, is it acknowledged explicitly that in the law it depends on a contingent choice which case features are relevant in case comparison? In the proposed model this acknowledgement becomes apparent in the introduction of a comparison basis affecting the outcomes of case comparison. As will be seen none of the discussed approaches explicitly deals with the contingency of the set of relevant case features, however.

4.1 Ashley's HYPO: adversarial case-based reasoning

The HYPO system represents one of the most important contributions to case-based reasoning in the field of Artificial Intelligence and Law. It has led to various publications (Rissland and Ashley 1987; Ashley 1989; Ashley 1990; Ashley 1991). The present discussion is based on Ashley's book (Ashley 1990).

HYPO is an implemented model of case-based reasoning, which can generate realistic arguments with cases. The reasoning it models is adversarial, in the sense that it creates arguments and counterarguments to argue from the opposing viewpoints of parties involved in a legal dispute.

The construction of adversarial case-based arguments relies on the comparison of cases, since this comparison determines which settled cases are relevant to cite. For the purpose of case comparison each case is represented by factors, whereby factors can either support a legal conclusion, or its opposite. In HYPO it is given once and for all which relevant factors there are, and the contingency of the set of relevant factors is not explicitly recognised in the model.

The relative strength of arguments is assessed by comparing cases. When the arguments have opposing conclusions, this assessment helps in deciding which argument is the stronger one. The relative strength of arguments is also used strategically to select the best cases to cite for a legal conclusion. Although HYPO can generate arguments citing settled cases and compare the arguments regarding their strength, the model does not tell which conclusions follow from the cited settled cases.

It is shown how HYPO can be reconciled with a multi-step account of reasoning with factors, in contrast with the one-step interpretation that is usually

given of it. Briefly, in this account the roles of HYPO's factors are captured formally in terms of their relevance for the disputed conclusion or its opposite.

4.1.1 Factors, arguments and counterarguments

In HYPO, cases are represented by factors. Factors are stereotypical, legally relevant facts that favour conclusions claimed by one of the two sides involved in a legal dispute, that is, either plaintiff or defendant (p. 26 and pp. 37-38). A factor may favour a conclusion claimed by plaintiff, for instance, in a case where a trade secret is claimed to have been misappropriated. If a factor favours a conclusion claimed by a party or its opponent, then from the perspective of the proposed model this factor supports or attacks this conclusion, respectively. Accordingly, the use of factors in HYPO will be discussed in the terminology introduced in the previous chapter, that is, in terms of support or attack. If in HYPO a factor favours a claim of trade secret misappropriation, for instance, then this will be captured by letting the factor support the conclusion that a trade secret has been misappropriated.

Furthermore, for simplicity the assumption is introduced here that each factor directly supports or attacks the conclusion that is at dispute. In other words, it is presently assumed that no intermediate conclusions are involved, and HYPO is accordingly treated as a one-step model of reasoning with factors. The factor that security measures were adopted to protect certain information, for instance, will directly support the conclusion that the information is a trade secret. No intermediate conclusion will be involved here, such as the conclusion that efforts were made to maintain secrecy. The present one-step account of HYPO is in agreement with the way in which the system is usually interpreted. Within the proposed model a more general account can also be given, however, involving multi-step reasoning with factors. This account will be presented at the end of Section 4.1.

Factors supporting a conclusion are tagged with a plus sign, while factors supporting the opposite conclusion are tagged with a minus sign. A legal conclusion for which one can identify supporting or attacking factors is, for instance, the conclusion that a trade secret has been misappropriated. The following are examples of factors that support or attack this legal conclusion:

- f1+: A competitive advantage was gained with the help of the information (Competitive-Advantage).
- f2-: The information merely pertained to customer business methods (Vertical-Knowledge).
- f3+: Information disclosures were subject to restriction (Disclosures-Subject-To-Restriction).
- f4-: The information was voluntarily disclosed (Secrets-Voluntarily-Disclosed).
- f5+: Security measures were adopted to protect the information (Security-Measures-Adopted).

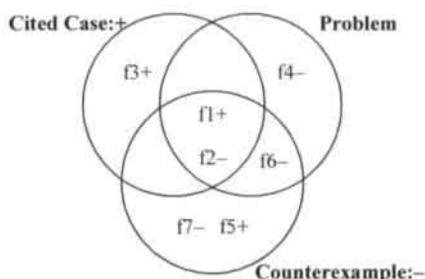
f6-: The information was developed solely by private means (Developed-By-Private-Means).

f7-: The information was disclosed in negotiations (Disclosure-In-Negotiations).

On the basis of the comparison of cases in terms of factors, HYPO can generate special case-based arguments, called *3-ply* arguments (p. 20 and pp. 62-75). These are short disputes involving three turns, or 'plies', of an argument citing a case. Briefly, the first ply involves citing and analogising the case, the second ply is a response to the first one, and the third a response to the response. In the figure below an example is given.

The cases in this figure are represented by circles containing factors. In the 3-ply-argument, a case **Cited Case** is cited first. In this case the conclusion was drawn that a trade secret was misappropriated, a conclusion indicated by the plus sign. It is claimed that the problem case should be decided likewise, and this claim is supported by analogising the cases, that is, pointing out relevant similarities. These relevant similarities are the two shared factors f1+ and f2-. The shared factor f1+ is relevant because it supports the conclusion for which **Cited Case** is cited. The shared factor f2- is a relevant similarity because it shows that the conclusion can follow in spite of its presence as a factor supporting the opposite conclusion. In other words, f2- is a weakness that could be overcome thanks to the presence of f1+.

f1+: Competitive-Advantage
f2-: Vertical-Knowledge
f3+: Disclosures-Subject-To-Restriction
f4-: Secrets-Voluntarily-Disclosed
f5+: Security-Measures-Adopted
f6-: Developed-By-Private-Means
f7+: Disclosure-In-Negotiations



The first citing and analogising ply is responded to in two ways. First, the cited case is distinguished by pointing out two relevant differences, which make the problem case provide less support for the conclusion. These relevant differences are the factors f3+ and f4-, because both make the problem case provide less support for the conclusion that a trade secret was misappropriated.

The figure also shows the second way of responding, namely by citing another case as a counterexample with the opposite conclusion. In the figure **Counterexample** is cited as such, and as indicated by the minus sign no trade secret was found to have been misappropriated there. Note that in addition to the similarities already found between the first cited case and the problem, this case also shares an additional factor f6- with the problem case. In other words, the counterexample has more relevant similarities with the problem case. Such a counterexample is said to be trumping or more on point, in accordance with the

measure of on pointness discussed later on. Briefly, on pointness is determined by the set of factors that a case shares with a problem case. The analogy with more on point counterexamples is considered to be stronger than those with less on point cases, and responses citing them are therefore evaluated as stronger as well.

In the third ply of the argument, the counterexample is in turn responded to by distinguishing it from the problem case. This is done by pointing out a difference which makes the counterexample provide more support for the opposite conclusion, namely the factor $f7-$. This factor helps explain the negative conclusion of the counterexample and it does not apply to the problem case.

The corresponding 3-ply argument could end up looking something like this:

⇒ Point for trade secret misappropriation:

Where: A competitive advantage was gained with the help of the information ($f1+$). Even though: The information merely pertained to customer business methods ($f2-$).

A trade secret was misappropriated.

Cite: **Cited Case**

⇐ Response against trade secret misappropriation:

Cited Case is distinguishable because:

In **Cited Case**, information disclosures were subject to restrictions ($f3+$). Not so in **Problem**. In **Problem**, the information was voluntarily disclosed ($f4-$). Not so in **Cited Case**.

Counterexamples:

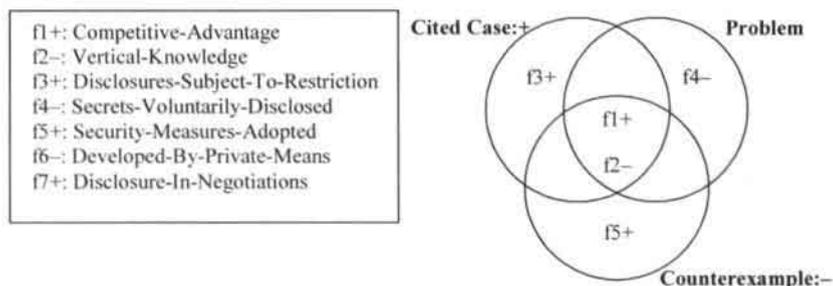
Counterexample is more on point ($f6-$) and held against trade secret misappropriation.

⇒ Response to the Response:

Counterexample is distinguishable because:

In **Counterexample**, the information was disclosed in negotiations ($f7-$). Not so in **Problem**

In the following figure a similar situation is sketched.



There are two differences with the situation above. First, the counterexample has no additional shared factors so it is not more on point. Instead it is merely as on point as the first cited case. Second, the difference $f7-$ is absent now, while the differences $f5+$ and $f4-$ make the problem case provide more support for the opposite conclusion drawn in the counterexample, thereby allowing for a kind of reasoning a fortiori. Accordingly, these differences cannot be used to distinguish the counterexample from the problem case. For this reason the response with **Counterexample** is considered stronger than the original point with **Cited Case**, although both settled cases are equally on point.

As these examples illustrate, the first ply in the 3-ply arguments is always a point involving an analogous cited case. The analogy is supported by pointing out relevant similarities between the cited case and the problem, that is, the factors shared by both cases. The second ply is a response to the first, which can be made in two ways. First, it can be that the cited case is distinguished by pointing out relevant differences, that is, factors that tend to make the cited case provide more support for the conclusion for which it was cited. A second way of responding is by citing counterexamples, that is, cases with an opposite conclusion. There are four kinds of counterexamples (pp. 163-166), among which the trumping counterexamples are most prominent. The third ply is a response to the counterexamples cited in the second ply, namely by distinguishing these counterexamples.

Summarising, the three plies of the 3-ply arguments are as follows (cf. pp. 62-63):

1. Citing a case and drawing an analogy with the problem case.
2. Responding by distinguishing the cited case, or by citing one or more counterexamples.
3. Responding by distinguishing the counterexamples.

In sum, HYPO can produce special arguments and counterarguments on the basis of expert background knowledge on supporting and attacking factors (p. 26). In particular, these arguments involve citing cases and subsequently distinguishing them. In addition to this, cited cases can be responded to by giving more on point counterexamples, where on pointness is measured as the overlap of shared factors. The background knowledge used for generating arguments is represented by tagging factors with a plus sign if they support a disputed conclusion, while factors supporting the opposite are tagged with a minus sign.

In HYPO it is not disputable whether or not a factor supports some conclusion or its opposite, however. More specifically, it is not possible that a factor supports or attacks the statement that another factor supports or attacks a conclusion. Recall that in the model proposed in this book, in contrast, it can be supported or attacked that one statement supports or attacks another, namely by

nesting statements of support or attack. Moreover, recall that these reason statements are in principle as contingent as other information, and that they accordingly play a role in case comparison as well.

At first sight there seems to be another difference regarding factors' roles between HYPO and the model proposed. In the present discussion of HYPO, for simplicity all factors are assumed to plead directly for or against the disputed conclusion. In other words, HYPO has been presented as a one-step model of reasoning with factors, in agreement with the usual interpretation of factors' roles within the system. Accordingly, HYPO's use of factors was not presented in terms of multi-step arguments with intermediate conclusions. Recall that in the model proposed in this book, in contrast, case comparison is done in terms of dialectical arguments, which in general consist of multiple steps with intermediate conclusions. As will be argued later on, this contrast is not as stark as it may seem, however, because HYPO's use of factors can be made compatible with a multi-step model of reasoning.

4.1.2 Dimensions and hypotheticals

An interesting feature of the HYPO system is that it can pose hypothetical fact situations to explore how conclusions are affected by changes in cases (pp. 84-86 and 148-155). This feature is briefly addressed in this section.

In HYPO, a variable called magnitude is associated with each factor, ranging between predefined extremes. For a factor stating that a competitive advantage is gained by misappropriating a trade secret, for instance, this magnitude is the size of advantage thus gained, in terms of saved investments or product development time. Magnitudes can also range over discrete values. A factor stating that there exists an express non-competition agreement, for instance, has as its magnitude the Boolean variable telling whether or not the employee entered into a non-competition covenant.

Factors and their magnitude are represented by Dimensions in HYPO, and the conditions under which a Dimension applies are called the Dimension's prerequisites. These are stated in an abstract case representation language involving Factual Predicates (p. 38, pp. 115-118 and Appendices D and F). Accordingly, to each Dimension there corresponds a list of Factual Predicates that have to be satisfied to make the Dimension apply. The Dimension stating that secrets were voluntarily disclosed, for instance, has among its Factual Predicates that there is a corporate plaintiff (p. 109).

Once the Factual Predicates that are satisfied in a case are known, a list of applicable Dimensions is created automatically (p. 103). From then on the HYPO program works with these Dimensions to perform symbolic comparisons of cases, to create arguments and to assess the relative strengths of these arguments.

HYPO exploits the magnitudes recorded in Dimensions in its 3-ply arguments, for example by distinguishing cases by the magnitude of a shared Dimension. Suppose, for instance, that a case is cited in support of a claim of

trade secret misappropriation. Suppose, moreover, that there is a shared Dimension stating that a competitive advantage is gained by trade secret misappropriation. Then the cited case can be distinguished if more competitive advantage was gained there than in the problem case, for instance because more product development time was saved.

Another application of magnitude is in the definition of a special kind of counterexample, namely the boundary counterexamples that can be cited in the second ply of the 3-ply arguments (pp. 164-165). Briefly, a boundary counterexample is a case where some factor's magnitude assumed a value that made the case extremely favourable along the corresponding Dimension, but where the opposite of the factor's conclusion was nevertheless drawn. In effect, a boundary counterexample tends to show that such a factor is not very significant.

Dimensions and their magnitudes can also play a role in another key feature of the HYPO system, namely that it can pose hypothetical variations of cases to explore how changes affect conclusions. Typically, the case that is to be modified is the problem case at hand (pp. 84 and 166), but other choices seem possible too (pp. 148-154).

Two typical ways of changing a problem are the following. One is by extending the problem with near miss Dimensions. Informally, these are Dimensions that almost apply to the problem, and they are defined as follows. To determine the magnitude of a factor associated with a Dimension, certain factual case information is required. If a factor states that a competitive advantage is gained by misappropriating a trade secret, for instance, the case must contain information concerning the size of the advantage gained in this way, in terms of saved investments or product development time. In the absence of such information the Dimension does not apply, because it is one of the prerequisites of the Dimension (pp. 38 and 40). If the absence of information on magnitude is the only reason why a Dimension does not apply, then the Dimension is said to be a near miss. In other words, the prerequisites of a near miss Dimension are all satisfied, except that which determines the magnitude of the associated factor.

A second way of hypothetically changing the problem is by taking an applicable Dimension, and change its magnitude. If a competitive advantage is gained by misappropriating a trade secret, for example, one could change the size of this competitive advantage in terms of saved investments.

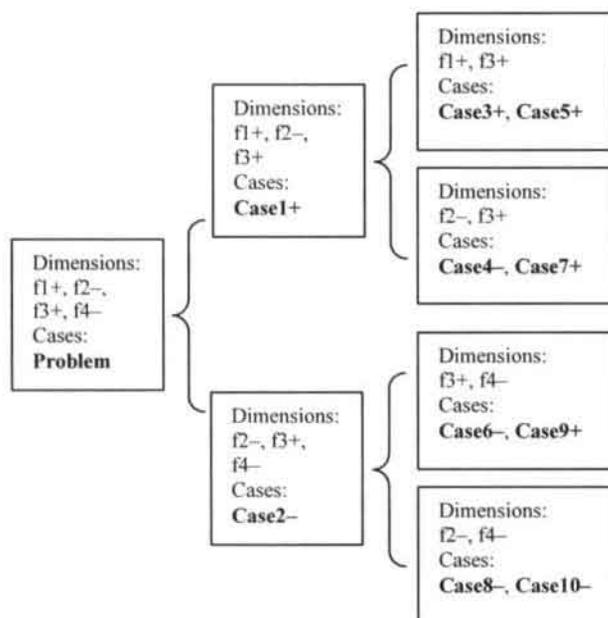
The model proposed in this book works with factors only, and not with dimensions. As a result, the interesting arguments with dimensions discussed above are not possible. In particular, there are no in-built heuristics to hypothetically change a problem case, and to explore the effect on conclusions of such a change. Other possibilities of dimensions cannot be exploited either, such as those discussed recently by Bench-Capon and Rissland (2001, p. 51).

In HYPO the applicable Dimensions can be derived from information about the Factual Predicates, which has to be entered by hand under assistance of a Case Editor (p. 102). On the basis of this information HYPO automatically infers whether or not Dimensions apply or are near misses. Such inferences are not complex and do not involve, for instance, chains of multiple reasoning steps. This is in contrast with the model proposed in this book, where the basic factors that apply are derived in conflict-free inference chains (Definition 4) that may involve multiple steps of reasoning.

4.1.3 On pointness, case comparison and argument evaluation

In HYPO, analogy between cases is defined by a measure called on pointness (pp. 128f.). Roughly, on pointness is the overlap of factors that a case shares with a problem case. In the first figure above, for instance, **Cited Case** shared the factors $f1+$ and $f2-$ with **Problem**. Two cases are said to be as on point if they share the same factors with a problem. One case is said to be more on point than a second, if the shared factors of the second are a proper subset of those of the first. In the first figure above, for instance, **Counterexample** was more on point than **Cited Case**, because it shared an additional factor $f7-$ with **Problem**. A case is most on point if there are no more on point cases.

Given a problem case, the decided cases are ordered with respect to their measure of on pointness. This (partial) ordering can be visualised by a tree called a Claim Lattice, of which there is an example in the following figure.

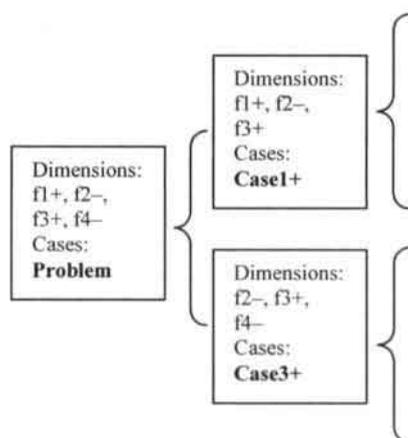


All nodes of the tree list a set of shared Dimensions and the corresponding set of cases. The problem case is at the root of the tree, together with all equally on

point cases that are found. The tree is meant to suggest that the cases in each node are more on point to the problem case than those in the node's children.

The most on point cases are those closest to the root. In the figure, **Case1+** is a most on point case where some conclusion was drawn, while **Case2-** is one with the opposite conclusion. Since both cases are most on point there are no trumping counterexamples to them when cited, since by definition a trumping counterexample is a more on point case. Accordingly, strong arguments for both opposing conclusions can be made here.

However, it is also possible that all most on point cases have the same conclusion, as in the following figure. In this figure, all most on point cases have a positive result, as indicated by the plus signs. None of them can be trumped with counterexamples. Accordingly, in this situation the argument for the conclusion is apparently stronger than that for its opposite.



The ordering of cases by on pointness in HYPO is comparable to the ordering by dialectical support used in the model proposed in this book, because both orderings are used to determine which settled cases are relevant for a problem case.

In addition to this, HYPO exploits the on pointness ordering to strategically select the best cases to cite. A more on point case is strategically better than a less on point one, because it is less likely to be trumped. Moreover, if two cases with the same conclusion are equally on point, then HYPO can evaluate their relative strengths on other grounds, for instance by preferring precedents that cannot be distinguished (cf. pp. 32-33). More specifically, if two equally on point precedents are cited for the same conclusion while one and only one of them is not distinguishable, then the point with the latter precedent is considered stronger.

From the perspective of the proposed model it also makes perfect sense to prefer settled cases that cannot be distinguished. The reason for this is that if a settled case is not distinguishable from a problem case, then the problem case

provides at least as much dialectical support for the conclusion as the settled case. This holds, for instance, for the counterexample used in the second figure above. In other words, a settled case that is not distinguishable would allow for reasoning a fortiori, and it would be relevant as a precedent in the proposed model as well. One can therefore say that HYPO's preference for settled cases without distinctions reveals a parallel with the proposed model. The parallel is that if in HYPO a settled case cannot be distinguished from a problem case, then according to the proposed model the problem case provides at least equal dialectical support for the conclusion.

There is an important difference between HYPO's on pointness as a criterion to compare cases, and the way in which case comparison is treated in the model proposed in this book. In the proposed model it is in principle contingent which factors are relevant for case comparison. More specifically, a comparison basis is introduced as a division between factors and non-factors, and the outcomes of case comparison are stated relative to this comparison basis. In HYPO the factors are a kind of expert knowledge, which comes from legal authorities such as judges and scholars. Accordingly, it is assumed given once and for all which relevant factors there are (cf. p. 158, where the set of factors is formally introduced). As a consequence, in HYPO the on pointness ordering is absolute rather than relative to the set of factors chosen.

In HYPO there is no way of determining whether a case-based argument actually makes its conclusion follow. This is in agreement with the perspective taken in designing HYPO, according to which problems in the legal domain hardly ever have one right answer (cf. pp. 28-29 and p. 203). As Ashley puts it (p. 28): "The law has no analytic model or procedure for resolving the competing factors, no weighting scheme or function that yields the "right answer" through some deductive or mathematical process." In line with this view, the on pointness metric for case comparison allows settled cases to become relevant even if they can be distinguished. Within the model proposed in this book, such distinguishable cases are not relevant as precedents. The reason for this is that distinguishable cases cannot be followed in the proposed model, since distinctions tend to make the settled case provide more support for the conclusion than the problem case. For instance, let a settled case be cited for a conclusion supported by a factor f_3^+ , which applies to the settled case and not to the problem case. Then in the proposed model this cited case is not a precedent, since f_3^+ is a distinction. By thus banning distinguishable cases as precedents, the proposed model can indeed account for the conclusions that follow by case comparison. At this point it remains an open question, though, whether this ban accurately reflects legal practice and the doctrine of *stare decisis*.

As said, in HYPO distinguishable settled cases can be cited in an argument, in contrast to the proposed model. This is because on pointness is defined in terms of the shared factors, while the differences between the cases are ignored. In the first figure above, for instance, there is a distinguishable counterexample that may be a most on point case as well (depending on the other settled cases).

That distinguishable cases can be cited is in accordance with the view that “A theory of argumentation deals with (...) persuasion, not logical demonstration (...)” (p. 6). In line with this view, the main purpose of HYPO is to produce realistic arguments with cases, and not to define the conclusions that follow from them. In the proposed model, in contrast, the main purpose is to determine which conclusions follow on the basis of a comparison with settled cases, while its secondary aim is to account for typical patterns of reasoning with cases, such as analogising and distinguishing.

4.1.4 Reasoning with factors in HYPO: one-step or multi-step?

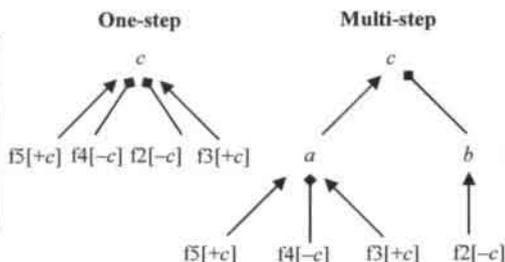
There appears to be widespread agreement in AI & Law that reasoning with factors in HYPO is one-step reasoning (Prakken and Sartor 1998, p. 249 and 267; Hage 1997, p. 186; cf. Aleven 1997, p. 211 on HYPO’s lacking of a factor hierarchy; Bench-Capon 1999, p. 40). From the perspective of the proposed model this one-step character of HYPO would mean that each factor directly supports the conclusion or its opposite.

Does HYPO’s architect share this opinion himself (Ashley 1990)? It is not easy to give a clear answer to this question, because HYPO’s presentation is not very transparent regarding the roles that it attributes to its factors. In HYPO these roles are presented as factors’ tendencies to strengthen or weaken an argument in favour of a legal claim (Ashley 1990, p. 26). From the perspective of the proposed model (see Section 3.6) these tendencies would mean that HYPO’s factors contribute to, or detract from, the conclusion that is claimed to hold. But do the contributing and detracting effects of HYPO’s factors necessarily imply a one-step reasoning model?

In the following it will be argued that the account of factors in HYPO is also compatible with multi-step reasoning, and therefore need not be taken to represent one-step reasoning only. The argument runs briefly as follows. The key idea is that the contributing and detracting effects of HYPO’s factors can be captured formally in terms of relevance for the conclusion or its opposite, as in the model proposed in the previous chapter (cf. Definition 6, Section 3.6). These relevances can occur as the result of one-step reasoning, in which factors directly support or attack the conclusion. This one-step account is in agreement with the way in which HYPO is usually interpreted in AI & Law. However, the same factors’ relevances can also occur as the result of multi-step reasoning.

Consider the example of the following figure.

c: Info-Trade-Secret
a: Efforts-To-Maintain-Secrecy
b: Info-Known-Or-Available
f2: Vertical-Knowledge
f3: Disclosures-Subject-To-Restriction
f4: Secrets-Voluntarily-Disclosed
f5: Security-Measures



In this figure there are four of HYPO's factors, each of which contributes to, or detracts from the conclusion (*c*) that certain information is a trade secret. The conclusion is contributed to by the factor (*f5*) that security measures were taken, and by the factor (*f3*) that disclosures were subject to restriction. The conclusion is detracted from by the factor (*f2*) that the misappropriated information is vertical knowledge, that is, the information pertains to customer business methods. Another detracting factor is the factor (*f4*) that the information was voluntarily disclosed. The contributing or detracting effect of each HYPO factor is made explicit in terms of the factor's relevance for the conclusion or its opposite, respectively (cf. Definition 6, Section 3.6). Relevance for the conclusion (*c*) is indicated by placing it between square brackets, together with a plus sign (i.e. [+*c*]). Likewise, relevance for the opposite conclusion is indicated with a minus sign (i.e. [-*c*]).

The effect of each HYPO factor can be accounted for by the one-step dialectical argument on the left, in which each factor directly supports or attacks the conclusion (*c*) that the information is a trade secret. The factor (*f5*) that security measures were taken, for instance, directly supports this conclusion and is accordingly also relevant for it. Likewise, the factor (*f2*) that the information is vertical knowledge attacks the conclusion, and is accordingly also relevant for the opposite conclusion. This one-step account of HYPO's factors corresponds to the way in which the reasoning in HYPO is usually interpreted.

The effects of the factors can also be accounted for using the multi-step dialectical argument on the right, however. Formally this becomes apparent by the fact that on the right the HYPO factors have the same relevances with respect to the conclusion as on the left. That the factor (*f5*) that security measures were taken contributes to the conclusion, for instance, can be explained by two support relations in the dialectical argument on the right. First, the intermediate conclusion (*a*) that efforts were made to maintain secrecy is supported by the factor (*f5*) that security measures were taken. Second, this intermediate conclusion (*a*) supports the conclusion (*c*) that the information is a trade secret. In a similar way it can be explained that the conclusion (*c*) is detracted from by the factor (*f2*) that the misappropriated information is vertical knowledge. The introduction of an intermediate conclusion (such as *a*) is, by the way, an example of the use of issues or abstract factors in CATO (Aleven 1997, pp. 50-53).

As this example shows, the contributing and detracting effects of HYPO's factors can be explained in terms of relevance for a conclusion or its opposite, respectively. Accordingly, HYPO's case comparison with factors can alternatively be accounted for with the help of the relation of relevance (Definition 6). These relevances can occur as the result of one-step reasoning, in accordance with the way HYPO is usually interpreted. However, the same relevances can also occur as the result of arguments involving multiple steps. As a result, HYPO's account of case comparison is compatible with multi-step reasoning as well.

Why, then, is HYPO generally associated with one-step reasoning? A possible answer to this question is that being a reason for or against a conclusion is the simplest way of being relevant as a factor. More precisely, a reason for a conclusion is the simplest example of a factor that is relevant for the conclusion, and a reason against the conclusion is the simplest example of a factor that is relevant for the opposite conclusion. Moreover, if reasoning with factors is treated as one-step reasoning, then only factors are involved that directly support or attack a conclusion as reasons. As a result, an account in terms of one-step reasoning is the simplest possible way to accommodate the contributing and detracting effects of factors in HYPO. As shown this does not necessarily mean, however, that within HYPO there would be no place for a multi-step account of reasoning with factors. In fact, the model proposed in the previous chapter shows a way to give such a multi-step account its place within HYPO.

4.2 Aleven's CATO: argumentation in case comparison

The CATO program uses a model of case-based argumentation in an instructional setting, to teach students basic skills of arguing with cases. It has been the subject of various publications (Aleven 1997; Aleven 1996; Aleven and Ashley 1997). Recently it was discussed from a jurisprudential perspective (Ashley 2002). The present discussion is based on Aleven's dissertation (1997, reviewed by Roth 1999a).

CATO has made contributions both in legal education and in legal case-based reasoning. Its main contribution to legal education was that it provided an instructional environment communicating a model of case-based reasoning, while the model itself was its main contribution to legal case-based reasoning.

The CATO model of case-based reasoning relies on a Factor Hierarchy, a body of background knowledge for making arguments with cases. Briefly, this knowledge is used to make special arguments on the significance of surface-level distinctions between cases, by interpreting the distinctions in the abstract. In the course of producing its special arguments, the program can strategically focus on certain interpretations (Focal Abstractions), which are selected on the basis of special criteria (Heuristic Criteria), such as the strength of an argument in the light of possible counterarguments. As it turns out, these special

arguments can be reconstructed within the model proposed in this book, namely as the introduction of a shift in the comparison basis.

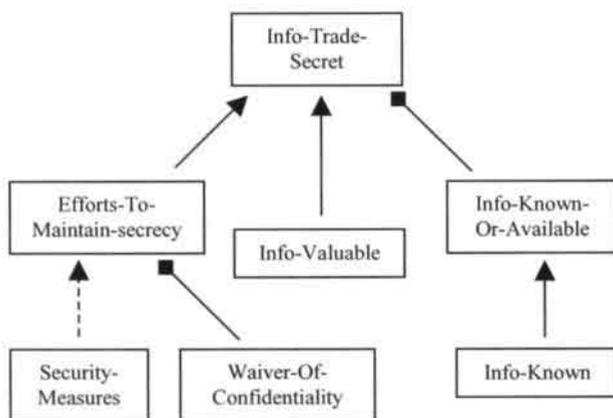
4.2.1 *The Factor Hierarchy*

The Factor Hierarchy is a body of case-independent background knowledge about how the relevant factors in some domain relate to each other. More specifically, it tells which abstract factors are supported as conclusions by other factors (pp. 45-46). The Factor Hierarchy is used to identify issues (pp. 50-53) in cases, such as whether certain information is a trade secret. More importantly, is it also used to reason about the significance of distinctions between cases, by inferring abstract interpretations from these distinctions. These lead to a special kind of argument in case comparison called downplaying and emphasizing, which will be discussed in more detail later on.

The Factor Hierarchy consists of links between factors, which represent a relation of support (p. 23). Links can be labelled positive or negative (p. 22), depending on whether the linked factors stand for conclusions favouring the same or opposite sides, respectively. The negative links will be called attack links, to facilitate the comparison later on with the model presented in the previous chapter.

Links can be weak or strong, indicating the level of support they represent. An example is in the following figure.

Factor Hierarchy (part)



In this figure a part of the Factor Hierarchy is shown. The factors are in rectangles, and the links are shown as arrows. Normal arrows represent support, and arrows ending in a solid square stand for attack. Nearly all links in the figure are strong, except the one represented by a dashed arrow. As explained

later, the relative strength of links is used to block inferences with weak links, in favour of inferences with strong links.

Factors at the bottom of the Factor Hierarchy are called base-level factors. An example shown in the figure above is the factor Security-Measures, stating that security measures were taken to protect the information that is claimed to have been misappropriated as a trade secret. At the top of the Factor Hierarchy one finds Legal Issues, the main issues involved in judicial reasoning in some domain (p. 46). The example in the figure above is the factor Info-Trade-Secret, which states that the information involved is a trade secret. In the middle of the Factor Hierarchy one finds Intermediate Legal Concerns, such as the factor Info-Known-Or-Available in the figure above. All factors in the Factor Hierarchy that are not base-level are called abstract factors.

Cases are represented in terms of base-level factors (pp. 41f. and 47f.), which have to be identified by hand from the textual descriptions of case facts (pp. 42-43). In other words, in CATO the applying factors are not derived from the non-factors that case descriptions typically present. This is in contrast with the model proposed in this book, where the applicable basic factors are derived from non-factors in combination with background knowledge.

The ultimate possible conclusions in the domain modelled are that a party is, or is not, liable for misappropriating a trade secret (p. 46), but these conclusions are not included in the Factor Hierarchy. The reason for this is that the Factor Hierarchy is not used to derive these ultimate conclusions. Instead the Factor Hierarchy mainly serves the purpose of interpreting base-level factors in terms of the abstract factors that they support. These interpretations are involved in the special arguments it can make on the significance of distinctions, namely the downplaying and emphasizing moves discussed later on.

As said there are weak and strong links in the Factor Hierarchy. Weak links stand for weak support for a conclusion, which can be blocked if there is strong support for an opposite conclusion. An example is in the figure above. As indicated by the dashed arrow, the factor Security-Measures provides weak support for the conclusion stated by factor Efforts-To-Maintain-Secrecy. The factor Waiver-Of-Confidentiality provides strong support for the opposite conclusion, though, as indicated by a arrow ending in a solid square. Accordingly, if the factors Security-Measures and Waiver-Of-Confidentiality both apply to a case, then the support provided by the former is blocked. As a result, the conclusion would not be supported that efforts were made to maintain secrecy.

In sum, the Factor Hierarchy is a body of background knowledge consisting of support relations between factors. These relations enable CATO to interpret the base-level factors in terms of the abstract factors they support, through inferences that can comprise multiple steps. If HYPO is presented as a one-step account of reasoning with factors, then CATO's multi-step arguments can be seen as an innovation over HYPO (cf. p. 12).

In other words, in CATO case comparison is done not just in terms of base-level factors, but also at the level of intermediate conclusions in the Factor Hierarchy. In this respect CATO resembles the model proposed in this book, where the entangled factor hierarchy contains all conclusions that can be supported by basic factors. A difference is, however, that in CATO the links between factors are assumed as self-evident background knowledge, 'grounded in the common sense of the legal claim' (p. 46). Accordingly, in CATO it cannot be supported or attacked that a factor supports a conclusion or its opposite.

In the model proposed in this book, in contrast, one can support or attack that a factor supports or attacks a conclusion, as becomes apparent, for instance, in the entanglement of dialectical arguments and of the factor hierarchy. In particular, this possibility of supporting or attacking reason statements is exploited in the account of case comparison as well.

Another difference between the two models is, that in CATO some conflicts among factors can be resolved. This resolution is done by blocking an inference with a weak link with an inference with a strong link. In the model proposed in this book, however, conflicts cannot be resolved in this way. The reason for this is, that this kind of priority information cannot be expressed in the case representation language. Instead conflicts are only resolved on the basis of a comparison with settled cases, using the comparison outcomes to decide which settled cases are relevant as precedents. In accordance with this, while there is no means of expressing priority information within the proposed model, the outcomes of case comparison can be expressed, namely with the help of dedicated case denotations.

4.2.2 Downplaying and emphasizing distinctions

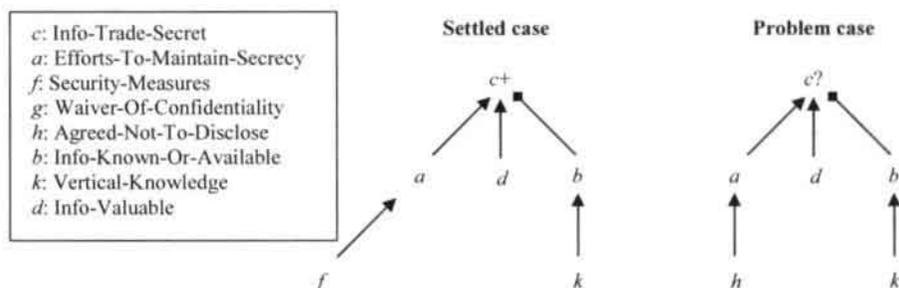
In CATO case comparison involves arguments on the issue whether a cited decided case is analogous to the problem case. This kind of argument is modelled by four Basic Argument Moves (p. 58): Analogizing, distinguishing, downplaying a distinction and emphasizing the distinction. Analogizing and distinguishing is done in essentially the same way as in HYPO (see the previous section). Briefly, analogizing is done by pointing out shared base-level factors as relevant similarities, while distinguishing involves unshared factors as relevant differences. Moreover, a difference is only relevant if it tends to make the decided case stronger for its conclusion than the problem. Such a relevant difference is called a distinction (p. 60).

Once a distinction has been found, the other two Basic Argument Moves can come into play as special arguments on its significance, namely downplaying the distinction and emphasizing it. To produce such arguments on the significance of a distinction, CATO can infer abstract interpretations of it with the help of the links encoded in its Factor Hierarchy. In making these arguments, it strategically focuses on certain abstract interpretations – Focal Abstractions – of the distinction.

Examples of arguments on a distinction's significance are in the figures below. The main issue in these examples is whether certain information represents a trade secret. The following abbreviations are used throughout the examples.

- c*: The information represents a trade secret (Info-Trade-Secret).
- a*: Efforts were made to maintain secrecy (Efforts-To-Maintain-Secrecy).
- f*: Security measures were taken to maintain secrecy (Security-Measures).
- g*: The confidentiality of the information was given up (Waiver-of-Confidentiality).
- h*: There was an agreement not to disclose the information (Agreed-Not-To-Disclose).
- b*: The information is known or available elsewhere (Info-Known-Or-Available).
- k*: The information merely pertains to customer business methods (Vertical-Knowledge).
- d*: The information is valuable (Info-Valuable).

There are two ways of downplaying a distinction, viz. by drawing an abstract parallel and by showing an opposite interpretation. An example of the first way of downplaying is in the next figure.



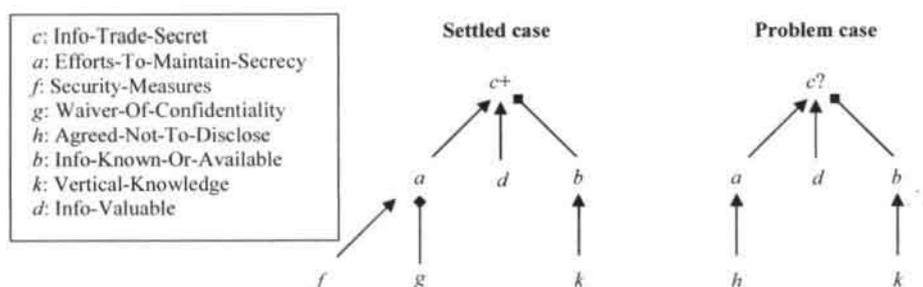
In this figure there is a settled case on the left where the conclusion *c* was drawn that the information was a trade secret, as indicated by the plus sign. On the right there is a problem case where *c* is an issue, as indicated by the question mark. As is evident from the figure, there are two differences between the cases. First, the factor *f* that security measures were taken applies to the settled case and not to the problem case. Second, the factor *h* that there was a nondisclosure agreement applies to the problem case and not to the settled case. Moreover, *f* is a distinction because it supports the intermediate conclusion *a* that efforts were made to maintain secrecy, so that *f* tends to make the settled case stronger for conclusion *c* than the problem case.

However, in the problem case *a* is supported as well, namely by the factor *h* that there was a nondisclosure agreement. In other words, both *f* and *h* can be

interpreted in the abstract as evidence for the intermediate conclusion a that efforts were made to maintain secrecy. (This holds true even if a were attacked in the problem case, provided that this attack is blocked in some way, for instance by h .) Factors such as h provide alternative support for an abstract interpretation of a distinction, in the case where it does not apply. Such alternative supporting factors are called *similar* to the distinction.

By focusing on an abstract interpretation such as a – a focal abstraction – a distinction like f appears to be not significant if a similar factor like h applies to the other case and supports the abstraction there. In other words, the result of downplaying is that after restating the distinction in terms of an abstraction, the abstraction is a similarity between the cases.

The second way of downplaying a distinction is by showing an interpretation opposite to that of the distinction. An example is in the following figure.



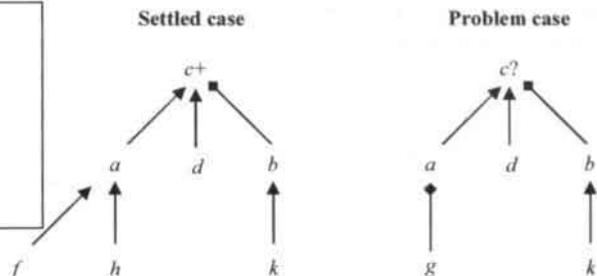
Again f is a distinction which can be interpreted in the abstract as evidence for the intermediate conclusion a that efforts were made to maintain secrecy. There is also evidence for the opposite of a that no such efforts were made, however, namely the factor g that there is a waiver of confidentiality (provided that its support is not blocked by f). In other words, the abstract interpretation a of f is in some sense undercut by a factor g with an opposite interpretation. Such factors are therefore called *undercutting* factors (p. 73)¹⁴.

By focusing on a focal abstraction like a , a distinction like f appears to be not significant if an undercutting factor applies to the same case as the distinction, and supports the opposite interpretation there. Accordingly, the result is that the abstract interpretation is not a distinction between the cases.

A distinction can also be emphasized, namely by drawing an abstract contrast. An example is in the following figure.

¹⁴ Note that this notion of undercutting is different from Pollock's (1995, pp. 41 and 86).

- e: Info-Trade-Secret
- a: Efforts-To-Maintain-Secrecy
- f: Security-Measures
- g: Waiver-Of-Confidentiality
- h: Agreed-Not-To-Disclose
- b: Info-Known-Or-Available
- k: Vertical-Knowledge
- d: Info-Valuable



Again the factor *f* that security measures were taken is a distinction between the cases, and again it is interpreted as evidence for focal abstraction *a* that efforts were made to maintain secrecy. As the figure shows, there is no support for *a* in the problem case. In other words, there is a contrast with regard to focal abstraction *a*. The contrast is even sharpened by the factor *h* that there was a nondisclosure agreement, because this factor provides additional support for *a* in the settled case. A factor like *h* is called a *corroborating* factor in the same case as that where the distinction applies. The contrast is also sharpened by the factor *g* that there was a waiver of confidentiality, because this factor provides support for the opposite of *a* in the problem case. Such a factor is called a *contrasting* factor in the other case than that where the distinction applies.

In sum, there not only exists a contrast with respect to distinction *f*, but also one at the more abstract level of *a*. As a result, if *f* is interpreted as evidence for focal abstraction *a*, then it appears to be a highly significant one. In other words, one may say that the emphasizing move draws attention to the fact that even after restating the distinction in terms of an abstraction, the abstraction is not a similarity between the cases.

As shown in the previous chapter (Section 3.10), such downplaying and emphasising arguments can also be accommodated in the model proposed in this book. As it turned out, one can regard arguments on the importance of a distinction as a local shift upwards of the comparison basis, to the effect that the comparison is done in terms of some more abstract basic factor. The situation after the shift is then compared to the situation before.

In addition to this, it was shown how one can downplay or emphasise a *similarity* within the proposed model. As it turned out, such arguments involve a local shift downwards of the comparison basis. After the shift the situation is compared to that before.

Having said that it must also be noted, however, that in the proposed model there are no heuristics (heuristic criteria, see below) to guide a strategic choice among several possible ways of downplaying or emphasising. There is more than one way to downplay a distinction, for instance, if the distinction supports more than one relevant abstract factor. It must also be noted that within the

formalism proper, it is not possible to reason with and about comparison bases, since the comparison basis has to be fixed in advance.

Nevertheless, CATO's special argument moves on the significance of a distinction can be accommodated semi-formally within the proposed model. In addition to this, comparable arguments involving a similarity can be accounted for. This demonstrates an important contribution of the research proposed in this book, namely that it is made explicit that in the law cases are compared relative to a contingent choice of factors, a comparison basis. However, within the formalism proper it is not yet possible to reason about the issue how one should choose the comparison basis.

4.2.3 *Heuristic criteria*

Not every abstract factor can be used as a focal abstraction, and to select among potential candidates CATO makes use of heuristic criteria (p. 68, p. 71 and p. 73). Broadly stated, the heuristic criteria help to select the strategically best interpretations on which to focus, given the current fact situation and the viewpoint of the arguer on whose behalf the reasoning is performed. The criteria are discussed in more detail in the following.

A general criterion for both downplaying and distinguishing is, that a focal abstraction must be linked directly to base-level factors (p. 69). This criterion excludes most Legal Issues appearing at the top of the Factor Hierarchy. The intuitive reason for the criterion is, that if a contrast or parallel is drawn that is too broad, then the argument is not convincing (p. 69).

Another general criterion is that the interpretation employed in an argument move must be supported in the case. Intuitively, an interpretation is supported if there is at least a minimal level of evidence for it (p. 65). More formally, a conclusion is supported in a case if two conditions are satisfied, which can be stated briefly as follows (cf. p. 66). First, the conclusion must be connected to some base-level factor in the case, following a path made up from links in the Factor Hierarchy. Second, the support for each conclusion in the path must not be blocked by factors supporting the opposite. Accordingly, in the example above of downplaying with an abstract parallel, there must not be a factor blocking the support for the factor *a* that efforts were made to maintain secrecy. Likewise, in the example above of downplaying with an opposite interpretation, there must not be a factor blocking the support for the opposite of *a* that no efforts were made to maintain secrecy. In sum, by demanding that the interpretation employed is supported in the case, one makes sure that the support for this interpretation cannot be blocked. This helps avoid questionable interpretations which can be responded to all too easily (pp. 76-77).

A heuristic criterion for focal abstractions used for emphasizing distinctions is, that either the focal abstraction is not supported in the case where the distinction does not apply, or else its opposite must be supported there. Moreover, in the latter situation there must be a closed world assumption to the effect that in the absence of support, the focal abstraction does not apply. In the

example of distinguishing above, the focal abstraction *a* was not supported in the problem case, while its opposite was. The reason for employing this criterion is to prevent that while at the surface level a contrast between the cases exists, there would be a parallel at the level of the focal abstraction.

Another criterion for focal abstractions used for emphasizing distinctions is, that there must also be a contrast among factors linked with them higher up in the Factor Hierarchy (except possibly among the Legal Issues at the top, such as conclusion *c* in the examples above). More specifically, a focal abstraction must be linked to more abstract factors higher up in the Factor Hierarchy which are all supported, and not blocked by factors supporting the opposite. In addition to this, no such more abstract factor is to be supported in the case where the distinction does not apply. This is to prevent that while there is a contrast at the level of the focal abstraction, there would be an abstract parallel higher up in the factor hierarchy.

More generally, the reason for these criteria for emphasizing is, that they help avoid interpretations that can easily be questioned, for example by making a downplaying move (cf. pp. 77-78). This also explains why the criteria do not hold for the Legal Issues connected with a distinction: According to the first general heuristic criterion discussed above, most of these Legal Issues cannot be used for drawing contrasts or parallels anyway, because these would be too broad to be convincing.

It may happen that there is more than one focal abstraction satisfying these heuristic criteria. To select among these possibilities, CATO uses three heuristic policies (pp. 69-70). A central notion thereby is coverage, which is measured in terms of similar or undercutting factors when downplaying, and in terms of corroborating and contrasting factors when emphasizing. The first policy is to prefer more abstract factors to more specific ones, as long as this increases coverage, and the second policy is to drop a more abstract factor if it does not have more coverage. The third policy is a rhetorical one. According to this policy, short and compact arguments are preferred, and repetitive arguments are avoided.

In the model proposed in this book such heuristic criteria are not formulated. Accordingly, in the proposed model there are no strategic guidelines for how the comparison basis should be shifted upwards to downplay or emphasise a distinction. More specifically, the proposed model does not tell along which path of support links the comparison basis should be shifted upwards, nor how high up to ascend along the path.

Another difference in this connection between CATO and the model proposed in this book is, that in the latter the conclusions are defined that follow on the basis of a collection of settled cases. In line with this, the outcomes of case comparison are defined 'from a bird's eye view', taking all relevant supporting and attacking statements into account. In this way a distinguishable case can never be a precedent, because distinctions prevent a problem case from

providing at least equal dialectical support for the conclusion. As a result, for the purpose of defining conclusions the proposed model leaves no room for a strategic choice among arguments in case comparison, since for this purpose all possibilities for arguing have already been taken into consideration anyway.

Having said that it must also be remarked, however, that the proposed model does not address a number of both practical and theoretical questions here. It is not clear, for instance, whether and how the model's definitions of comparison outcomes and precedents can be implemented efficiently in a working computer program. A more theoretical question in this connection is, whether the definition of precedents is perhaps too restrictive by banning distinguishable cases as candidate precedents. As remarked earlier in the discussion of HYPO, it remains an open question whether this ban accurately reflects legal practice and the doctrine of *stare decisis*.

4.3 Prakken and Sartor's formal dialogue game

Prakken and Sartor (1998) deal with case-based reasoning as a kind of dialectical argumentation involving arguments and counterarguments. In particular, their model allows for argument moves typically found in case comparison, such as analogising and distinguishing cases. This kind of reasoning with cases is treated as a kind of premise introduction (p. 260) or theory construction (cf. Prakken 2000, p. 52). Analogising, for instance, is captured by introducing a new rule into the dispute, on the basis of a rule selected from a settled case (broadening, see below).

In the formalisation, Prakken and Sartor make use of the argument-based system published in Prakken and Sartor (1996) and Prakken and Sartor (1997a). Their account of case-based reasoning was published earlier at the ICAIL Conference (Prakken and Sartor 1997b). The present discussion is based on an elaborate paper in Artificial Intelligence and Law (Prakken and Sartor 1998).

In Prakken and Sartor's model, the smallest building blocks of argumentation are rules. Rules that are subject to exceptions are called defeasible and are represented by \Rightarrow . Rules are used to express that factors support conclusions or their opposites, and to be able to make statements about them they are preceded by a name.¹⁵ If a factor f supports conclusion c , for instance, a rule named r can be used as to represent this as follows: $r: f \Rightarrow c$. Different factors can be combined into one conjunctive reason, as for instance in the rule $r_1: a \wedge b \Rightarrow c$.

The language has weak (\sim) and classical (\neg) negation to model argument defeat by undercutting and rebutting, respectively (see below). Informally, a weakly negated literal like $\sim b$ is to be read as 'there is no evidence that b '. If a rule has such a weakly negated literal among its conditions, then the corresponding condition can be ignored for the purpose of applying the rule, unless

¹⁵ In Verheij 2000a, pp. 45-48 and 58-59 some disadvantages of this naming technique are mentioned. Verheij proposes a technique that makes use of nesting, the method that was adopted in the present work as well.

the factor whose presence is weakly negated is found to be derivable after all. The application of the rule $r_2: \sim b \Rightarrow c$, for instance, can only be blocked if b can be derived as a conclusion of some argument.

Rules are ordered by a priority relation, denoted \prec . The priority of one rule over another is expressed by using the corresponding rule names. That rule r_3 is to be preferred over r_2 , for instance, is expressed by $r_2 \prec r_3$. It is a special feature of the system that the priority relations are debatable, and can be established on the basis of other rules. Thus the rule $r_5: reason \Rightarrow r_2 \prec r_3$, for instance, says that on the basis of the information expressed by *reason*, the rule r_3 is to be preferred over the rule r_2 .

Facts, rules and information on the rule ordering are the premises from which the reasoning starts, and together they make up what is called an ordered theory.¹⁶ The consequences of an ordered theory are determined in a dialogue game involving argument moves.

Arguments are formed by chaining (facts and) rules, ignoring weakly negated conditions. By chaining a fact f and the rules $f \Rightarrow a$ and $a \wedge \sim b \Rightarrow c$, for instance, one gets an argument with conclusions (f and) a and c . Arguments can defeat other arguments, and this can happen in three ways. The first is by *undercutting*, which happens if an argument has a conclusion whose weak negation occurs as a condition of a rule in another argument. An argument with a rule $r_2: a \Rightarrow b$, for instance, undercuts an argument with a rule $r_4: \sim b \Rightarrow c$.

The second way of defeating is by *excluding*, which happens if the conclusion of an argument holds that a rule in another argument is not applicable. If a rule r is not applicable, then this is formalised as $\neg appl(r)$. Accordingly, an argument with a rule $r_2: a \Rightarrow \neg appl(r_3)$, for instance, excludes an argument with a rule r_3 .

The third way to defeat an argument is by *rebutting*. An argument rebuts another if the two arguments contain rules whose conclusions are in head-to-head conflict with each other, while the rule in the second argument does not have priority over the one in the first. An argument with a rule $r_6: b \Rightarrow \sim c$, for instance, rebuts an argument with a rule $r_7: d \Rightarrow c$, provided that r_7 does not have priority over r_6 .

Arguments can defeat each other. If no priority relation holds between two rules, for instance, then it is possible that arguments using these rules rebut each other. If one argument defeats another but not vice versa, though, then the first strictly defeats the second. The relation of strict defeat is used in the definition of the dialogue game, to reflect the dialectical asymmetry between the proponent of a claim and the opponent. More specifically, the proponent's arguments are required to be strictly defeating, while those of the opponent may be just defeating (p. 257).

¹⁶ Formally, facts are treated as (strict) rules with empty conditions (pp. 254-255). For the present discussion this is irrelevant.

The final assessment of the status of a claim is done in the form of a dialogue game between two players, viz. the proponent of the claim and an opponent. Briefly, this game determines whether a conclusion can be upheld against all possible defeating arguments, and it tells whether or not the conclusion is justified on the basis of a given ordered theory. A conclusion is justified if there is a winning strategy to win any possible dialogue, starting with an argument for the conclusion. In short, a dialogue is a sequence of moves involving arguments and counterarguments, where each counterargument defeats the argument of the last move (p. 257). A subtlety here is that the proponent's arguments have to be strictly defeating while those of the opponent can be just defeating, a requirement that reflects the dialectical asymmetry between the players. A dialogue is won by a player if the other player cannot advance a defeating argument against the last move.

More formally, a winning strategy takes the form of what is called a dialogue tree (pp. 258f.), that is, a tree of moves involving all arguments that defeat any argument advanced by the proponent of the conclusion. A conclusion is justified, then, if and only if there is a justified argument for it. An argument is justified if and only if there is a dialogue tree with the argument at its root and won by the proponent.

4.3.1 Representing cases

The representation of cases reflects their dialectical structure, that is, the arguments pro and con that appear in them. This is accomplished by representing settled cases as collections of rules from which the arguments can be constructed (p. 256).

The extracted rules in effect determine the factors that are relevant in comparing cases. Suppose, for instance, that a decision is interpreted as the rule that if a migrated employee does not work for a domestic company, then the employee has to pay income tax in the country where he works. The result of this interpretation is that if the issue is where a migrated employee must pay income tax, then one relevant factor is that the employee does not work for a domestic company.

It is not clarified how the rules are to be extracted from settled cases, however. In particular, it is not made clear how the level of generality of the rules is to be determined (cf. the discussion on rule extraction in Chapter 2). Since the rule antecedents determine the relevant factors in Prakken and Sartor's model, their model in effect presupposes one unique choice of relevant factors, without recognising explicitly that several other choices of factors are in principle possible, too. In the model proposed in this book, in contrast, it is explicitly acknowledged that in the law it depends on a contingent choice which relevant factors there are. In particular, the relevant factors can be chosen at any level of abstraction. This becomes apparent in the use of a comparison basis relative to which cases are compared. More specifically, the outcomes of case comparison

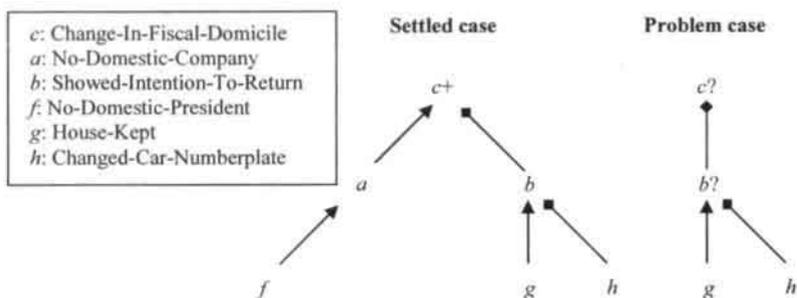
are stated to the comparison basis, while the same holds for the conclusions derived from precedents.

There is another difference with the proposed model in connection with the choice of factors relevant in case comparison. As said, in Prakken and Sartor's model this choice is determined by the rules with which cases are represented. As a practical matter these rules could be obtained, for instance, from judges' opinions. As a result of this practice of representing cases by rules, it is not possible in Prakken and Sartor's model to distinguish between the factors that are relevant for case comparison, and other information concerning cases. More specifically, it is not possible to derive the relevant factors from other case features, whilst treating these features themselves as irrelevant for case comparison. In the proposed model, in contrast, the applying factors are derived explicitly from non-factors in combination with background knowledge, that is, case information that is in general not relevant for case comparison.

As said, factors can be combined into conjunctive reasons, but if appropriate their supporting or attacking effect as individual reasons can also be expressed by separate rules. The reason for this is in the observation that if several factors each favour a conclusion, their combination need not do so (pp. 271-272). Suppose, for instance, that one has that a and b each support conclusion c . Then this can be expressed by the rules $r_3: a \Rightarrow c$ and $r_4: b \Rightarrow c$, but these do not automatically imply the rule $r_5: a \wedge b \Rightarrow c$.

The conflict resolving potential of settled cases is exploited by interpreting them as sources of priority information on conflicting rules (p. 256; cf. also Bench-Capon 1999, pp. 36-38, who combines these priorities with an ordering by specificity). In addition to this, the model also allows for making explicit the grounds for the pieces of priority information thus obtained. If a case with conclusion c contains a rule $r_3: a \Rightarrow c$ and a rule $r_2: b \Rightarrow \neg c$, for instance, then the decision is interpreted as saying that r_3 has priority over r_2 . This can be expressed, for instance, by the rule $r_5: reason \Rightarrow r_2 < r_3$, where *reason* expresses the grounds why r_3 's factor a outweighs r_2 's factor b . An example of a reason for a priority relation between two rules is, that one rule comes from a more on point case than another (p. 271).

If a conclusion depends on a multi-step argument, it may happen that several different sets of priority rules explain the same conclusion. Accordingly, it is not clear then how new problems are to be resolved on the basis of a settled case. An example is in the following figure. In this example a person has migrated to work abroad, and the question is whether this changed his fiscal domicile, that is, the country where this person must pay income tax.



In this figure there is a settled case on the left where the conclusion c was drawn that there was a change in fiscal domicile, as indicated by the plus sign. On the right there is a problem case where the intermediate conclusion b is an issue, that is, the issue whether the person showed an intention to return home. Since the intermediate b is an issue the conclusion c is one as well, as indicated by the question mark that follows it.

The resolution of these issues depends on the way in which the decision in the settled case is interpreted in terms of priority information. There are a number of possibilities. First, one may hold that the rule $r_3: a \Rightarrow c$ has priority over the rule $r_2: b \Rightarrow \neg c$, while b is derivable in the settled case. Assuming that b is derivable in the settled case, the rule $r_7: g \Rightarrow b$ would have priority over the rule $r_6: h \Rightarrow \neg b$. Second, one may hold that the attacking b is not derivable in the settled case, because it is successfully attacked by h . This would mean that the rule $r_6: h \Rightarrow \neg b$ has priority over the rule $r_7: g \Rightarrow b$. Another possibility is to maintain that both priorities hold.

If one chooses the first possibility, this would mean that b is derivable in the problem case. As a result, the opposite conclusion $\neg c$ would then follow. If the second possibility is chosen, however, then b would not be derivable in the problem case. Accordingly, the conclusion $\neg c$ would then not follow then.

This example illustrates that the interpretation of a decision in terms of rule priorities is a non-trivial step which essentially introduces additional information about a case into a dispute. In general several different interpretations of decisions can be given, each having a different precedential effect on other cases. It may well be that also Prakken and Sartor are aware of the possibility of multiple interpretations, since they speak in this connection of “the simplest (...) formalizations” (p. 256). However, they do not explicitly go into this freedom of interpretation that is inherent to the use of multi-step arguments.

In the model proposed in this book, in contrast, no such differences in interpretation are possible. The reason for this is that the conflict resolving potential of settled cases does not have to be made explicit in any form, because the model uses comparison outcomes to derive conclusions in problem cases. More specifically, in the proposed model a direct comparison of cases establishes that there is at least equal dialectical support for some disputed conclusion, and then the conclusion can follow immediately. No intermediate step is

required to interpret the decision in terms of explicit conflict resolving information, because the conclusion that was drawn already implicitly indicates how a conflict is to be resolved.

In the situation in the figure above, for instance, the question whether $\neg c$ follows in the problem case does not depend on any explicit priority information regarding the settled case. Instead all that matters is whether or not the judge in the settled case actually made a decision regarding the intermediate b . Of course it may be difficult to tell from the opinion whether or not this intermediate conclusion was actually drawn, but in general it will be even harder to find explicit priority information. Suppose therefore simply that it was decided that b in the settled case. Then obviously b follows in the problem case, because both cases provide equal dialectical support for b . As a result, $\neg c$ then follows in the problem case as well. If it was not decided that b , though, then $\neg c$ does not follow in the problem case (and neither does c because this positive conclusion is not even supported in the problem case).

In sum, Prakken and Sartor's approach relies on an interpretation of decided cases in terms of explicit conflict resolving information. This interpretation in effect introduces additional information on cases into the dispute. The present model of reasoning by case comparison, in contrast, does not involve such an interpretation, because all that matters there in this respect is which decisions have actually been made by judges. As noted above, however, it is not always unproblematic to tell from judges' opinions which intermediate decisions were actually made before arriving at a final judgement.

4.3.2 Reasoning with cases as premise introduction

In Prakken and Sartor's system there are four ways of reasoning with settled cases (pp. 242-245), viz. analogizing, distinguishing, following and overruling. In their formalisation they focus on analogizing and distinguishing, however. These two ways of reasoning are treated as a kind of premise introduction (p. 261) or theory construction (cf. Prakken 2000, pp. 51f.). More specifically, they involve the introduction of new rules on the basis of rules of settled cases.

Analogizing is treated as selecting a rule from the settled case, and making this rule applicable to the problem case by omitting unfulfilled conditions. The resulting rule is then said to broaden the original one (p. 262). This kind of argument is intended to capture HYPO-style analogising on the basis of shared factors, and in effect it comes down to claiming that the missing conditions were not necessary to arrive at the conclusion. Suppose, for instance, that from the settled case a rule $r_1: a \wedge b \Rightarrow c$ is selected. Suppose, moreover, that only a is present in the problem case, while b is absent. Then the selected rule can be broadened into the rule $r_3: a \Rightarrow c$, which can be applied in the problem case.

Distinguishing is only possible after a rule has been broadened. This kind of argument is meant to capture HYPO-style distinguishing by pointing out relevant unshared factors. It can take place in two ways, strong and weak. Strong distinguishing is treated as the introduction of a rule which states as its

conditions that the ones omitted from the broadened rule are not explicitly satisfied, and which has an opposite conclusion. In effect, strong distinguishing comes down to claiming that the omitted conclusions are necessary to arrive at the conclusion, because in their absence the opposite conclusion follows. Among other things one then needs a priority consideration on the broadening and the distinguishing rule, in order to determine whether the analogizing argument is strictly defeated by the distinguishing one, or vice versa.

Formally, strong distinguishing is captured by weakly negating the omitted conditions and including them in the distinguishing rule. As an example of strong distinguishing, suppose as above that the rule $r_1: a \wedge b \Rightarrow c$ is broadened first into the rule $r_3: a \Rightarrow c$. Then the rule $r_2: \sim b \Rightarrow \sim c$ strongly distinguishes the broadening r_3 . A priority consideration on r_2 and r_3 is then needed to tell whether the analogizing or the distinguishing move is strictly defeating (or neither of them).

Weak distinguishing is treated as the introduction of a rule which concludes to the inapplicability of the broadening rule on the basis of that rule's unfulfilled conditions. In effect this way of distinguishing comes down to claiming that the missing conditions are necessary to arrive at the conclusion, because in their absence the remaining reasons do not support it. The broadening r_3 , for instance, is weakly distinguished by the rule $r_2: \sim b \Rightarrow \neg appl(r_3)$.

This way of reasoning with cases differs in two closely related respects from that of the previous chapter. First, cases are represented by rules in Prakken and Sartor's model, and reasoning with cases is accordingly treated in terms of manipulations of these rules. In the proposed model, in contrast, reasoning with cases involves the use of comparison outcomes rather than rules to derive conclusions. A second difference in this connection is, that the model proposed in this book requires the establishment of a comparison outcome before a settled case can be followed. There must be at least as much dialectical support for a conclusion in the problem case, and only then can one draw the conclusion. In Prakken and Sartor's model, in contrast, there is no intermediate step of case comparison. Instead their analogizing move, for instance, provides an argument that pleads directly for the disputed conclusion.

In Prakken and Sartor's account, the question which conclusions are justified in the problem case can be answered in two ways, depending on which broadening and distinguishing rules are taken into consideration. If one takes as a starting point only broadening and distinguishing rules that have actually been introduced in the course of the debate (p. 266), then it is possible that a conclusion follows by analogising with a broadening rule. If one takes as a starting point all broadening and distinguishing rules that can possibly be introduced (p. 265), however, then no conclusion whatsoever can follow on the basis of an analogy. The reason for this is that if a broadening is introduced in the course of the debate, then this broadening can always be distinguished (p. 266).

In other words, if all opportunities to distinguish are seized, then one cannot justify conclusions by drawing analogies. This is in agreement with the model proposed in this book, where the problem case is required to provide at least as much dialectical support for the disputed conclusion as the settled case. Accordingly, in the model proposed in this book there must not be distinction either between a settled case and a problem case, since this would prevent the settled case from being a precedent. As in Prakken and Sartor's model, the typical example of such a distinction is a reason for the disputed conclusion which occurs in the settled case and not in the problem case.

However, in the proposed model there are more ways of analogising and distinguishing than in Prakken and Sartor's model. While in the latter model cases can only be analogised and distinguished with factors that directly support some conclusion, in the proposed model one can also analogise and distinguish with factors that are relevant in a less direct way. One may use factors, for example, which support or attack the directly relevant factors, or which support or attack that another factor supports or attack a conclusion. Moreover, since these less directly relevant factors can also play a role in establishing outcomes of case comparison, they can thereby affect the relevance of settled cases as precedents, too.

According to Prakken and Sartor it is no real drawback that analogies cannot justify conclusions, though, because in actual legal dispute there are procedural rules which attribute the judge the discretion to evaluate the arguments advanced by parties (p. 267). They propose to treat the judicial evaluation of arguments simply as another kind of premise introduction in the form of priority considerations.

In the model proposed in this book one could also accommodate judicial discretion to evaluate arguments in case comparison. Here one could give the judge the authority to determine the comparison basis relative to which cases are compared. As explained above in connection with the downplaying move in CATO, in this way one can make distinctions irrelevant for the purpose of case comparison. In other words, a judge could make cases comparable by choosing an appropriate comparison basis, and in this way he could determine the final outcome of a dispute. It is left for future research to extend the proposed model by working out this kind of reasoning with the comparison basis. At this point all that is noted is, that this extension would provide an important and interesting contribution of the model, which moreover seems to flow naturally from it because of the contingency of the comparison basis.

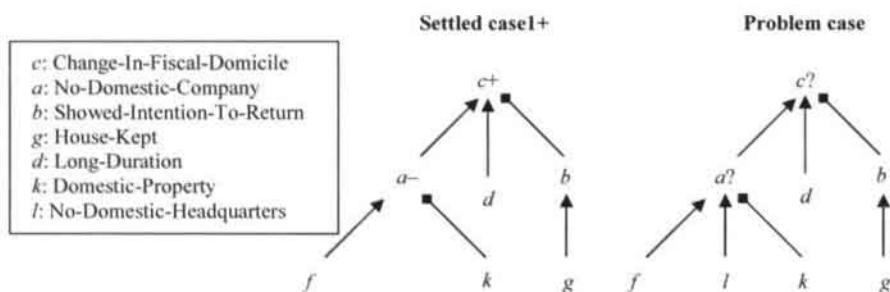
4.3.3 *Case comparison and on pointness*

Prakken and Sartor want to deal with case comparison in terms of HYPO's criterion of on pointness (p. 267). In HYPO, on pointness is measured in terms of the overlap of shared factors (Section 4.1 above). More specifically, one case is more on point than another if the first shares more relevant factors with the problem case than the second.

A difference is, however, that while in HYPO on pointness is used to select the best cases to cite, in Prakken and Sartor's system it induces priority relations on rules. If one rule is selected from a more on point case than a second rule, then the first rule normally has priority over the second (p. 271). Note that this means that the priorities induced from the on pointness ordering depend on the problem case at hand. Prakken and Sartor's model deals with the rule priorities induced from comparison outcomes rather than with these outcomes themselves. This is in line with the fact that cases are represented by rules in their system, rather than being treated as a set of facts or applicable factors.

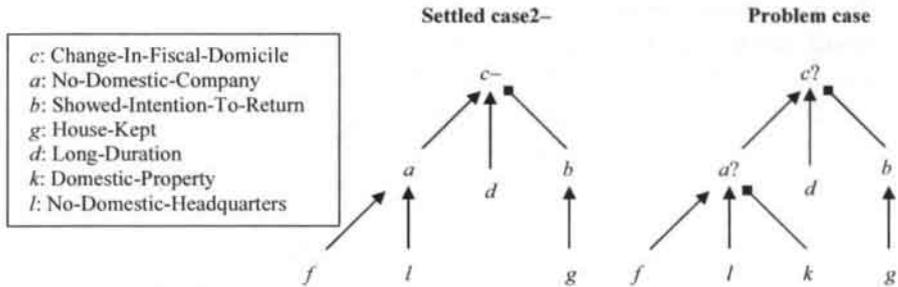
Another difference is that in HYPO a case can be cited for its ultimate conclusion only (even though the system can be presented as a multi-step account of reasoning with factors, see Section 4.1). Moreover, only two opposing ultimate conclusions are possible in HYPO, namely that the misappropriated information is, or is not, a trade secret. Accordingly, in HYPO the on pointness comparison is not done relative to the conclusion at stake, but in terms of all applicable factors. In Prakken and Sartor's system, in contrast, cases can be cited for different conclusions at a time, and in general different sets of factors are relevant for these conclusions. Accordingly, in Prakken and Sartor's system the on pointness comparison is done relative to the conclusion for which the settled case happens to be cited ('local' comparison, p. 269). In particular, they want to allow cases to be compared relative to intermediate conclusions as well. Note that in this respect there is no difference with the model proposed in this book, where the comparison of cases is also done relative to conclusions, and where also settled cases may be precedents for intermediate conclusions.

A problem with Prakken and Sartor's version of on pointness is, however, that it involves factors which are not given but have to be derived. An example is in the following figure. Again the main issue is whether a person's fiscal domicile has changed after migration.



In this figure there is a settled case on the left where *c* was drawn as a conclusion, as indicated by the plus sign. The opposite of the intermediate conclusion *a* was drawn as a conclusion, though, as indicated by the minus sign. On the right there is a problem case where these conclusions are issues, as indicated by the question marks. The arrows in the settled case represent the rules from

which the arguments of that case can be constructed, and the arrows in the problem case stand for rules obtained from precedents. As the figure shows there is a rule in the problem case that *l* supports *a*, while this rule does not occur in the settled case. This rule was obtained from another settled case where the opposite of *c* was drawn as a conclusion, as in the following figure.



Since the two settled cases have opposite conclusions, the question is which of the two settled cases is more on point than the other. There are two possible methods to answer this question. First, one can consider only factors that are directly relevant for the issue at stake, and here these factors are *a*, *b* and *d*. Following this method, the on pointness order of the settled cases depends on the question whether *a* is derivable in the problem case. If *a* is derivable in the problem case, then the second settled case is more on point than the first. If *a* is not derivable, though, then the cases are equally on point. The second method to measure on pointness is to look at the relevant facts of the cases rather than the intermediate *a*. If that is done then neither of the settled cases is more on point than the other, because the set of facts shared with one settled case does not include the set of facts shared with the other. The reason for this is that among the facts of the problem case, *k* is shared with the first settled case and not with the second, while *l* is shared with the second settled case and not with the first.

One can choose different methods to measure on pointness, and it obviously depends on the method chosen how cases are ordered with respect to on pointness. Prakken and Sartor do not make a choice between the different methods, however, and accordingly they do not define comparison outcomes in terms of on pointness (p. 270). Instead they simply assume that the parties have already agreed upon some criterion for on pointness. Accordingly, it is assumed that the on pointness ordering is represented by a set of defeasible rules in the ordered theory from which the debate starts (pp. 270-271). In other words, in Prakken and Sartor's system the outcomes of case comparison are assumed as given.

This is in contrast with the model proposed in the previous chapter, where the outcomes of case comparison are precisely defined, namely in terms of the level of dialectical support for a conclusion. The intermediate conclusions in these arguments need not be derivable, because all that counts is the relative

level of dialectical support for them. In other words, in the model proposed in this book the comparison of cases is separated from the derivation of conclusions. The outcomes of case comparison are defined independently from the conclusions that may follow, and the derivation of conclusions is done on the basis of these given outcomes. By separating case comparison from the derivation of conclusions, the difficulties indicated above are avoided.

There is another difference in this connection between Prakken and Sartor's model and the model proposed in this book. In the latter model a comparison basis is introduced explicitly as a set of factors relevant for comparing cases. Accordingly, the outcomes of case comparison are defined and stated relative to this comparison basis, while the same holds for the conclusions that follow from settled cases. As said Prakken and Sartor do not make a choice for a particular method of comparing cases, however. Moreover, they assume the outcomes of case comparison as given. In accordance with this, they do not make explicit that the outcomes of case comparison depend on the question which factors are deemed relevant for the purpose of comparing cases.

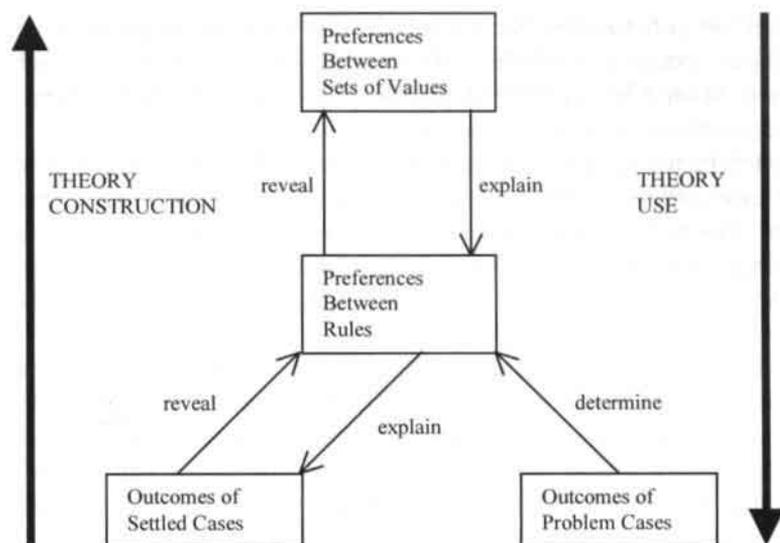
4.4 Bench-Capon and Sartor's approach incorporating theories and values

In a series of recent publications (Bench Capon 2000; 2001; forthcoming), Bench-Capon and Sartor have proposed an approach in which reasoning with cases is treated as the construction and use of theories. The present discussion is based mainly on their paper at the ICAIL conference (Bench-Capon and Sartor 2001).

In their model, theories are intended to explain decided cases in terms of the values promoted by the legal system as a whole, and to help predict the outcomes of new cases. More specifically, these theories contain statements of preferences among values, which are abduced from rule priorities derived directly from decisions.

The reasoning with theories can be construed as consisting of three elements. The first element is constructing one or more theories that explain decisions taken in prior cases. This is done using a number of theory constructors, which tell how theories can be extended with new information. An important example is the constructor which allows for the derivation of rule preferences from decisions. The second element of reasoning with theories is using them in an attempt to establish a legal conclusion in a new case. One can use a rule preference stated by a theory, for instance, to resolve a conflict that occurs in a new case. The third element is evaluating competing theories, to select the most persuasive ones among them. Important criteria in this connection are explanatory power and consistency.

The following is a slightly adapted version of a figure taken from Bench-Capon and Sartor (forthcoming). It shows two elements of Bench-Capon and Sartor's approach, namely theory construction and theory use.



The large arrow pointing upwards stands for theory construction, while that pointing downwards stands for theory use. At the bottom of the figure one finds uninterpreted cases, where the decided cases are on the left. As the figure suggests, the decisions are interpreted in terms of rule preferences, just as in Prakken and Sartor's approach discussed above. In Bench-Capon and Sartor's model, however, another step is taken by interpreting the rule preferences in turn as preferences on the values promoted by the rules. The point is that once these value preferences have been established, new rule preferences can be derived from them. These new rule preferences are then in turn used to resolve new cases.

4.4.1 Reasoning with values

In Bench-Capon and Sartor's model decided cases are represented by sets of factors and the conclusions assigned to these sets (2001, p. 13). The representation of current cases is done in terms of factor sets too. In cases where the pursuit of wild animals is interrupted by someone, for instance, a conclusion could be that there is a legal remedy against the person who interrupted the chase. A relevant factor for this conclusion could be that the person chasing an animal was on his own land.

As in other systems, factors support conclusions.¹⁷ The factor stating that a person was on his own land when chasing an animal, for instance, supports the conclusion that there is a legal remedy against someone who interrupted the chase. An essential feature of Bench-Capon and Sartor's model in this connection is, however, that it can explain *why* certain factors

¹⁷ Actually, in Bench-Capon and Sartor's model the factors only favour sides. Their factors can very well be treated as supporting or attacking conclusions, however.

support certain conclusions. Such explanations are done in terms of the values that the legal system as a whole seeks to promote. An example of such a value is that there should be security of private possession. Another example is that socially desirable economic productivity should be encouraged.

To each factor such a value is linked, namely the value that is promoted by deciding cases where it applies in accordance with the conclusion supported by that factor. The factor stating that a person was on his own land while chasing a wild animal, for instance, promotes the security of private possession, by supporting the conclusion that there is a legal remedy against anyone who interrupts the chase.

Factors supporting the same conclusion can be combined into conjunctive reasons to form rules. Rules are thus a directed connection between factor sets and conclusions. Consider, for instance, the factor that a person was on his own land while chasing a wild animal, and consider the factor that the person was pursuing his livelihood. Suppose, moreover, that both factors support the conclusion that there is a legal remedy against anyone who interrupts the chase. Then these two factors can be combined in to the rule that if someone chases a wild animal on his own land, and is pursuing his livelihood, then there is a legal remedy against anyone who interrupts the chase.

It is not possible that a factor supports (or backs) a rule stating that another factor supports some conclusion, because rules cannot be nested. The reason for this is that rules are ordered pairs of factor sets and conclusions (2001, p. 13), and therefore cannot appear as conclusions of other rules. Accordingly, it not possible either that a rule is supported or attacked by some factor. As a result, attack among rules is only possible by rebuttals (2001, p. 13), which is the reason why undercutting exceptions cannot be dealt with in Bench-Capon and Sartor's model.

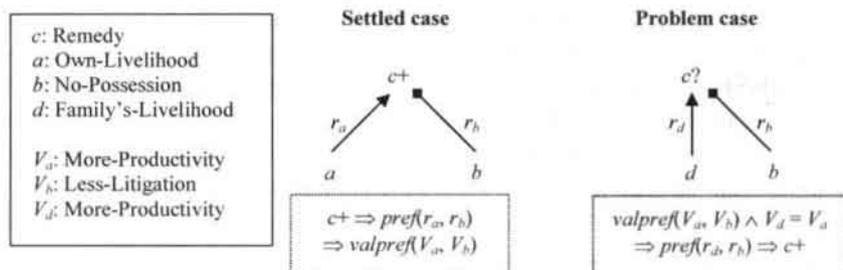
With each rule a set of values is associated, containing precisely those values that are promoted by the factors appearing in the rule antecedent. Suppose, for instance, that a rule is formed from two factors. Suppose, moreover, that one of these factors promotes the value of security of private possession, while the other promotes the value of desirable economic productivity. Then the rule formed from both factors can be said to advance both values promoted by the factors. Accordingly, the set of values associated with the rule contains precisely these two values.

Rules with opposite conclusions attack each other. If one rule has as its conclusion that there is a legal remedy against a trespasser while another says that there is no remedy, for instance, then these two rules are mutually attacking. Such conflicts can sometimes be resolved on the basis of a preference relation between the rules. If rule r_1 is preferred over r_2 , for instance, then this is written $pref(r_1, r_2)$. If two rules attack each other while there is no preference relation between them, then both rules defeat each other. If one is preferred over another, however, then the first defeats the second, but not the other way round. The defeat relation between rules comes into play, for instance, when deciding

whether a certain theory can explain a case. More specifically, if the theory contains a rule which is defeated by another rule, then this theory may not be capable of explaining cases where the conclusion of the defeated rule was drawn.

An essential feature of Bench-Capon and Sartor's model is that rule preferences are not assumed as given but are related to the sets of values associated with rules. More specifically, there is a preference relation between sets of values which stands in one-to-one correspondence with the preferences on the rules. Suppose, for instance, that V_1 and V_2 are sets of values such that the first is preferred over the second. Suppose, moreover, that V_1 and V_2 are the sets of values associated with the rules r_1 and r_2 , respectively. Then r_1 is preferred over r_2 if and only if V_1 is preferred over V_2 . Such a preference between value sets is expressed as $valpref(V_1, V_2)$ (2001, p. 13).

The correspondence between preferences on rules and value sets allows for interesting ways of reasoning with cases, as in the following simple example.



On the left there is a settled case where conclusion c was drawn, holding that there was a remedy against a trespasser. On the right there is a problem case where this conclusion is an issue, as indicated by the question mark. Note that the difference between both cases is that c is supported by a in the settled case, while it is supported by d in the problem case. Here the factor a is that the person chasing the animals was pursuing his own livelihood, while the factor d is that the chasing person pursued his family's livelihood. In both cases the conclusion (c) that there is a remedy is attacked by a factor (b) that the person interrupting the chase had no possession of the animal. According to the model described in the previous chapter, the differences between both cases would make them incomparable, so that no conclusion would follow in the problem case.

In Bench-Capon and Sartor's approach, though, conclusion c can follow in the problem case, depending on the values involved. The line of reasoning along which this is established is briefly indicated in the figure above, in stippled boxes below the dialectical arguments. This reasoning involves rules r_a and r_b and r_d and r_b in the settled and problem case, respectively, where each rule connects a factor with its corresponding conclusion. Rule r_a , for instance, says that a supports c as a reason, while rule r_b says that b supports $\neg c$.

First, the decision in the settled case is interpreted as evidence that r_a is preferred over r_b , just as in Prakken and Sartor's approach described above. Note that accordingly, the difficulty may arise that for multi-step arguments several combinations of preferences can explain the same outcome. This freedom to construct several theories is a central feature of Bench-Capon and Sartor's model, however, and it therefore appears not to be much of a problem. For Prakken and Sartor's model this freedom of interpretation is somewhat more problematic, since there the focus is more on the conclusions that follow than on the theories that can be constructed.

Suppose that r_a promotes the value set V_a , that r_b promotes the value set V_b and that r_d promotes the value set V_d . As indicated in the figure, the value set V_a contains one value, namely economic productivity. Moreover, the value set V_b contains the value that there is less litigation. In addition to this, suppose that r_d promotes the same value set V_a as r_a (that is, $V_a = V_d$), namely economic productivity. Then the preference between the rules r_a and r_b reveals that value set V_a (economic productivity) is to be preferred over value set V_b (less litigation). Moreover, since r_d promotes the same value set V_a as r_a , this means that rule r_d is to be preferred over r_b as well. As a consequence, the conclusion c follows in the problem case too. In other words, despite the apparent difference between the cases, the settled case's decision can be adhered to in the problem case. This shows how value considerations can add information on rule priorities, to help decide problem cases on the basis of settled ones.

Note that there is a hidden assumption behind this kind of reasoning, namely that conflicts among values are always resolved in the same way, irrespective of the context of the cases in which they appear. In other words, Bench-Capon and Sartor assume their value preferences to be context independent. The question then is, of course, to what extent their assumption of context independence is substantiated by legal practice.

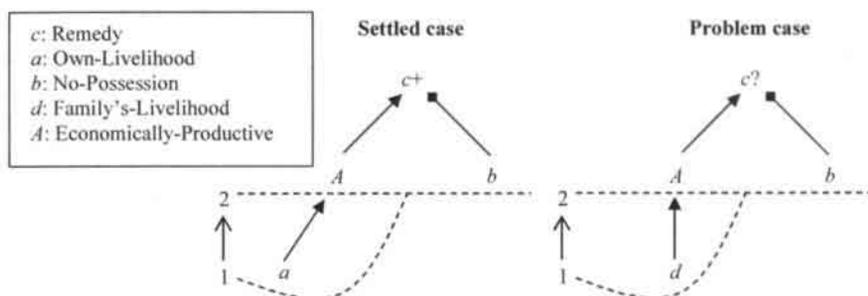
Bench-Capon and Sartor's value-based reasoning with cannot be modelled straightforwardly in the approach described in the previous chapter. The reason for this is that preferences among sets of factors cannot be expressed in the case representation language used in that model. Moreover, this language has no dedicated means to express certain states of affairs about values, such as that one value is preferred over another.

One can speculate, though, on ways in which this kind of reasoning with values could be accommodated in the model proposed in this book. One way to do this might be the following, which is akin to the downplaying move in CATO described above.

The starting point is the observation that if different factors promote the same values, then in Bench-Capon and Sartor's approach one can abstract from differences caused by these factors. If factors a and d , for instance, promote the same values V_a , then a case with a is comparable to a case with d instead of a . This suggests that what actually happens here is, that different factors with the same value are subsumed by a more abstract factor, and that the case compari-

son is done in terms of this more abstract factor. This can be treated by letting a concrete factor support the abstract factor that subsumes it.

More specifically, let A be an abstract factor subsuming the factors a and d with the value set (V_a) containing economic productivity as a value worth striving for. For instance, A can be that the person chasing the animals was economically productive, which also promotes the value of economic productivity. Then the picture above can be adapted as follows.



As the figure shows, in both cases the abstract factor A applies because it is supported by a factor that it subsumes. The comparison basis is visualised by dashed lines. If the comparison is done at level 1, then the cases are incomparable. However, if the comparison is done at level 2 of the abstract factor A , then the two cases are comparable and the settled case can be followed, just as in Bench-Capon and Sartor's model.

This analysis shows a formal connection between Bench-Capon and Sartor's value-based approach and the downplaying move in CATO. The connection is that in both approaches an abstract interpretation is inferred from a factor (this was already hinted at by Prakken 2000, p. 55). Accordingly, Bench-Capon and Sartor's analysis in terms of value preferences can in principle be accommodated in the present model in the same way as CATO's downplaying move, namely as a shift upwards of the comparison basis.

In Bench-Capon and Sartor's model it is not possible, however, to account for reasoning with values as the introduction of a more abstract comparison basis. The reason for this is that their model does not recognise explicitly that in the law it depends on a contingent choice which factors are relevant for case comparison, in contrast with the model proposed in this book. Although the latter model does not provide formal means to reason about the choice of relevant factors, it does make the contingency of the choice explicit by introducing a comparison basis, and by stating the outcomes of case comparison relative to this set.

4.4.2 Which conclusions follow?

In Bench-Capon and Sartor's approach the focus is on the construction, use and evaluation of theories as the main activities in reasoning with cases. Their

model acknowledges explicitly that several different theories can be constructed from the body of decided cases. Each theory can make different conclusions follow in a new case, according to a logic telling which cases and outcomes are explained by the theory (2001, p. 15). Moreover, as theories will typically be constructed on behalf of opposing parties, they will often make opposite conclusions follow in a new case.

The question then is, of course, which conclusions eventually come to hold on the basis of the decided cases. Bench-Capon and Sartor's system cannot answer this question in general, since their logic can only tell which conclusions follow on the basis of one theory. As a result, one would have to choose among competing theories to tell which conclusions come to hold.

They therefore propose to evaluate competing theories with respect to a number of intuitively appealing measures. One measure is the explanatory power of theories, which roughly coincides with the number of cases explained by the theory. Another measure is the consistency of rule and value preferences stated by the theory.

However, Bench-Capon and Sartor do no more than indicate these and some other considerations which may make one theory preferable over another. The corresponding measures for theory evaluation are not formalised any further. As a result, in their system there is no way of deciding which conclusions eventually come to hold on the basis of the decided cases. From Bench-Capon and Sartor's perspective this may be no real drawback, though, given their focus on theory construction and theory use. In the model proposed in this book, in contrast, the focus is on the question which conclusions are derivable from a given set of decisions, rather than on the question which theories can be constructed. A remaining question thereby is, of course, to what extent value considerations actually determine judges' decisions in a compelling way.

4.5 Branting's model of *ratio decidendi*

A computational approach to case-based reasoning was developed by Branting in a number of publications (Branting 1991; 1994; 2000). The present discussion is based on the last publication (Branting 2000).

Branting's model integrates rule-based reasoning with case-based reasoning, and is implemented as a computer program named GREBE (for Generator of Recursive Exemplar-Based Explanations). The purpose of the model is to generate explanations of legal conclusions in a problem case, thereby using rules and settled cases. Settled cases are seen as exemplars, that is, example fact situations warranting an abstract legal conclusion. Explanations can conflict in the sense that they yield opposing conclusions, but there is a measure of relative strength in terms of which they are compared.

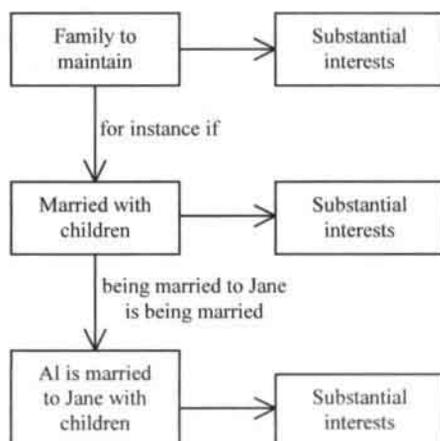
4.5.1 Warrants and explanations

As said, Branting attempts to integrate rule-based reasoning with case-based reasoning. Accordingly, in his model both rules and cases can act as warrants

for conclusions in problem cases (p. 10; cf. Toulmin 1958, pp. 98f.). Warrants are used to form explanations of conclusions in settled cases and problem cases. The explanations are constructed by back-chaining from the conclusion to the facts of a case, where each step corresponds to a warrant from a case or rule.

Explanations can be represented as warrant-reduction graphs (p. 12), of which there is an example in the following figure.

Warrant reduction graph



The figure represents the justification of the conclusion that an employee has substantial interests in keeping his job. The arrows in this figure represent warrants, while facts and conclusions are in rectangles. Horizontal arrows stand for warrants for the conclusion that an employee has substantial interests in keeping his job. Vertical arrows are intended to link these warrants together, which is suggested graphically by links between the condition parts of the linked warrants. These links also represent a kind of warrants, called reduction warrants. Briefly, reduction warrants authorise the reduction of some abstract warrant to a more concrete one. This explains why the vertical arrows for reduction warrants point downwards instead of upwards.

As the figure shows, there are different warrants for the conclusion that an employee has substantial interests in keeping his job, and the conditions of these warrants range from abstract to highly concrete. At the top of the figure the conditions warranting the conclusion are that the employee has a family to maintain, while at the bottom one finds some specific facts of the concrete case of an employee named Al. These conditions correspond to warrants at different levels of generality, connected by the reduction warrants just mentioned.

As will be discussed in the next subsection, only the warrants of lowest generality are relevant for the purpose of generating case-based explanations. In the example this is the warrant stating as sufficient conditions the facts that Al is

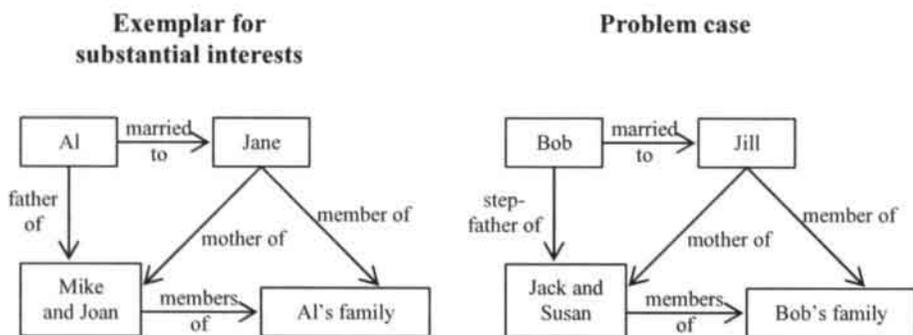
married to Jane and that they have children. Such lowest generality warrants are called exemplars (pp. 3 and 50) or precedent constituents (p. 66).

From the perspective of the previous chapter, warrants are comparable to reason statements holding that one statement supports another. The difference is, however, that in Branting's model the warrants cannot be put at issue. More specifically, it is not possible to support or back up a warrant (pp. 58-60) with reasons, and neither is it possible to attack one.

The explanations in Branting's model are comparable to the dialectical arguments of the previous chapter, in that both represent the reasoning structure underlying a conclusion. The difference is, however, that Branting's explanations only involve warrants supporting conclusions, while no attacks of conclusions are included. This is in agreement with the fact that cases play different roles in the two models. In Branting's approach cases are used as example fact situations warranting a conclusion. In the model proposed in this book, in contrast, cases are treated as a source of information on how certain conflicts are to be resolved. Accordingly, in the present model the representation of cases reflects their dialectical structure.

4.5.2 Case comparison

As said cases are used as exemplars, that is, the smallest collections of case facts that warrant certain conclusions (cf. pp. 58 and 66). Exemplars are a kind of warrants which can be used in building explanations for settled cases or for problem cases, and whose conditions are concrete case facts. These case facts are expressed as (binary) relations that hold between objects or persons (pp. 71-74). The following figure presents an example (not taken from Branting).



On the left there is an exemplar for the conclusion that the employee has substantial interests in keeping his job. Note that the conclusion itself is not included, since the exemplar contains only facts about relations between family members. On the right there is a problem case where this conclusion is an issue. In the rectangles one finds some persons involved in the cases, while the arrows stand for the relations that hold between them. Note that unlike the relations of

support and attack introduced in the model proposed, the relations in exemplars hold between objects and persons and not between statements. The type of statement relevant for case comparison is that some relation holds between objects or persons, such as that Al is the father of Mike and Joan.

Case comparison involves structure matching, which comes down to finding a match between an exemplar and a problem case, such that as many relations as possible still hold in the latter (pp. 83-89). More specifically, a 1-to-1 mapping is constructed from the objects in the exemplar to those in the problem, thereby preserving as many relations as possible.¹⁸ In the figure above, for instance, one would map the husbands and wives to one another, as well as the children and the families. As the figure shows only one relation of the exemplar remains unmatched under this mapping, namely the relation that Al has with the children as their father. This is because Bob is the stepfather of the children, and not their father.

Note that by focusing on relations between objects or persons, Branting's model abstracts from the concrete objects or persons appearing in cases. From the perspective of the previous chapter this in effect means that only these relations are relevant for the purpose of case comparison, while this is not true of the objects or persons between which the relations hold. More precisely, for case comparison one does not consider instantiated relations, such as that Al is married to Jane. Instead one uses generalised relations for that purpose, such as that the employee is married to a person.

An interesting feature of Branting's model is that case comparison can also be done in terms of intermediate conclusions that have actually been drawn in court (pp. 77-82). Suppose, for instance, that in reaching the conclusion that employee Al had substantial interests in keeping his job, from the case facts the intermediate conclusion was drawn that he has a family to maintain. Then the court's reasoning in effect shows that having a family to maintain is what matters for the issue whether one has substantial interests in keeping one's job. More specifically, the court's reasoning shows that if the employee has a family to maintain, then this supports the conclusion that he has substantial interests in keeping his job.

In accordance with the observation that the court drew the intermediate conclusion that the employee had a family to maintain, the exemplar above for having substantial interests can be restated in terms of this intermediate conclusion. Then case comparison can be done in terms of the restated exemplar. See the following figure.

¹⁸ This requirement is relaxed somewhat by allowing semantically similar relations as well (pp. 84 and 108).

**Exemplar for
substantial interests**

Problem case



As the figure shows there is only one relation left in both the restated exemplar and the problem case, namely the relation that the employee maintains his family. Moreover, this relation supports the conclusion that the employee has substantial interests in keeping his job. After the restatement of the exemplar for substantial interests, there therefore is a perfect match between the settled case and the problem case. As a result, the settled case now yields a very strong explanation for the conclusion that the employee has substantial interests in keeping his job. Recall that originally one relation was unmatched, namely the relation that Al had with the children as their father.

From the perspective of the model proposed in this book, this restating of exemplars comes down to interpreting an explicit intermediate decision by a judge as saying that the statements supporting this conclusion do not matter for the purpose of case comparison. The intermediate decision that Al has a family to maintain, for instance, can in this view be interpreted such that for case comparison it makes no difference whether or not one has to maintain one's own children. In other words, in this kind of reasoning the comparison basis is chosen at a higher level of abstraction, in accordance with an intermediate conclusion explicitly articulated by a judge.

This discussion shows both a similarity and a difference between Branting's approach and the model proposed in this book. The similarity is that in both models a court's reasoning can be interpreted as guiding the way cases are to be compared, and in particular at what level of abstraction the comparison should be made. In the proposed model this is possible because the comparison basis is introduced as a contingent set of case features deemed relevant for case comparison. In principle any comparison basis can be chosen there, and one can in particular let one's choice be guided by actual judicial reasoning.

Here the difference with Branting's approach is, however, that in Branting's model the reasoning by a prior court determines which case features are relevant for case comparison. As a result, in Branting's model it does not have to be made explicit that the outcomes of case comparison depend on the set of case features deemed relevant for that purpose. Accordingly, this dependence does not have to be made explicit either for the output of the system (see below). This is in contrast with the model proposed in the previous chapter, where a comparison basis is introduced explicitly as a set of case features relevant for comparing cases, and where the outcomes of case comparison and the conclusions that follow are stated relative to this comparison basis.

4.5.3 *Output*

By chaining rules and exemplars, Branting's implementation GREBE produces explanations for legal conclusions (p. 63). These explanations are ordered by an estimate of their strength (p. 107).

Case-based explanations are stronger to the extent that the match between the cases is better, that is, the smallest proportion of relations remain unmatched under the mapping from the exemplar to the problem case (pp. 95, 98 and 109). On the basis of this measure the available case-based explanations are classified into one of four categories, ranging from "very strong" to "weak".

The generated explanations can be conflicting in the sense that they explain opposing conclusions, as they typically will in hard cases (p. 111). However, in such situations no explicit choice among the different possibilities is made. If there is an explanation for the conclusion that a dismissal can be voided and one against, for instance, then no clear answer follows. Instead the two explanations are both presented in an order determined by their relative strength.

More generally, all that happens in Branting's approach is that different explanations are presented and ordered by strength. In other words, conflicts between explanations are not explicitly resolved by the model. It is not surprising that in Branting's model there is no way of resolving conflicts, however, because as said above, cases are not seen primarily as sources of information on how certain conflicts are to be resolved. Instead they are treated as examples of fact collections where some legal conclusion held. This is in contrast with the model proposed in this book, where the resolution of conflicts is the main purpose for which settled cases are used. Accordingly, in the proposed model the dialectical structure of cases is made explicit in their representation.

In connection with the strength of case-based explanations, there is another difference between the two approaches. In Branting's model the strength is measured in terms of a structural match between cases, that is, the proportion of relations unmatched under a 1-to-1 mapping. In effect this means that statements of relations are considered relevant for the purpose of case comparison, whilst abstracting from the specific objects or persons between which the relations obtain. The question how abstractly these relevant relations are to be stated is in turn determined by the reasoning of a prior court (see above). In other words, in Branting's model a specific choice is made of what counts as relevant in comparing cases, and what does not. In accordance with this, the strength of case-based explanations is not stated relative to a contingent division between relevant and irrelevant case features. In the proposed model, in contrast, it is explicitly acknowledged that in the law it depends on a choice which case features are relevant for comparing cases. Accordingly, the outcomes of case comparison are stated relative to what is called a comparison basis, a contingent division between factors and non-factors.

4.6 Hage's reason-based account of reasoning with cases

Hage gives an account of case-based reasoning in Reason-Based Logic (Hage 1997, Chapters IV and V; Verheij 1996, Chapters 2 and 3). Reason-Based Logic is a rich logical system which involves the use of rules and principles in legal reasoning, and which assigns a central place to the notion of reasons for or against conclusions. Reasons play a role, for instance, when drawing conclusions on the balance of sets of reasons for and against, which is possible when weighing information is available (Hage 1997, pp. 139-141 and 152-153).

It is an important characteristic of Reason-Based Logic that rules and principles are entities with which legal consequences can be created, and that they do not describe these consequences and the conditions under which they can come into existence. Accordingly, rules and principles are denoted by terms in the logical language, and are not expressed by sentences (Hage 1997, pp. 134-136). The conditional structure of rules and principles can be preserved thanks to a special convention in Reason-Based Logic. Briefly, by this convention states of affairs are entities denoted by terms, where the terms are directly obtained from the sentences expressing these states of affairs (reification, Verheij 1996, pp. 30-31).

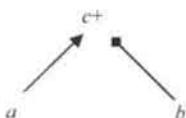
Another characteristic of Reason-Based Logic is its elaborate account of arguments involving rules or principles. With regard to rule-applying arguments, for instance, a subtle distinction is made between the applicability of a rule on the one hand, and the rule's application on the other. In addition to this it is possible to have exceptions to rules or principles, to the effect that conclusions cease to hold on the basis of these rules or principles (Hage 1997, pp. 137-138, 141-143 and 150-151). The exception mechanisms of Reason-Based Logic are akin to the mechanism of attacking that a statement is a reason for or against some conclusion, which plays a role in the account of case comparison proposed in this book. In Reason-Based Logic it is in principle possible to accommodate such mechanisms as well, but as it turns out Hage does not exploit them in his account of case comparison.

In Hage's reason-based account of case-based reasoning, case comparison is done in terms of the reasons for or against the disputed conclusion. From the perspective of the model proposed in this book, reasons for and against conclusions are comparable to supporting and attacking statements, respectively. Since for case comparison Hage only considers reasons for and against the disputed conclusion, in effect his account only captures one-step reasoning.

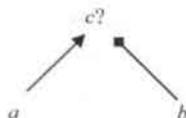
Hage distinguishes five situations, three of which are illustrated by the simple examples presented next.

c: Dismissal-Can-Be-Voided
a: Substantial-Interests
b: Criminal-Record

Settled case



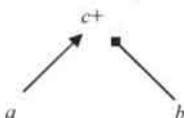
Problem case



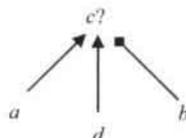
In this example exactly the same reasons appear in both cases, while conclusion *c* was drawn in the settled case. Accordingly, the conclusion *c* follows in the problem case as well.

c: Dismissal-Can-Be-Voided
a: Substantial-Interests
b: Criminal-Record
d: Working-Atmosphere-Not-Affected

Settled case



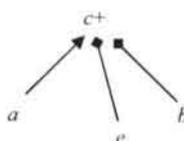
Problem case



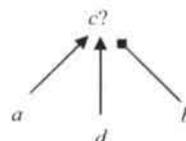
In this example there is an extra reason for the disputed conclusion *c* in the problem case, while for the rest the cases are equal. As a result, the conclusion follows in the problem case.

c: Dismissal-Can-Be-Voided
a: Substantial-Interests
b: Criminal-Record
d: Working-Atmosphere-Not-Affected
e: Forged-Diploma

Settled case



Problem case



In this example there is an additional reason against the conclusion *c* in the settled case, while the conclusion was nevertheless drawn. Accordingly, the conclusion *c* follows in the problem case too.

More generally, the five situations distinguished by Hage can be summarised as follows. A settled case where a conclusion was drawn can be followed if the following three conditions hold:

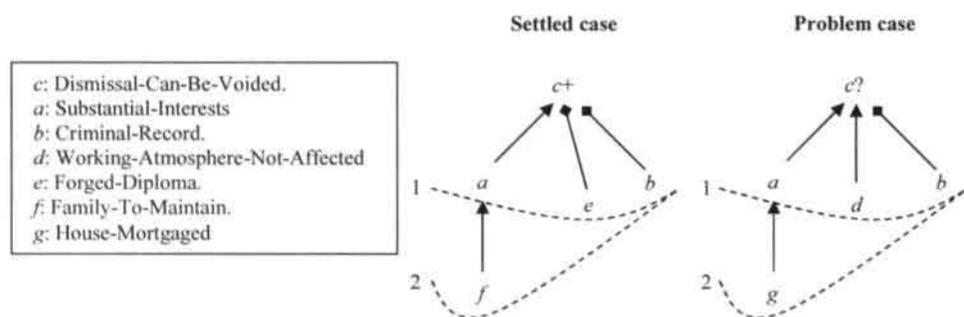
1. In the problem case, the reasons for the conclusion form a (strict or non-strict) superset of the set of corresponding reasons in the settled case; and
2. in the problem case, the reasons against the conclusion form a (strict or non-strict) subset of the set of corresponding reasons in the settled case; and
3. the conclusion was drawn in the settled case.

There are three major differences between Hage's account and the model proposed in this book. First, the formal theory proposed in this book does not merely account for the conclusions that follow by case comparison, but it also accommodates the reasoning patterns along which the conclusions follow, such as analogising and distinguishing. In Hage's account, in contrast, the focus is on the conclusions that follow by comparison with settled cases.

A second major difference is that Hage's account only involves one-step reasoning for the purpose of comparing cases. In the proposed model, in contrast, the dialectical arguments involved in case comparison can consist of multiple steps with intermediate conclusions. Moreover, for case comparison it does not matter whether or not these intermediate conclusions can actually be derived. Since in Hage's account only one-step reasoning is involved in case comparison, however, all the reasons that support or attack the conclusion have to be collected in advance, including those reasons that need to be derived as intermediate conclusions.

A third major difference is that while entanglement of factors can in principle be captured in Hage's Reason-Based Logic, it does not play a role in his account of case comparison. As remarked above the mechanism of attacking that a statement is a reason, for instance, could be modelled in Reason-Based Logic as the exclusion of a principle giving rise to the reason (Hage 1997, pp. 137-138 and 150). As said, in Hage's account mechanisms like this are not involved in case comparison, however. Instead they are used merely to establish in advance which facts are reasons at all, after which case comparison is done in terms of these reasons.

Using the terminology of the previous chapter, it accordingly seems that Hage chooses to lay the comparison basis at the level of statements that directly support or attack the conclusion. In the following figure this is visualised by the dashed line indicated by the number 1.



As a result of this choice underlying Hage's account, the intermediate conclusions would have to be derived first, before they can be involved in the comparison of cases. In the figure above, for instance, this would hold for the

intermediate conclusion (*a*) that the employee has substantial interests in keeping his job.

From the perspective of the previous chapter Hage's account of case comparison is only one choice among many possibilities, however. It is also possible to choose a comparison basis at a deeper level. In the figure this is visualised by the dashed line drawn at level 2. If such a deeper comparison basis is chosen, then the intermediate conclusions do not need to be derived before using them in case comparison, because all that matters then is the level of dialectical support for them.

Since in Hage's account the comparison basis is always chosen at the level of the directly supporting and attacking statements, it is not necessary there to state conclusions relative to this particular choice. This is in contrast with the model proposed in this book, which explicitly acknowledges that in the law it depends on a contingent choice at what level cases are to be compared. Accordingly, the outcomes of case comparison are stated relative to the comparison basis, and the same holds for the conclusions derived from precedents.

In spite of these differences between Hage's account and the model proposed in this book, there also exists a strong abstract parallel between both approaches. The parallel is that Hage's approach can also be seen as one in which cases are compared regarding their dialectical support for the conclusion. The conclusion of a settled case can be drawn, if the case at hand provides at least as much dialectical support for a conclusion as the settled case. In Hage's account the conditions under which this comparison outcome holds are stated as set inclusions involving sets of reasons. Similarly, in the proposed model the outcomes of case comparison are stated in terms of set inclusions involving sets of relevant basic factors (Definition 8). A difference thereby is, however, that in the proposed model the relevances of basic factors can also occur as the result of multi-step dialectical arguments that can involve entangled factors, in contrast to Hage's account.

4.7 Henderson and Bench-Capon's dynamic evolution of case law

Henderson and Bench-Capon (Henderson and Bench-Capon 2000; 2001) take the view that case law evolves dynamically through a series of decisions. Each decision can force one to review one's analysis of some domain of case law, and to reinterpret past decisions in the light of this reviewed analysis. More specifically, decisions can reveal distinctions not previously made, for instance the distinction between medical and non-medical professionals (2000, p. 27).

The process of reviewing the body of case law is modelled in a way that can be described briefly as follows. The starting point is a static common sense ontology of the domain under consideration. This ontology consists of a simple abstraction hierarchy of the concepts used to describe cases in the domain, together with a set of attributes supplied for every concept (except for the concept at the top of the abstraction hierarchy). The concept 'legal profes-

sional', for instance, is more abstract than the concept 'barrister'. Both are attributed a 'professional qualification', but for the barrister the additional attribute of 'working for a solicitor' is introduced. Moreover, for the barrister the legal professional's attribute of 'working for a client' is cancelled.

The ontology is static and cannot be revised in the course of the argumentation. This reflects Henderson and Bench-Capon's view that such a revision is more useful to clarify the reasoning of a body of decided cases than to argue for a point in a current case (2001, p. 68). In this respect Henderson and Bench-Capon's approach is comparable to the model proposed in this book, where a fixed comparison basis is postulated, yielding an entangled factor hierarchy as a static domain analysis.

The ontology is used to establish analogies between cases on the basis of a shared abstraction (2001, p. 64). A case with a solicitor, for instance, is analogous to a case with a barrister, since both are instances of the more abstract category of legal professionals. Here the main difference with the model proposed in this book is, however, that analogies between cases are not stated relative to the ontology with which they are established.

The number of decided cases increases as new decisions come in, and at each point in time new rules are formulated that can explain the whole body of cases decided up to that point. The rule conditions are stated in terms of combinations of (complements of) concepts selected from the ontology. Thus if A, B and C are concepts, then a possible rule states as conditions A and not-(B and not-C). This means that a case must be an instance of A if the rule's conclusion is to be expected to hold in the case, while at the same time it must be an instance of C if it happens to be an instance of B.

The rules are defeasible in the sense that not all decisions have to be explained by them (2000, p. 26). However, a rule is better to the extent that more cases are explained (coverage, 2001, p. 61), and a lower proportion of cases is assigned the wrong conclusion (precision, 2001, p. 61).

As said arguments by analogy involve shared abstractions selected from the ontology. These arguments also draw attention to shared attributes that are introduced for the more concrete concepts, and are not inherited from the abstraction. If a case with a medical consultant is compared with that of a homeopath, for instance, then a shared introduced attribute could be that they are both employed by a client. Typically, such introduced attributes do not apply to cases not involved in the comparison, so that they are characteristic of the cases considered. Similarly, attention is drawn to attributes cancelled from the abstraction. For the cases of the medical consultant and the homeopath, for instance, such a cancelled attribute could be that the medical professional is employed by a hospital. Cancelled attributes like this typically do apply to cases not involved in the comparison, so that their absence is characteristic of the cases under consideration.

The most striking difference between Henderson and Bench-Capon's approach and that proposed in this book is, that the whole body of decided cases is used as a source of general rules. These rules are then used to draw conclusions in problem cases, rather than the settled cases from which the rules are derived. In the model proposed, in contrast, the conclusions are derived directly from the cases and on the basis of an individual comparison among them. If a problem case provides at least as much dialectical support for a conclusion as a settled case, then the decision can be adhered to and the conclusion can follow immediately. As said this is not the way in which conclusions would follow in Henderson and Bench-Capon's approach, however. In this respect their approach is more akin to the rule extraction method described in Chapter 2.

Another difference with the model proposed in this book is, that analogies are supported by pointing out case specific features. More precisely, Henderson and Bench-Capon focus on concept attributes that are typical of the cases under consideration, and on attributes whose absence is typical. To facilitate this kind of case comparison they make use of an ontology enriched with attributes (2000, pp. 30f.). It is an interesting question how one could accommodate attributes to factors in the model proposed in this book, to enrich the account of case comparison that it provides.

4.8 Other research

In this section some other research contributions to case-based reasoning are discussed briefly and compared to the model proposed. The discussion is done in the following order: CABARET (Rissland and Skalak 1991; 1992), BankXX (Rissland, Skalak and Friedman 1996), IVS (Oskamp 1998) and Combrink-Kuiters' jurimetric approach (Combrink-Kuiters 1998).

CABARET (for CASe-BASed REasoning Tool) is an implemented hybrid system which combines case-based reasoning with rule-based reasoning (Rissland and Skalak 1991, 1992). The system focuses on statutory interpretation and it operates on a sub-domain of tax law, namely that of home office deduction. Rules of tax law and regulations contain open-textured predicates, and the system uses case-based reasoning to determine whether or such a predicate is satisfied. From the perspective taken in the present work, the application of rules could be modelled as the introduction of conditional statements from the background information. The reasoning with these conditional statements can then be integrated with case-based reasoning (see Definition 15). Specifically, in the proposed model one could use reasoning by case comparison to establish that the condition of a rule is derivable, and then apply the rule to arrive at its conclusion. Accordingly, there appears to be not so much of a difference as regards the integration of rule-based and case-based reasoning.

CABARET can create arguments from different user-specified points of view (1991, pp. 857f.). In order to perform that task it contains heuristics to

decide whether to use cases or rules, and to find the cases that are appropriate in some line of argumentation. The model proposed in this book, in contrast, assumes all relevant settled cases to be available and accordingly contains no heuristics for finding them, while the relevant support or attack relations are assumed to be given as well. Such heuristics will have to be formulated, of course, if the proposed model is actually implemented in a computer program.

CABARET does not tell which conclusions eventually come to hold on the basis of the arguments advanced. This is in contrast with the model proposed in this book, where for each case a set of derivable conclusions is defined.

BankXX (Rissland, Skalak and Friedman 1996) is a system that models the process of heuristically retrieving information that can be used in legal arguments. This information includes cases, legal theories and domain-specific information. Interestingly, the information on cases is simultaneously represented in different ways (p. 15), with links between the different case representations. For instance, one representation involves the facts appearing in a case, another tells in which other cases it is cited. The links between the different representations help the system to retrieve the cases relevant for creating an argument.

Apparently BankXX focuses on the heuristics of finding relevant pieces of information (such as cases) and on grouping these pieces together into arguments. In BankXX there is no way of deciding whether the arguments can actually make their conclusions follow, however. In the model proposed in this book, in contrast, it is clear which conclusions follow in a case, but on the other hand no heuristics are provided to retrieve the relevant settled cases and background knowledge upon which these conclusions are based.

A case-based reasoning system with an eye on legal practice is Oskamp's Information Supply system for Sentencing (IVS, Informatievoorziening Voor Straftoemeting, Oskamp 1998). Oskamp's system operates in Dutch penal law, where it aims to help judges exchange information on the severity of their sentences. Since the severity of sentences is not subject to the usual prohibition of reasoning by analogy in penal law (art. 1 Dutch Penal Code), this information may be useful as a guideline in deciding new cases.

The system deals with roughly two categories of criminal offences. The first are frequently occurring stereotypical minor offences, while the second are the relatively rare cases in which a severe crime was committed. For the first category the system provides a statistical overview of the sentences imposed by other judges, thus helping the judge to determine the severity of punishment. For the second category the user can obtain a listing of the considerations behind similar cases, where similarity is a number calculated using a mathematical algorithm (the *k*-Nearest Neighbour algorithm). This is useful since judges do not frequently encounter these severe situations, so that it is hard for them to rely on their own experience in similar cases.

Oskamp's system deals with the severity of punishment as a non-Boolean variable conclusion. This is in contrast with the model proposed in this book, where conclusions are Booleans which can either hold or not. Another difference is that in Oskamp's system similarity is measured in terms of a numerical value, while in the proposed model the outcomes of case comparison do not involve numbers. A third difference is that Oskamp's system only lists the considerations behind decisions, and does not tell which conclusions eventually follow on the basis of decided cases.

Combrink-Kuiters (Combrink-Kuiters 1998) contributed to the field of jurimetrics. In jurimetrics one attempts to improve the understanding of judicial decision making as a kind of human behaviour. To this end Combrink-Kuiters attempts to find a quantitative measure of how important factors are in arriving at decisions. This measure is then used to compute for each case a quantitative strength for a conclusion, and to compare cases with respect to this strength (cf. pp. 218f.). This comparison can then yield a prediction for cases that have not yet been decided (pp. 5-6).

An interesting feature of Combrink-Kuiters' approach is that she considers the judge names relevant as factors, so that the judges involved play a role in the comparison of cases regarding their strength (cf. pp. 113-114). This is motivated by the observation that for the decisions taken these judge names appear to be relevant determinants, which is apparently corroborated by the empirical results obtained. From the perspective of the model proposed in this book, this would mean that factors stating judge names would have to be included in the comparison basis. Note that such factors would be extremely case-specific, of course, and differ much from the generalised statements that are used as factors in systems like HYPO (Section 4.1) and CATO (Section 4.2).

4.9 Concluding remarks

The foregoing discussion can be used to illustrate two ways of categorising approaches to case-based reasoning. First, one can distinguish the approaches by the method of employing cases that they focus on. Roughly, this distinction is between the approaches relying on the extraction of rules or theories from cases, and those which emphasise the comparison of cases as a method of deriving conclusions. An example of the former kind is the theory construction approach of Bench-Capon and Sartor, while examples of the latter kind are Ashley's HYPO and the model proposed in this book.

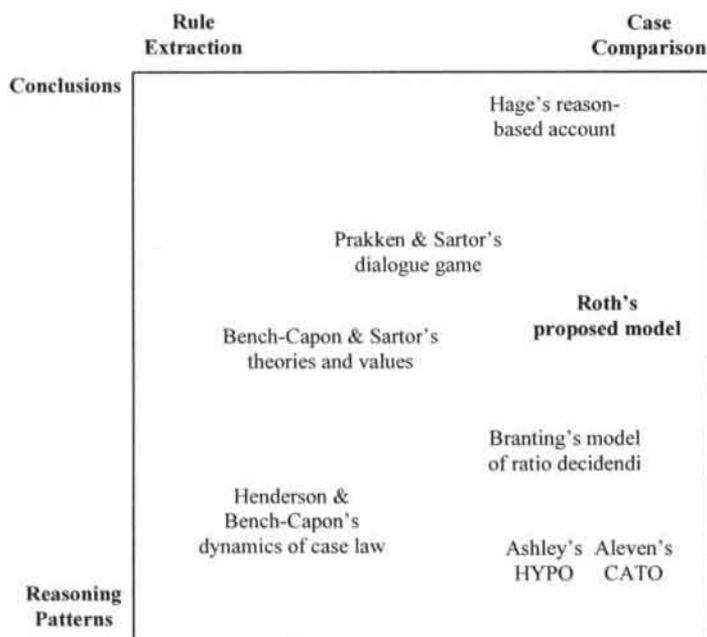
The second distinction between approaches to case-based reasoning is roughly that between approaches that focus on the conclusions that follow from precedents, and those which primarily aim to account for reasoning patterns like analogising and distinguishing, along which the conclusions are established. Examples of approaches focusing on conclusions are Hage's account and the model proposed in this book, while Aleven's CATO is an example of a system that aims at modelling patterns of reasoning.

These two categorisations of approaches to case-based reasoning were already addressed in Chapter 2. In Section 2.2 two methods of employing cases were described and labelled the 'rule extraction method' and 'case comparison method', respectively. Moreover, in Section 2.3 a distinction was made between the question which conclusions follow, and along which patterns of reasoning it is established that they follow.

Approaches to case-based reasoning cannot easily be categorised according to the two divisions sketched, because they tend to have characteristics of more than one category. Prakken and Sartor's approach, for instance, combines elements of rule extraction with elements of case comparison. While in their model cases are represented by rules, case comparison is modelled in terms of on pointness. Moreover, their model addresses the question which conclusions follow and along which patterns of reasoning this is achieved. The conclusions that follow are captured by a formal dialogue game, while the patterns of reasoning are modelled as manipulations on rules.

Although a strict categorisation of approaches to case-based reasoning is not possible, one can attempt to arrange them by their relative emphasis on a particular method of employing cases, and by their relative focus on conclusions rather than reasoning patterns. An example of such an attempt is in the following figure.

Positioning approaches to case-based reasoning



In this figure one finds a square whose sides represent extremes of the two divisions sketched. At the extreme right one has approaches with emphasis on case comparison, while at the left there is more emphasis on rule extraction. At the top the main focus of attention is on the conclusions that follow from precedents, while at the bottom the reasoning patterns involved are more important.

As the figure shows, the model proposed in this book is positioned at the right of the square, because it lays a strong emphasis on case comparison as a method of adhering to decisions. The figure also indicates that in the proposed model there is a little more emphasis on the conclusions that follow from precedents, than on reasoning patterns. In the proposed model there is a little more attention for reasoning patterns than in Prakken and Sartor's, because the model can accommodate more reasoning patterns with cases (downplaying and emphasising, Section 3.9). Hage's account occupies an extreme position at the top, because it focuses on conclusions and contains no typical patterns of reasoning with cases, such as analogising and distinguishing.

Aleven's CATO system is ranged to the right of Ashley's HYPO, because CATO introduced some special argument moves in case comparison. The opposite arrangement could also be defended, though, because CATO lacks HYPO's Dimensions and therefore cannot capture the corresponding ways of comparing cases.

Henderson and Bench-Capon's approach is positioned below Bench-Capon and Sartor's because the latter has a more elaborate logic to tell whether a theory explains a case. At the same time Henderson and Bench-Capon's approach is placed a little more to the left than Bench-Capon and Sartor's, because in the former approach general rules are extracted from cases, while in the latter cases are a source of conflict resolving information.

Branting's model is on the lower right because it uses an elaborate computational definition of similarity involving structure matching, and because conflicts among its explanations are not resolved.

Prakken and Sartor's model is somewhere above the middle of the square, to indicate that it accounts both for conclusions that follow and for case-based reasoning patterns, and that it combines elements of rule extraction and case comparison.

One final remark can now be made on the relative position in the research field occupied by the model proposed in this book. In the proposed model two aspects of case-based reasoning are addressed. The conclusions that follow by case comparison are accounted for, as well as the reasoning patterns along which the conclusions follow. Prakken and Sartor's model also addresses both these aspects, but in other related research the focus is on only one of them. In HYPO, for instance, the primary aim is to produce realistic arguments with cases rather than to answer the question which conclusions follow on the basis

of the arguments. As said Prakken and Sartor's model deals both with reasoning patterns and the conclusions that follow. As an account of reasoning by case comparison their model has the drawback, however, that the outcomes of case comparison are not formally defined, and therefore have to be assumed in advance.

In the proposed model only case comparison is formalised as a method of adhering to decisions. Much related research starts from rule extraction as a method of following settled cases, however. In Prakken and Sartor's model, for instance, the settled cases are treated as sources of general rules. And in Henderson and Bench-Capon's approach the whole body of decided cases is summarised into a set of rules of case law. As said the proposed model deals with the case comparison method rather than with rule extraction. It is an interesting question for future research whether and how a connection can be established between rule extraction and case comparison as methods of adhering to decisions.

Chapter 5: Evaluation

In this chapter the results of the present research are presented, in the form of answers to the research questions of Chapter 1. Then the proposed model of reasoning by case comparison is critically evaluated. The evaluation addresses the model's strengths and weaknesses in relation to relevant other research. The weaknesses also suggest possible subjects for future research.

5.1 Results

In Chapter 1 a number of groups of research questions were posed. This section briefly discusses how one can answer these research questions within the proposed model.

- How can case comparison be formalised as a criterion to determine whether a settled case can be followed? In particular, how can the dialectical structure of cases be used in their comparison?

Reasoning by case comparison is formalised as a kind of reasoning *a fortiori*. In the formalisation, case comparison concerns the relative dialectical support for a conclusion (Section 3.6). If the problem case provides at least as much dialectical support for a conclusion as a settled case, then the settled case can be followed and its conclusion can hold.

The dialectical structure of cases is made explicit by tree-like structures called dialectical arguments (Section 3.1). Dialectical arguments involve statements supporting or attacking other statements. In addition to this, it is possible to support or attack that one statement supports or attacks another.

- Which conclusions can be drawn by comparison with settled cases, and along which patterns of reasoning? In particular, what if precedents contradict each other?

The conclusions that can be drawn depend on the case base, that is, the set from which the settled cases are selected that can be relevant as precedents (Section 3.7). A settled case in the case base is relevant with respect to a problem case, if and only if the problem case provides at least equal dialectical support for the conclusion.

Reasoning patterns typically found in this connection are analogising and distinguishing (Section 3.9). Analogising involves pointing out significant similarities, that is, shared factors that affect the dialectical support for the conclusion. Distinguishing involves pointing out significant distinctions, that is, factors that tend to make the problem case provide less dialectical support for the conclusion than the settled case.

If there are precedents for opposing conclusions, the precedents have been decided in a contradictory way (Section 3.7). Accordingly, in that situation no conclusion follows by comparison with settled cases (Section 3.8).

- How does the method of case comparison relate to other approaches to case-based reasoning? In particular, how is case comparison formalised in these approaches? Is there an account of the conclusions that follow?

The proposed model has focused on one particular method of adhering to decisions, namely reasoning by case comparison. The model is unique in that the comparison outcomes and the derivable conclusions are defined relative to what is called a comparison basis (Chapter 4). This is a set of factors with the property that the statements supporting them are not relevant for the purpose of case comparison, so that all that matters is whether or not they apply to a case. The comparison basis is postulated as a contingent set, and no criteria are proposed to guide a choice among different possibilities.

Case comparison is captured in most work by indicating the argument moves or reasoning patterns that are possible, such as analogising and distinguishing (*e.g.*, HYPO, Ashley 1990, see Section 4.1; CATO, Aleven 1997, see Section 4.2). In this work the conclusions that follow by case comparison are not specified. In other work the conclusions are defined, but the outcomes of case comparison are assumed to be given (Prakken and Sartor 1998, see Section 4.3). In the proposed model the outcomes of case comparison are precisely defined, as relations concerning the dialectical support for the conclusion. In addition to this, there is an account of the conclusions that follow by case comparison.

5.2 Evaluation of the model

In this section the model is critically evaluated, thereby focusing on its strong points and weaknesses in relation to other relevant research. Suggestions for future research are also made.

A strong point of the model is that it integrates the case-based reasoning research tradition within AI & Law in a formal modelling approach, thereby retaining the strengths of both.

A strong point shared with the case-based reasoning tradition (Rissland and Ashley 1987; Ashley 1990; Aleven 1997) is that cases occupy a central place as its object of reasoning within the proposed model, rather than for instance rules derived from them. Another shared strong point is that a number of reasoning patterns typically found in case comparison are accounted for, such as analogising and distinguishing (*cf.* Ashley 1990, pp. 25f.; Aleven 1997, pp. 58f.).

A strong point shared with the formal modelling tradition (Prakken 1997; Hage 1997; Verheij 1996) is that the conclusions following from precedents are precisely defined for each case. In this respect the proposed model goes a step further than systems like HYPO (Ashley 1990), CATO (Aleven 1997) or

GREBE (Branting 2000), because these systems merely propose arguments pro and con. Moreover, this feature of the model makes it possible to predict the decisions that will actually be made in court, and the model's predictive value in this respect makes it especially relevant for the domain of law.¹⁹ Another strong point shared with the formal modelling tradition is that a number of concepts pertaining to the case comparison method are made formally precise. The concept of a precedent, for instance, is formally defined with the help of the concept of dialectical support for a conclusion.

A strong point of the model's account of case comparison is, that in comparing cases it is not only possible to state reasons for or against conclusions, but also to support or attack that a statement is a reason. If it is stated that having a family to maintain supports having substantial interests in keeping one's job, for instance, then this can be attacked by the statement that the person's wife has a good income. Attacking a support or attack statement is akin to what is known as an undercutting exception (Hage 1997, p. 166; Verheij 1996, pp. 200-201; Pollock 1995, pp. 41 and 86; Pollock 1987, p. 485), while supporting a support or attack statement a conclusion is comparable to Toulmin's backing of a warrant (1958, pp. 98f.). A step forward to deal with these mechanisms was to treat it as a statement itself that a conclusion is supported or attacked (cf. Verheij's DEFLOG 2000c, pp. 5f.; see also Verheij 1999a, pp. 45f.). Accordingly, these mechanisms were captured by using nested constructions in the case representation language, while they also appeared as the entanglement of factors in the entangled factor hierarchy. With respect to this entanglement of factors in case comparison, the proposed model is more general than systems like CATO (Aleven 1997) or HYPO (Ashley 1990).

Another strong point is that the model explicitly acknowledges that in the law it depends on a contingent choice which case features are relevant for the purpose of case comparison (see Section 2.5 on the issue of relevance). In particular, the relevant case features can be stated at different levels of abstraction. In other words, the 'problem of importance' discussed in legal theory (Burton 1985, pp. 31f. and Part II; cf. also Section 2.5) can at least partly be addressed. It is possible, for instance, to include the employee's sex as a relevant factor in dismissal cases, but one may also leave this out of consideration. Likewise it is possible, for instance, to count as relevant that an employee caused considerable damage to company property, but one may also focus on more concrete things, such as that the employee cracked up a company car. Formally this is captured by using a comparison basis as a contingent division between factors and non-factors. By choosing different comparison bases one can obtain different outcomes of case comparison and derive different conclusions. Accordingly, the comparison outcomes and the derived conclusions are stated relative to the comparison basis. As said before, the explicit use of a comparison basis is a unique feature of the proposed model, which is lacking in

¹⁹ Cf. Brüninghaus and Ashley 2003 on predicting decisions.

other approaches to case-based reasoning. In the HYPO system (Alevén 1997), for instance, the outcomes of case comparison are not defined relative to the set of factors pertaining to the domain, but instead presented as if they were absolute.

Another strong point of the model is that the factors that apply to a case can be derived from non-factors in the case and in the background, and therefore do not have to be given in advance. If it is a factor that considerable damage was done to company property, for instance, then this can be derived from the supporting statement that a car was cracked up, or that an expensive pipe was dropped. In other systems, in contrast, the factors have to be identified by hand (as in CATO, Alevén 1997), or are simply assumed to be given (Prakken and Sartor 1998).

A final strong point of the proposed model is that conflicts in problem cases are resolved directly with the help of decisions taken in settled cases (cf. Section 2.4 on case based reasoning as conflict resolution). If case comparison yields the outcome that the problem case provides equal or more dialectical support for the disputed conclusion than a settled case, then the settled case can be followed and its conclusion can hold. In this respect the proposed model may seem similar to the HYPO system (Ashley 1990), but here it must be recalled that in HYPO no conclusions are attached to settled cases (see Section 4.1). Since precedents can be followed immediately by copying their conclusions, it is not necessary to interpret decisions in terms of explicit conflict resolving information. This is in contrast to some other systems (Prakken and Sartor 1998; Bench-Capon and Sartor 2000; 2001), which involve an interpretation of decisions in terms of rule priorities. As argued in Section 4.3, many different such interpretations are possible when a decision depends on a multi-step argument. Moreover, different interpretations can then yield different conclusions in problem cases.

Next to these strong points, the proposed model also has some weaknesses.

A weakness of the model is that while for one-step dialectical arguments the model reflects reality fairly well, it does not seem very realistic to assume that lawyers actually compare multi-step dialectical arguments when reasoning with cases. In other words, for multi-step arguments it seems that the model's account of case comparison does not reflect the way in which cases are actually compared in legal practice. It is an interesting subject for future research how the model could be made more realistic in this respect.

Another weakness of the proposed model is that it does not deal with explicit priority or weighing information. In this respect Prakken and Sartor's model (1998) and Hage's Reason-Based Logic (1997), for instance, provide a more complete account. Moreover, in principle the conflict resolving potential of settled cases can also be captured by interpreting them in terms of weighing or priority information. That weighing information is not dealt with means that no conclusions are derived solely on the balance of the reasons pro and con.

Instead such information is assumed to be implicit in the way cases have been decided. Accordingly, conclusions are derived directly from settled cases, once it has been established by case comparison that they are relevant as precedents. As far as conflicts can be resolved with precedents it is no limitation that explicit conflict resolving information is not dealt with. In practice there are other sources of conflict resolving information as well, however, such as when it is self-evident from plain common sense how a conflict is to be resolved (cf. Aleven 1997, p. 65). Since priority information cannot be dealt with in the proposed model, principles involving the hierarchy among courts (*Lex Superior*) or the relative recency of decisions (*Lex Posterior*) cannot be taken into account in a straightforward way. To take these principles into account in the proposed model, one would have to order the precedents in the case base by their importance for the case at hand.

Another weakness is that no numerical value is assigned to factors, because the case representation language does not provide the means to express that a factor has a certain value. Accordingly, in the model no dimensional comparisons of cases are made in terms of the values of factors. In this respect the HYPO system (Ashley 1990) is richer than the proposed model, since in that system case comparison is done in terms of dimensions rather than factors. They are used, for instance, to distinguish a cited case with regard to the magnitude of a shared factor. The absence of dimensions surely is a limitation of the model, as Bench-Capon and Rissland have stressed again recently (2001). The reason for this is that dimensions allow for a finer grained comparison of cases than factors, because their support for a conclusion can range over different positions (pp. 41-42). Moreover, one and the same dimension can support opposing conclusions, depending on the case's position along a dimension (p. 42; cf. also Rissland and Ashley 2002, p. 70, who introduce the term "HYPO-style factor" for a dimension that can support opposing conclusions). From the perspective of the proposed model the need for a dimensional approach becomes apparent by the fact that the values of factors can influence the dialectical support for a conclusion. If having caused considerable damage attacks that the dismissal can be voided, for instance, then the dialectical support for this conclusion will decrease as the size of the damage increases. It is an interesting subject for future research how this kind of reasoning could be accommodated in the present account of case comparison.

Another weakness of the proposed model is that it only involves the mechanisms of support and attack as part of the dialectical structure of cases, where attack of a conclusion is treated as support for the opposite conclusion. Other common logical mechanisms cannot be expressed in a direct way, for instance the common conjunction (and) and disjunction (or). For instance, several factors never combine into one conjunctive reason. As a result, the model does not capture rules with cumulative conditions, such as the rule that if someone owes a duty of reasonable care and acts contrary to this duty, then there is liability for negligence. Systems that employ a richer logic seem

stronger in this respect, such as that of Prakken and Sartor (1998), Hage (1997) or Branting (2000). However, using only support and attack it is possible to express other logical mechanisms as well. Take for instance the rule on liability for negligence, which states as conditions that (1) there is a duty of reasonable care, and (2) that someone acts contrary to this duty. This rule could also be expressed by saying that (1) owing a duty of reasonable care supports that (2) having acted contrary to the duty supports the conclusion that there is liability for negligence. By restating common logical mechanisms like this in terms of support and attack, it is in principle possible to accommodate them in the present model as well. The only problem then is how the restatement is to be done, because often there are different reasonable alternatives. An alternative formulation of the rule on negligence, for instance, may be obtained simply by interchanging the conditions. The rule would then say that (2) having acted contrary to a duty of reasonable care supports that (1) owing the duty supports the conclusion that there is liability for negligence.

Another weakness of the proposed model is that it does not formalise reasoning with hypothetical cases as a separate reasoning strategy. In the model no hypothetical variants are considered, for instance, to test the viability of a legal claim. In this respect the model differs from the HYPO system (Ashley 1990), which is particularly strong in its account of reasoning with hypotheticals. However, in the proposed model it is relatively straightforward to accommodate reasoning with hypothetical cases. This could be done, for instance, by considering the conclusions that would follow from a hypothetical modification of a problem case, or by adding a hypothetically modified settled case to the case base. The fact that hypotheticals are not accommodated also means that there is no account for forward-looking or consequentialist case-based reasoning, in which the effect of a decision on hypothetical other cases is taken into consideration (cf. MacCormick 1997, pp. 129f.). More specifically, in consequentialist case-based reasoning one avoids taking decisions of which the effect on other cases would be undesirable, for instance in the light of values or goals promoted by the system of law. From the perspective of women's emancipation, for instance, it may under certain circumstances be undesirable to decide that a woman employee's dismissal cannot be voided, because that would have an unwanted effect for other dismissal cases. Notwithstanding its importance for legal practice and for legal theory, consequentialist case-based reasoning has so far been a neglected subject in AI & Law. To the present author's knowledge no theoretical or practical systems exist that capture this reasoning strategy. It may therefore be a promising subject for future research.

Another weakness of the proposed model is that the procedural aspects of legal reasoning are not addressed (*e.g.*, Hage, Leenes and Lodder 1994). The conclusions following from precedents are defined on the basis of case facts and background information, and are not determined by arguments advanced within the context of a legal procedure. However, in actual legal procedure it is conceivable, for instance, that a party does not seize all its opportunities, such as

distinguishing a cited case. If a dismissed person who is not very esteemed cites a dismissal case where the dismissal of a highly esteemed employee could be voided, for instance, then the person's employer may fail to draw attention to this distinction between the cases.

Another weakness of the model is that it deals exclusively with the case comparison method of following cases. Recall from Chapter 2, though, that next to case comparison one also has rule extraction as a method of adhering to decisions. Briefly, the difference between case comparison and rule extraction is that in the former method a conclusion follows directly by adherence to a decision, while the latter method involves an intermediate step where rules are extracted from settled cases. With regard to this step of rule extraction the approach by Henderson and Bench-Capon (2000; 2001) is the most advanced, since in that approach the rules are dynamically revised in the light of each new decision. In most approaches the rule extraction step is skipped and the rules are assumed to be given in advance (Prakken and Sartor 1998; Branting 2000). As said the rule extraction method is not dealt with by the proposed model either, however.

The most serious weakness of the model is that within the formalism proper, one cannot deviate from conclusions derived from precedents. In other words, it is not possible to depart from *stare decisis* as it is presently formalised, viz. as a kind of reasoning *a fortiori*. In actual legal procedure, however, a court may be attributed the authority for such deviations (cf. Raz 1979, pp. 183f. on deviations from precedent; see also Section 2.8). Moreover, the court may have good grounds to decide against the principle of *stare decisis*, namely if this seems desirable in the light of values or goals promoted by the system of law. In Section 2.8 two possibilities were distinguished, which can be rephrased using the terminology of the proposed model.

First, it may happen that a court intends to enforce a change in the comparison basis, for instance because it considers this desirable in the light of legal values or goals. The procedural authority to introduce such changes will typically be attributed to the highest courts only, and accordingly they often have consequences for all other legal decision makers. As an example of such a change, suppose that from the perspective of women's emancipation a court believes that in dismissal cases the employee's sex should be included as a relevant factor. Then if the court has the appropriate authority, it could simply rule that this factor always has to be taken into consideration in deciding dismissal cases. In situations like this the decision is intended to change the comparison basis, while in addition a new precedent is set in accordance with this change.

Second, it may be that a precedent gives rise to a conclusion that is considered undesirable in the concrete case at hand. Then the judge may want to make an exception to the principle of *stare decisis* by setting aside the precedent, but without questioning the comparison basis. It may be, for instance, that the judge decides that if a dismissed employee is near retirement, then the consequences

of a dismissal would be disproportionate to the goal served by it. In situations like this the decision is not intended to set a new precedent, while moreover the comparison basis is not questioned and therefore remains untouched.

In sum, several departures are conceivable from *stare decisis* as it is formalised in the proposed model. These departures could be grounded, for instance, in the values and goals that the system of law seeks to promote. One possibility could be that a court aims at enforcing some desirable change in the comparison basis, such that the outcomes of case comparison are changed accordingly. Another possible departure is that an exception to *stare decisis* is made for a single case, without putting the comparison basis at issue. None of these possibilities is presently included in the formalism proper, because that presupposes a given comparison basis and precisely defines the conclusions that follow.

Changes in the comparison basis can be accommodated straightforwardly in the proposed model, however, as demonstrated in Section 3.9 on downplaying and emphasising. Accordingly, the model can also cope with the value- or goal-based considerations in case-based reasoning that are captured by some related work in AI & Law, for instance the approach by Bench-Capon and Sartor (2000; 2001) or by Prakken (2000) (cf. also Hage 2000, where values and goals are used for comparison of statutory regulations instead of cases). As shown in the previous chapter, it seems that the account given by Bench-Capon and Sartor, for example, can be reformulated as a scheme according to which a more abstract comparison basis is derived from values.

Resuming, the proposed model has the following strong points:

- The model integrates AI & Law research on case-based reasoning in a formal modelling approach, combining the strengths of formal modelling with those of research on case-based reasoning.
- A strong point shared with research on case-based reasoning is that cases occupy a central place as the objects of reasoning. Another shared strength is that a number of common reasoning patterns in case comparison are accommodated, such as analogising and distinguishing.
- A strong point shared with the formal modelling approach is that the conclusions that follow by case comparison are precisely specified, and that concepts such as 'precedent' are made formally precise.
- It is possible to support or attack that a statement supports or attacks a conclusion, and these mechanisms are included in the formalisation of case comparison.

- It is explicitly acknowledged that in the law it depends on a contingent choice which case features are relevant for the purpose of case comparison. This is captured formally by what is called a comparison basis.
- The factors that apply to cases can be derived from non-factors and the background, and do not have to be given in advance.
- For the purpose of following a case it is not necessary in the model to resolve all conflicts with regard to intermediate conclusions. Moreover, decisions do not have to be interpreted in terms of priority or weighing information.

However, the proposed model also has some weaknesses:

- While for one-step argumentation the model reflects reality fairly well, for multi-step argumentation it does not seem realistic that lawyers would actually compare dialectical arguments when reasoning with cases.
- The model does not capture reasoning with explicit weighing or priority information. In particular, conflicts among precedents are not resolved on the basis of considerations like the hierarchy among courts, or the order in time in which cases have been decided.
- Cases are not compared in terms of numerical values attributed to factors. In other words, no dimensional comparison of cases is done.
- Common logical mechanisms like conjunction and disjunction cannot be accommodated in a direct way, since the model only recognises the mechanisms of support and attack as part of the dialectical structure of cases.
- Reasoning with hypothetical cases is not formalised as a separate reasoning strategy. In particular, consequentialist or forward-looking case-based reasoning is not accommodated.
- The procedural aspects of reasoning with cases are not addressed by the model.
- Only case comparison is formalised as a method of adhering to decisions. Rule extraction is not accommodated.

- It is not possible to depart from stare decisis as it is presently formalised. In particular, changes in the comparison basis are not made within the formalism proper, for instance on grounds derived from values or goals underlying the system of law.

With this summary of strengths and points of criticism, the present thesis about case-based reasoning in the law is concluded.

Summary

This thesis is about case-based reasoning in the law. Case-based reasoning is presented as a method of adhering to decisions taken in settled cases, on the basis of a comparison with the case at hand. If the case at hand provides at least as much dialectical support for the conclusion as a settled case, then the settled case can be followed and its conclusion can be drawn. The thesis sets forth a formal theory of the method of reasoning by case comparison, and it compares the theory to existing approaches to case-based reasoning.

The work has aimed at answering roughly two groups of research questions. The first group concerns the method of reasoning by case comparison itself:

- How can case comparison be formalised to determine whether a settled case can be followed? Which conclusions can be drawn by comparison with settled cases? Along which patterns of reasoning do the conclusions follow?

The second group of research questions concerns the proposed model's relation with existing approaches to case-based reasoning in Artificial Intelligence and Law (AI & Law):

- How does the method of case comparison relate to other approaches to case-based reasoning? To what extent is case comparison formalised in these approaches? Is there an account of the conclusions that follow?

Both groups of questions have been addressed in the five chapters of the present thesis, the contents of which can be summarised briefly as follows.

In Chapter 1 the research context is sketched within which the present work on case-based reasoning was done. The work is positioned within the field of AI & Law, a field where a variety of legal activities is investigated from an Artificial Intelligence perspective, ranging from legal decision making to drafting bills.

In Chapter 2 some issues concerning case-based reasoning in the law are discussed that are especially relevant for the present thesis, both from a legal-theoretical and an AI & Law perspective. Two methods of adhering to decisions are presented, namely case comparison and rule extraction. In the case comparison method the conclusion of a settled case is copied directly after a comparison with the case at hand. In the rule extraction method the decision is first summarised into a general rule, after which the rule is applied to the case at hand.

The chapter goes on to discuss some relevant issues concerning the case comparison method. One issue is which conclusions can be established on the basis of this reasoning strategy. Another issue is along which patterns of reasoning one arrives at the conclusions. Patterns of reasoning typically found

in case-based reasoning are 'analogising' and 'distinguishing'. The former involves pointing out similarities between a decided case and the case at hand, while in the latter the relevant differences between the cases are stressed.

Another relevant issue that is discussed in Chapter 2 is, that the outcomes of case comparison are determined by the set of case features that are deemed relevant for the purpose of comparing cases. In particular, it is important at which level of generality the relevant case features are to be stated.

In Chapter 3 a formal theory is developed for reasoning by case comparison. The main objective of the theory is to specify precisely the conclusions that follow on the basis of the case comparison method. A secondary aim is to account for reasoning patterns involved in case comparison, such as analogising and distinguishing.

In the chapter, formal conditions are stated for when there is an analogy between a problem case and a settled one. Intuitively, these conditions hold that the problem case provides at least as much support for the disputed conclusion as the settled one. Then by a variant of reasoning *a fortiori*, the conclusion can hold in the problem case as well.

To determine whether the settled case can be followed in the problem case, a comparison is made of the argumentation relevant for the conclusion. This argumentation is made explicit by what is called a dialectical argument. Dialectical arguments include statements of fact upon which conclusions can be founded, but also intermediate conclusions. Moreover, in contrast to other work on case comparison, the model does not take for granted which statements are reasons for or against conclusions. Accordingly, it is possible to support or attack that one statement supports or attacks another, which results in a kind of entanglement of the dialectical arguments. In the proposed model this entanglement is also taken into account in case comparison, which becomes apparent in the use of an entangled factor hierarchy.

An important feature of the present model is, that it is acknowledged explicitly that in the law it depends on a contingent choice which case features are relevant in comparing cases. In accordance with this, the relevant case features are given by what is called a comparison basis, and the comparison outcomes are stated relative to this comparison basis.

In Chapter 4 it is discussed how the proposed model relates to other approaches to case-based reasoning. One issue thereby discussed is whether the conclusions that follow from settled cases are defined, another issue is whether there is an account of the reasoning patterns along which these conclusions follow. A third issue is whether adherence to decisions is treated as rule extraction or as case comparison. As seen, none of the discussed approaches presents a formal account of the case comparison method that includes the both the conclusions that follow, and the reasoning patterns along which they follow. A fourth issue is whether it is acknowledged that in the law it depends on a contingent choice which case features are relevant for case comparison. As

seen, none of the discussed approaches explicitly deals with the contingency of the set of relevant case features.

In Chapter 5 the theory is critically evaluated. A strong point of the model is that the conclusions that follow by case comparison are specified. The chapter discusses a number of reasoning patterns that are not yet captured by the proposed model. Among these one has, for instance, reasoning on the basis of the hierarchy among courts. The evaluation also discusses how such patterns of reasoning may be dealt with within the model proposed. In addition a number of subjects for future research are recommended, like modelling reasoning with hypothetical cases.

This concludes a brief summary of the contents of the present thesis. The main contributions of the thesis can be summed up as follows.

- Case comparison is formalised as a method of adhering to decisions. The outcomes of case comparison are formally defined, and there is an account of the conclusions that follow by case comparison.
- The dialectical structure of cases is made explicit by tree-like structures called dialectical arguments. Dialectical arguments involve statements supporting or attacking other statements. In addition to this, it is possible to support or attack that one statement supports or attacks another, giving rise to a kind of entanglement of the dialectical arguments. It is an innovation of the present model that this entanglement plays a role in case comparison, which becomes apparent in the use of an entangled factor hierarchy.
- The model explicitly acknowledges that in the law it depends on a contingent choice which case features are relevant in case comparison. Formally this is done by introducing a comparison basis as a contingent set of factors relevant for case comparison, and by defining the comparison outcomes and the derivable conclusions relative to the comparison basis.
- Typical reasoning patterns in case comparison are accommodated, namely analogising, distinguishing, downplaying and emphasising. The latter two patterns are accounted for as a change in the comparison basis, whereby it is an innovation is that downplaying and emphasising are not only possible for distinctions, but for similarities as well.

Samenvatting

Dit proefschrift gaat over casusgebaseerd redeneren in het recht. Casusgebaseerd redeneren wordt gepresenteerd als methode om vast te houden aan rechterlijke beslissingen, op basis van een vergelijking met de voorliggende casus. Als het onderhavige geval tenminste even veel dialectische steun geeft aan de conclusie van een beslecht geval als dat beslechte geval zelf, dan kan de beslechte casus worden gevolgd en kan de overeenkomstige conclusie worden getrokken. In het proefschrift wordt een formele theorie ontwikkeld van deze methode van redeneren op basis van casusvergelijking, waarna de theorie wordt vergeleken met bestaand werk op het gebied van casusgebaseerd redeneren.

Het doel van het onderzoek was het beantwoorden van ruwweg twee categorieën onderzoeksvragen. De eerste categorie betreft de methode van het redeneren op basis van casusvergelijking:

- Hoe kan casusvergelijking worden geformaliseerd als criterium om te bepalen of een beslecht geval kan worden nagevolgd? Welke conclusies kunnen worden getrokken op basis van een vergelijking met beslechte gevallen? Volgens welke redeneerpatronen volgen die conclusies?

De tweede categorie van onderzoeksvragen betreft de relatie die het voorgestelde model heeft met bestaand onderzoek naar casusgebaseerd redeneren binnen Kunstmatige Intelligentie en Recht (KI & R):

- Hoe verhoudt zich de methode van redeneren op basis van casusvergelijking met ander onderzoek naar casusgebaseerd redeneren? Tot op welke hoogte wordt casusvergelijking geformaliseerd in dit andere werk? Wordt er rekenschap gegeven van de conclusies die volgen?

Beide categorieën van onderzoeksvragen zijn aan bod gekomen in de vijf hoofdstukken van dit proefschrift, waarvan de inhoud als volgt kort kan worden samengevat.

In Hoofdstuk 1 wordt de wetenschappelijke context geschetst waarbinnen het onderzoek is uitgevoerd. Het werk wordt gepositioneerd binnen het veld der Kunstmatige Intelligentie & Recht, een veld waarin een veelheid van juridische activiteiten wordt onderzocht vanuit het perspectief van de Kunstmatige Intelligentie, variërend van het nemen van rechterlijke beslissingen tot het schrijven van wetten.

In Hoofdstuk 2 worden enkele problemen bediscussieerd met betrekking tot casusgebaseerd redeneren in het recht, die zowel vanuit rechtstheoretisch alsook KI & R oogpunt bijzonder relevant zijn voor het hier besproken onderzoek. Twee methoden om vast te houden aan beslissingen worden gepresenteerd, namelijk casusvergelijking en regelextractie. Bij de casusverge-

lijkmingsmethode wordt de conclusie van een beslecht geval direct gekopieerd na een vergelijking met het voorliggende geval. Bij de regelextractiemethode wordt de beslissing eerst samengevat in een algemene regel, waarna de regel wordt toegepast op het voorliggende geval.

Het hoofdstuk gaat verder in op enkele relevante problemen die de casusvergelijkingsmethode opwerpt. Eén probleem is welke conclusies kunnen worden onderbouwd op basis van deze redeneermethode. Een ander probleem is langs welke redeneerpatronen men al dan niet tot de conclusies komt. Typische redeneerpatronen in casusvergelijking zijn het aanwijzen van relevante overeenkomsten tussen casus, alsook het maken van relevante onderscheidingen.

Een ander belangrijk probleem waarop Hoofdstuk 2 ingaat is, dat de uitkomsten van casusvergelijking afhankelijk zijn van de verzameling casuskenmerken die geacht worden relevant te zijn voor casusvergelijking. In het bijzonder blijkt het een belangrijk probleem hoe algemeen de termen moeten zijn waarin de relevante casuskenmerken dienen te worden omschreven.

In Hoofdstuk 3 wordt een formele theorie ontwikkeld over het redeneren op basis van casusvergelijking. Het belangrijkste doel van de theorie is het exact specificeren van de conclusies die volgen op basis van de casusvergelijkingsmethode. Een nevendoeel is het behandelen van redeneerpatronen typisch voor casusvergelijking, zoals het aanwijzen van overeenkomsten tussen casus, of het maken van onderscheidingen.

In dit hoofdstuk worden de formele condities opgesteld waaronder er sprake is van een analogie tussen een probleemgeval een reeds beslist geval. Intuïtief houden deze condities in dat het probleemgeval minstens even veel steun geeft aan de omstreden conclusie als het besliste geval. Dan kan, via een variant van redeneren a fortiori, de conclusie ook volgen in het probleemgeval.

Om te bepalen of een beslissing kan worden nagevolgd, wordt een vergelijking gemaakt van de argumentatie die mogelijk is in alle betrokken casus, voorzover deze argumentatie relevant is voor de omstreden conclusie. Deze relevante argumentatie in casus wordt expliciet gemaakt door middel van wat wordt genoemd een dialectisch argument. Dialectische argumenten bevatten niet enkel de gestelde feiten waarop de conclusies kunnen steunen, maar ook eventuele tussenconclusies. Bovendien neemt het model, in tegenstelling tot ander werk over casusvergelijking, niet voetstoots aan welke gedane beweringen redenen zijn voor of tegen conclusies. Het is dan ook mogelijk om met redenen te ondersteunen of aan te vallen dat de ene bewering de andere steunt of aanvalt, welke mechanismen resulteren in een soort verstrengeling van de dialectische argumenten. In het voorgestelde model wordt deze verstrengeling ook meegenomen in casusvergelijking, hetgeen duidelijk wordt uit het gebruik van een verstrengelde factorenhiërarchie.

Een belangrijk en bijzonder kenmerk van het huidige model is, dat het expliciet erkent dat het in het recht contingent is welke casuskenmerken relevant zijn voor casusvergelijking. De relevante casuskenmerken worden dan

ook expliciet gegeven in termen van wat wordt genoemd een vergelijkingsbasis, en de vergelijkingsuitkomsten worden aangegeven als relatief ten opzichte van die vergelijkingsbasis.

In Hoofdstuk 4 wordt de relatie bediscussieerd die het voorgestelde model heeft met ander onderzoek naar casusgebaseerd redeneren. Een discussiepunt is daarbij of de conclusies die volgen op basis van besliste gevallen worden gedefinieerd, een ander punt is of de redeneerpatronen worden behandeld langs welke de conclusies volgen. Een derde punt is of het vasthouden aan beslissingen behandeld wordt als regelextractie of als casusvergelijking. Zoals gezien wordt er nergens in het bediscussieerde onderzoek een formele behandeling gegeven van de casusvergelijkingsmethode die zowel de conclusies die volgen specificiert, alsook de redeneerpatronen waarlangs de conclusies volgen. Een vierde punt van discussie is of er expliciet wordt erkend dat het in het recht contingent is welke casuskenmerken relevant zijn voor casusvergelijking. Zoals gezien wordt er in het bestaande werk nergens expliciet met deze contingentie rekening gehouden.

In Hoofdstuk 5 wordt de theorie kritisch geëvalueerd. Een sterk punt van het model is dat de conclusies die volgen op basis van casusvergelijking worden gespecificeerd. In dit hoofdstuk komt ook een aantal redeneerpatronen aan bod die nog geen plaats hebben binnen het voorgestelde model. Een voorbeeld daarvan is het redeneren op basis van de hiërarchische relaties tussen rechters. De evaluatie gaat ook in op de vraag hoe zulke redeneerpatronen eventueel toch behandeld zouden kunnen worden binnen het voorgestelde model. Tenslotte wordt in dit hoofdstuk een aantal onderwerpen genoemd die voor toekomstig onderzoek interessant zijn, zoals het modelleren van het redeneren met hypothetische casus..

Na deze korte samenvatting van de inhoud van dit proefschrift, kunnen de belangrijkste bijdragen ervan als volgt worden opgesomd.

- Casusvergelijking wordt geformaliseerd als methode om beslissingen na te volgen. De uitkomsten van casusvergelijking worden formeel gedefinieerd, en de conclusies die volgen op basis van casusvergelijking worden aangegeven.
- De dialectische structuur van casus wordt expliciet gemaakt door middel van boom-achtige structuren die dialectische argumenten worden genoemd. In dialectische argumenten steunen beweringen elkaar, of vallen elkaar aan. Bovendien is het mogelijk om te steunen of aan te vallen dat de ene bewering de andere steunt of aanvalt, wat aanleiding geeft tot een soort verstrengeling van de dialectische argumenten. Het is een vernieuwend element van het huidige model dat deze verstrengeling ook een rol speelt in casusvergelijking, hetgeen aan het licht treedt in het gebruik van een verstrengelde factorenhierarchie.

- Het model erkent expliciet dat het in het recht contingent is welke casuskenmerken relevant zijn voor casusvergelijking. Dit wordt formeel bereikt door een vergelijkingsbasis te introduceren als een contingente set factoren die relevant zijn voor casusvergelijking, en door de uitkomsten van casusvergelijking te definiëren als relatief ten opzichte van de vergelijkingsbasis.
- Een aantal voor casusvergelijking typische redeneerpatronen krijgt ook een plaats binnen het model, te weten het aanwijzen van relevante overeenkomsten, het maken van onderscheidingen en het bagatelliseren en benadrukken van overeenkomsten of verschillen tussen casus. De laatste twee redeneerpatronen worden behandeld als een verandering in de vergelijkingsbasis, waarbij het vernieuwend is dat ook overeenkomsten kunnen worden gebagatelliseerd of benadrukt.

References

- Aleven, V. (1996). How Different Is Different? Arguing about the Significance of Similarities and Differences. *Proceedings of the Third European Workshop: EWCBR-96*, pp. 1-15.
- Aleven, V. and Ashley, K.D. (1997). Evaluating a Learning Environment for Case-Based Argumentation Skills. *Proceedings of the Sixth International Conference on Artificial Intelligence and Law*, pp. 170-179. New York: ACM Press.
- Aleven, V. (1997). *Teaching Case-Based Reasoning Through a Model and Examples*. Ph.D. dissertation University of Pittsburgh.
- Alexy, R. (1978). *Theorie der juristischen Argumentation*. Frankfurt am Main: Suhrkamp Verlag. (In German)
- Alexy, R. (1985). *Theorie der Grundrechte*. Baden-Baden: Nomos Verlagsgesellschaft. (In German)
- Ashley, K.D. (1989). Toward a Computational Theory of Arguing with Precedents: Accommodating Multiple Interpretations of Cases. *Proceedings of the Second International Conference on Artificial Intelligence and Law*, pp. 93-102. New York: ACM Press.
- Ashley, K.D. (1990). *Modeling Legal Argument: Reasoning with Cases and Hypotheticals*. Cambridge: MIT Press/Bradford Books.
- Ashley, K.D. (1991). Reasoning with cases and hypotheticals in HYPO. *International Journal of Man-Machine Studies*, Volume 34, pp. 753-796.
- Ashley, K.D. (1992). Case-Based Reasoning and its Implications for Legal Expert Systems. *Artificial Intelligence and Law*, Volume 1, pp. 113-208.
- Ashley, K.D. (2002). An AI model of case-based legal argument from a jurisprudential viewpoint. *Artificial Intelligence and Law*, Volume 10, pp. 163-218.
- Bankowski, Z. (1997). Rationales for Precedent. *Interpreting Precedents. A comparative study* (eds. MacCormick, D.N. and Summers, R.S.). Aldershot: Dartmouth Publishing Company Limited.
- Bareiss, E.R. (1988). *Exemplar-Based Knowledge Acquisition – a Unified Approach to Concept Representation, Classification, and Learning*. San Diego: Academic Press.
- Bench-Capon, T.J.M. (1999). Some Observations on Modelling Case Based Reasoning With Formal Argument Models. *Proceedings of the Seventh International Conference on Artificial Intelligence and Law*, pp. 36-42. New York: ACM Press.
- Bench-Capon, T.J.M. and Sartor, G. (2000). Using Values and Theories to Resolve Disagreement in Law. *Legal Knowledge and Information Systems: Jurix 2000* (eds. Breuker, J., Leenes, R. and Winkels, R.), pp. 73-84. Amsterdam: IOS Press.
- Bench-Capon, T.J.M. and Sartor, G. (forthcoming). A Model of Legal Reasoning with Cases Incorporating Theories and Values. To appear in: *Artificial Intelligence and Law* (Special Issue).
- Bench-Capon, T.J.M. and Rissland, E.L. (2001). Back to the Future: Dimensions Revisited. *Legal Knowledge and Information Systems: Jurix 2001* (eds. Verheij, B., Lodder, A.R., Loui, R. and Muntjewerff, A.J.), pp. 41-52. Amsterdam: IOS Press.
- Bench-Capon, T.J.M. and Sartor, G. (2001). Theory Based Explanation of Case Law Domains. *Proceedings of the Eighth International Conference on Artificial Intelligence and Law*, pp. 12-21. New York: ACM Press.

- Bench-Capon, T.J.M. (2002). Representation of Case Law as an Argumentation Framework. *Legal Knowledge and Information Systems: Jurix 2002* (eds. Bench-Capon, T.J.M., Daskalopulu, A. and Winkels, R.), pp. 103-112. Amsterdam: IOS Press.
- Bench-Capon, T.J.M. and Sartor, G. (2003). A Model of Legal Reasoning with Cases Incorporating Theories and Values. To appear in: *Artificial Intelligence and Law* (Special Issue).
- Branting, L.K. (1991). Building explanations from rules and structured cases. *International Journal of Man-Machine Studies*, Volume 34, pp. 797-837.
- Branting, L.K. (1994). A Computational Model of Ratio Decidendi. *Reasoning with Rules and Precedents. Artificial Intelligence and Law*, Volume 2, pp. 1-31.
- Branting, L.K. (2000). *Reasoning with Rules and Precedents. A Computational Model of Legal Analysis*. Dordrecht: Kluwer Academic Publishers.
- Brüninghaus, S. and Ashley, K.D. (1999a). Toward Adding Knowledge to Learning Algorithms for Indexing Legal Cases. *Proceedings of the Seventh International Conference on Artificial Intelligence and Law*, pp. 9-17. New York: ACM Press.
- Brüninghaus, S. and Ashley, K.D. (1999b). Bootstrapping Case-Based Development with Annotated Case Summaries. *Proceedings of the Third International Conference on Case-Based Reasoning*, pp. 59-73. Berlin: Springer Verlag.
- Brüninghaus, S. and Ashley, K.D. (2003). Predicting Outcomes of Case-Based Legal Arguments. *Proceedings of the Ninth International Conference on Artificial Intelligence and Law*, pp. 233-242. New York: ACM Press.
- Burnham, W. (1995). *Introduction to the law and legal system of the United States*. St. Paul: West Publishing Co.
- Burton, S.J. (1985). *An introduction to law and legal reasoning*. Boston: Little, Brown and Company.
- Combrink-Kuiters, C.J.H. (1998). *Kennis van Zaken*. Ph.D. dissertation Erasmus Universiteit Rotterdam: Sanders Instituut/Gouda Quint.
- Crombag, H.F.M. (1994). On the artificiality of artificial intelligence. *Artificial Intelligence and Law*, Volume 2, pp. 39-49.
- Cross, R. (1977). *Precedent in English Law*. Third edition. Oxford University Press.
- Dancy, J. (1985). *Introduction to contemporary epistemology*. Oxford: Basil Blackwell Limited.
- Debaene, S., Van Kuyck, R. and Van Buggenhout, B. (1999). Legislative Techniques as Basis of a Legislative Drafting System. *Proceedings of the Twelfth Annual Jurix Conference* (eds. Van den Herik, H.J., Moens, M.F., Bing, J., Van Buggenhout, B., Zeleznikow, J., Grüters, C.A.F.M.), pp. 23-34.
- Dworkin, R. (1986). *Law's Empire*. Cambridge, Massachusetts: The Belknap Press of Harvard University Press.
- Dworkin, R. (1977). *Taking Rights Seriously*. Fourth impression. London: Gerald Duckworth & Co.
- Gamut, L.T.F. (1991). *Logic, Language and Meaning. Volume I: Introduction to Logic*. Chicago: The University of Chicago Press.
- Golding, M.P. (2001). *Legal Reasoning*. Ontario: Broadview Press.
- Goodhart, A. (1930). Determining the *ratio decidendi* of a case. *Yale Law Review*, Volume 40, pp. 161f.

- Goodhart, A. (1959). The ratio decidendi of a case. *The Modern Law Review*, Volume 22, no. 5, pp. 117f.
- Hage, J.C. (1996). A Theory of Legal Reasoning and a Logic-to-Match. *Artificial Intelligence and Law*, Volume 4, pp. 199-273.
- Hage, J.C. (1997). *Reasoning with Rules. An Essay on Legal Reasoning and Its Underlying Logic*. Dordrecht: Kluwer Law and Philosophy Library.
- Hage, J.C. (2000). Goal-Based Theory Evaluation. *Legal Knowledge and Information Systems. Jurix 2000: The Thirteenth Annual Conference*, pp. 59-72.
- Hage, J.C. (2001). Formalizing legal coherence. *Proceedings of the Eighth International Conference on Artificial Intelligence and Law*, pp. 22-31. New York: ACM Press.
- Hage, J.C., Leenes, R. and Lodder, A.R. (1994). Hard Cases: A Procedural Approach. *Artificial Intelligence and Law*, Volume 2, pp. 113-167.
- Hage, J.C., Schölsels, R.J.N. and Wolleswinkel, R. (2001). *Recht, vaardig en zeker. Een inleiding in het recht*. Den Haag: Boom Juridische Uitgevers. (In Dutch)
- Hammond, K.J. (1989). *Case-Based Planning – Viewing Planning as a Memory Task*. San Diego: Academic Press.
- Hart, H.L.A. (1983). Positivism and the Separation of Law and Morals. *Essays in Jurisprudence and Philosophy*, Essay 2, pp. 49-87. Oxford: Clarendon Press.
- Hart, H.L.A. (1994). *The Concept of Law*. Second edition. New York: Oxford University Press.
- Hare, R.M. (1963). *Freedom and Reason*. Oxford University Press.
- Henderson, J. and Bench-Capon, T.J.M. (2000). A model of the development of distinctions in case law. *Legal Knowledge and Information Systems: Jurix 2000* (eds. Breuker, J., Leenes, R. and Winkels, R.), pp. 23-34. Amsterdam: IOS Press.
- Henderson, J. and Bench-Capon, T.J.M. (2001). Dynamic Arguments in a Case Law Domain. *Proceedings of the Eighth International Conference on Artificial Intelligence and Law*, pp. 60-69. New York: ACM Press.
- Herik, H.J. van den (1991). *Kunnen computers rechtspreken?* Arnhem: Gouda Quint. (In Dutch)
- Kolodner, J. (1993). *Case-Based Reasoning*. San Mateo: Morgan Kaufmann Publishers Inc.
- Koton, P. (1988). *Using Experience in Learning and Problem Solving*. Ph.D. dissertation MIT, Department of Computer Science.
- Levi, E.H. (1949). *An Introduction to Legal Reasoning*. University of Chicago Press.
- Llewellyn, K.N. (1960). *The Common Law Tradition*. Boston: Little, Brown and Company.
- Llewellyn, K.N. (1989). *The case law system in America*. University of Chicago Press.
- Lodder, A.R. (1998). *DiaLaw: On legal justification and dialog games*. Ph.D. dissertation Universiteit Maastricht.
- Lodder, A.R. and Huygen, P.E.M. (2001). A Simple Tool to Structure the Information Exchange between Parties in Online Alternative Dispute Resolution. *Legal Knowledge and Information Systems: Jurix 2001* (eds. Verheij, B., Lodder, A.R., Loui, R. and Muntjewerff, A.J.), pp. 117-129. Amsterdam: IOS Press.
- Loui, R.P. (1997). Eliding the Arguments of Cases. *Draft*. Downloaded from <http://www.cs.wustl.edu/~loui/ratjur.txt>.
- Loui, R.P. and Norman, J. (1995). Rationales and Argument Moves. *Artificial Intelligence and Law*, Volume 3, pp. 159-189.

- MacCormick, D.N. (1987). Why Cases have *Rationes* and what these are. In: *Precedent in Law* (ed. Goldstein, L.), pp. 155-182. Oxford: Clarendon Press.
- MacCormick, D.N. and Summers, R.S. (eds.) (1995). *Interpreting Precedents. A comparative study*. Aldershot: Dartmouth Publishing Company Limited.
- MacCormick, D.N. (1997). *Legal Reasoning and Legal Theory*. Oxford: Clarendon Press.
- Matthijssen, L. (1999). *Interfacing between Lawyers and Computers. An Architecture for Knowledge-based Interfaces to Legal Databases*. Ph.D. dissertation Katholieke Universiteit Brabant.
- Mommers, L. (2002). *Applied legal epistemology. Building a knowledge-based ontology of the legal domain*. Ph.D. dissertation Universiteit Leiden.
- Nieuwenhuis, J.H. (1976). Legitimatie en heuristiek van het rechterlijk oordeel. *Rechtsgeleerd Magazijn Themis*, pp. 494-515. (In Dutch)
- Oliphant, H. (1928). A return to stare decisis. *American Bar Association Journal*, Volume 14. Reprint in: *American Legal Theory* (ed. Summers, R.S. 1992), pp. 143-153. Aldershot: Dartmouth Publishing Company Limited.
- Oskamp, E.W. (1998). *Computerondersteuning bij straftoemeting*. Ph.D. dissertation Rijksuniversiteit Leiden. Deventer: Gouda Quint. (In Dutch)
- Oskamp, A. and Lodder, A.R. (2002). *Informatietechnologie voor juristen*. Deventer: Kluwer. (In Dutch)
- Peczenik, A. (1997). The Binding Force of Precedent. *Interpreting Precedents. A comparative study* (eds. MacCormick, D.N. and Summers, R.S.), pp. 461-480. Aldershot: Dartmouth Publishing Company Limited.
- Peczenik, A. and Hage, J.C. (2000). Legal knowledge about what? *Ratio Juris*, Volume 13, pp. 325-345.
- Pollock, J.L. (1987). Defeasible Reasoning. *Cognitive Science*, Volume 11, pp. 481-518.
- Pollock, J.L. (1995). *Cognitive Carpentry: A Blueprint for How to Build a Person*. Cambridge: MIT Press.
- Prakken, H. (1993). *Logical Tools for Modelling Legal Argument*. Ph.D. dissertation Vrije Universiteit Amsterdam.
- Prakken, H. (1997). *Logical Tools for Modelling Legal Argument. A Study of Defeasible Reasoning in Law*. Dordrecht: Kluwer Law and Philosophy Library.
- Prakken, H. (2000). An exercise in formalising teleological case-based reasoning. *Legal Knowledge and Information Systems: Jurix 2000* (eds. Breuker, J., Leenes, R. and Winkels, R.), pp. 49-57. Amsterdam: IOS Press.
- Prakken, H. and Sartor, G. (1996). A Dialectical Model of Assessing Conflicting Arguments in Legal Reasoning. *Artificial Intelligence and Law*, Volume 4, pp. 331-368.
- Prakken, H. and Sartor, G. (1997a). Argument-based extended logic programming with defeasible priorities. *Journal of Applied Non-classical Logics*, Volume 7, pp. 25-75.
- Prakken, H. and Sartor, G. (1997b). Reasoning with Precedents in a Dialogue Game. *Proceedings of the Sixth International Conference on Artificial Intelligence and Law*, pp. 1-9. New York: ACM Press.
- Prakken, H. and Sartor, G. (1998). Modelling Reasoning with Precedents in a Formal Dialogue Game. *Artificial Intelligence and Law*, Volume 6, pp. 231-287.

- Prakken, H. (2000). An exercise in formalising teleological case-based reasoning. *Legal Knowledge and Information Systems: Jurix 2000* (eds. Breuker, J., Leenes, R. and Winkels, R.), pp. 49-57. Amsterdam: IOS Press.
- Raz, J. (1979). *The Authority of Law*. Oxford University Press.
- Rissland, E.L. and Ashley, K.D. (1987). A Case-Based System for Trade Secrets Law. *Proceedings of the First International Conference on Artificial Intelligence and Law*, pp. 60-66. New York: ACM Press.
- Rissland, E.L. and Ashley, K.D. (2002). A note on dimensions and factors. *Artificial Intelligence and Law*, Volume 10, pp. 65-77.
- Rissland, E.L. and Skalak, D.B. (1991). CABARET: rule interpretation in a hybrid architecture. *International Journal of Man-Machine Studies*, Volume 34, pp. 839-887.
- Rissland, E.L., Skalak, D.B. and Friedman, M.T. (1996). BankXX: Supporting Legal Arguments through Heuristic Retrieval. *Artificial Intelligence and Law*, Volume 4, pp. 1-71.
- Roth, A.C. (1999a). Teaching Case-Based Reasoning Through a Model and Examples. Book review of *Teaching Case-Based Argumentation Through a Model and Examples*, Ph.D. dissertation by Vincent Aleven, University of Pittsburgh, 1997. Review in: *BNVKI newsletter*, Volume 16, no. 2, 1999, pp. 47-52.
- Roth, A.C. (1999b). 'Zomerweelde: a court for practising procedural criminal law'. Lecture by Jürgen Wöretshofer, held on June 2nd 1999 in Maastricht. Report in: *BNVKI newsletter*, Volume 16, no. 3, 1999, pp. 76-77.
- Roth, A.C. (2000a). 'The 'E-T11' project: an automated aid for assessing the quality of statutory regulations'. Lecture by Luuk Matthijssen, held on May 15th 2000 in Tilburg. Report in: *BNVKI newsletter*, Volume 17, no. 3, 2000, pp. 69-70.
- Roth, A.C. (2000b). New reasoning patterns in analogical legal case-based reasoning: an informal investigation. *Proceedings of the Thirteenth Annual Jurix Conference* (eds. Breuker, J., Leenes, R. and Winkels, R.), 113-122.
- Roth, A.C. (2001a). A Typology of Moves Involved in Case Comparison. *Proceedings of the 8th International Conference on Artificial Intelligence and Law*, pp. 129-138.
- Roth, A.C. (2001b). 'Possible Applications of Argumentation Research'. Lecture by Henry Prakken, held on May 11th 2001 in Utrecht. Report in: *BNVKI newsletter*, Volume. 18, no. 3, 2001, pp. 71-72.
- Roth, A.C. (2001c). The Dialectics of Case Comparison: a Formal Reconstruction. *Proceedings of the Fourteenth Annual Jurix Conference* (eds. Verheij, B., Lodder, A.R., Loui, R.P. and Muntjewerff, A.J.), 29-40.
- Roth, A.C. (2002). 'Extended Conceptual Retrieval'. Lecture by Radboud Winkels, held on January 29th 2002 in Amsterdam. Report in: *BNVKI newsletter*, Volume 19, no. 1, 2002, pp. 19-20.
- Russell, S.J. and Norvig, P., eds. (2003). *Artificial Intelligence: A Modern Approach*. Upper Saddle River, New Jersey: Pearson Education, Inc.
- Simpson, A.W.B. (1959). The *ratio decidendi* of a case. *The Modern Law Review*, Volume 22, no. 5, pp. 453f.
- Skalak, D.B. and Rissland, E.L. (1992). Arguments and Cases: An inevitable Intertwining. *Artificial Intelligence and Law*, Volume 1, pp. 3-44.
- Span, G.P.J. (1992). *LITES: een intelligent tutorsysteem voor juridisch onderwijs*. Ph.D. dissertation Universiteit Maastricht. (In Dutch)

- Span, G.P.J. (1993). LITES: an Intelligent Tutoring System for Legal Problem Solving in the Domain of Dutch Civil Law. *Proceedings of the Fourth International Conference on Artificial Intelligence and Law*, pp. 76-81. New York: ACM Press.
- Stone, J. (1947). *The province and function of law*. London: Stevens & Sons Limited.
- Stone, J. (1968). *Legal system and lawyers' reasonings*. Stanford University Press.
- Summers, R.S. (1997). Precedent in the United States (New York State). *Interpreting Precedents. A comparative study* (eds. MacCormick, D.N. and Summers, R.S.). Aldershot: Dartmouth Publishing Company Limited.
- Toulmin, S.E. (1958). *The Uses of Argument*. London: Cambridge University Press.
- Valente, A. (1995). *Legal Knowledge Engineering. A modelling Approach*. Ph.D. dissertation University of Amsterdam.
- Verheij, B. (1996). *Rules, Reasons, Arguments. Formal studies of argumentation and defeat*. Ph.D. dissertation Universiteit Maastricht.
- Verheij, B. (1999a). Automated Argument Assistance for Lawyers. *Proceedings of the Seventh International Conference on Artificial Intelligence and Law*, pp. 43-52. New York: ACM Press.
- Verheij, B. (1999b). Logic, Context and Valid Inference. Or: Can there be a Logic of Law? *Proceedings of the Twelfth Annual Jurix Conference* (eds. Van den Herik, H.J., Moens, M.F., Bing, J., Van Buggenhout, B., Zeleznikow, J., Grütters, C.A.F.M.), pp. 109-121.
- Verheij, B. (2000a). Book review: Formalism and interpretation in the logic of law. *Artificial Intelligence and Law*, Volume 8, pp. 35-65.
- Verheij, B. (2000b). Dialectical Argumentation as a Heuristic for Courtroom Decision-Making. In: *Rationality, Information and Progress in Law and Psychology* (eds. Koppen, P.J. van and Roos, N.H.M.), pp. 203-226. Maastricht: Universitaire Pers.
- Verheij, B. (2000c). DEFLOG - a logic of dialectical justification and defeat. *Draft*. Downloaded from <http://www.metajur.unimaas.nl/~bart/publications.htm>.
- Verheij, B. (2003). DefLog: on the Logical Interpretation of Prima Facie Justified Assumptions. *Journal of Logic and Computation*, Volume 13, no. 3, pp. 319-346.
- Verheij, B. and Hage, J.C. (1994). Reasoning by analogy: a formal reconstruction. *Legal knowledge based systems. The relation with legal theory* (eds. Prakken, H., Muntjewerff, A.J. and Soeteman, A.), pp. 65-78. Lelystad: Koninklijke Vermande.
- Vey Mestdagh, C.N.J. de (1997). *Juridische Kennissystemen. Rekentuig of Rekenmeester?* Ph.D. dissertation Rijksuniversiteit Groningen. Deventer: Kluwer. (In Dutch)
- Wiarda, G.J. (1999). *Drie typen van rechtsvinding*. Zwolle: Tjeenk Willink. (In Dutch)
- Winkels, R., Bosscher, D., Boer, A., Hoekstra, R. (2000). Extended conceptual retrieval. *Proceedings of the Thirteenth Annual Jurix Conference* (eds. Breuker, J., Leenes, R. and Winkels, R.), 85-97.
- Zeleznikow, J. and Stranieri, A. (1995). The Split Up system: Integrating neural networks and rule based reasoning in the legal domain. *Proceedings of the Fifth International Conference on Artificial Intelligence and Law*, pp. 185-194. New York: ACM Press.
- Zeleznikow, J. and Stranieri, A. (1997). Knowledge Discovery in the Split-Up Project. *Proceedings of the Sixth International Conference on Artificial Intelligence and Law*, pp. 89-97. New York: ACM Press.

Index

- adherence to decisions, 8
- AI & Law, 1, 18, 113, 158, 162
- Aleven, 10, 115
- Alexy, 2
- analogue. *See* analogising and distinguishing
- analogue. *See* analogising and distinguishing, 86
- analogy: between cases, 11; two perspectives on, 18
- arguments, 104, 116, 125
- Ashley, 10, 103
- attack. *See* support and attack
- background, 41
- backing: of warrant. *See* warrant
- Bankowski, 8
- Bench-Capon, 17, 134
- Branting, 140
- Brüninghaus, 159
- Burton, 11
- case: hypothetical, 98, 108; problem, 21, 69; representation, 40; settled, 21, 69
- case base, 71
- CATO, 115
- coherentism, 101
- common law, 7
- comparison basis, 50; contingency of, 52
- comparison outcomes: definition, 65; invariance under negation, 67; inversion under negation, 67; transitivity, 67; types of, 63
- conclusions: CBR-derivable, 77; CF-derivable, 44; labelled, 83
- conflict: of precedents, 72; of reasons, 38; resolution, 15
- connective: conjunction, 99; disjunction, 99; negation, 39; support, 38
- Cross, 11
- dialectical arguments, 30; in case comparison, 47
- dialectical support, 47
- dialogues, 2, 124
- discretion, 131
- distinguishing. *See* analogising and distinguishing
- downplaying and emphasising, 86, 118
- Dworkin, 2
- entangled factor hierarchy, 53, 118
- entanglement, 31, 77, 86, 148
- exception: to stare decisis, 22, 164; undercutting, 31
- exemplar, 140
- fact, 38
- factor: applying, 56; basic, 51; grounded, 56; relevant, 56
- Golding, 8
- Goodhart, 12
- GREBE, 140
- Hage, 17, 146
- Hart, 1
- heuristic criteria, 122
- HYPO, 103
- Kolodner, 1, 71
- language. *See* case representation
- legal: realism, 9; security, 8; theory, 1
- Levi, 11
- Llewellyn, 1
- Lodder, 2
- MacCormick, 7, 9, 98
- material implication, 39
- modus ponens, 39; closure, 44
- negation. *See* connective
- non-monotonic: inferences, 39; logics, 2
- Oliphant, 9
- on pointness, 91, 110, 131
- Pollock, 31, 159
- Prakken, 1, 15, 124
- precedents, 69
- premise, 42, 77

priority: between rules, 17, 125
procedural law, 162
ratio decidendi, 9, 12, 140
Raz, 100
reason, 38
reasoning: a fortiori, 19; by case
 comparison, 77; patterns of, 14,
 86; with support and attack, 42,
 77
rebuttal, 124, 136
relevance: for case comparison, 20,
 50; of settled cases, 68; relation
 of, 56
Rissland, 3, 11, 103, 109
rule: extraction, 9, 151
rules, 146
Sartor, 124, 134
sentence: atomic, 38; conditional,
 38
Skalak, 3
stare decisis: departures from, 22,
 74, 163; reasons for, 8
state of affairs, 37
statement: compound, 38;
 elementary, 38; opposite, 39
statutory law, 7, 41
Stone, 9, 13
support. *See* support and attack
support and attack, 37
theories, 134
Toulmin, 31
values, 134
Verheij, 2, 30, 31
warrant: backing of, 31; reduction
 graph, 141
weighing: reasons, 17, 99;
 similarities against differences,
 19
Wiarda, 7

Curriculum Vitae

Bram Roth was born in The Hague on March 11th, 1968. From 1980 till 1986 he attended secondary school, after which he started studying Technical Physics at the Technical University of Eindhoven. After an interruption of a few years, he took up his studies in 1993 and graduated in 1995, "with appreciation from the Faculty".

After his graduation he followed an academic course at the Catholic University of Nijmegen, to become a physics teacher. He got his teaching qualification in 1997. Until 1998 he worked as a physics teacher at several schools for secondary education.

In 1998 he started working as a researcher and teacher at the Faculty of Law of the Universiteit Maastricht. Since then he has published regularly on his research subject case-based reasoning in the law, whilst taking part as a teacher in both theoretical and practical education.

