Chapter 9
A Logical Analysis of Burdens of Proof

Henry Prakken and Giovanni Sartor

1. Introduction

The legal concept of burden of proof is notoriously complex and ambiguous. Various kinds of burdens of proof have been distinguished, such as the burden of persuasion, burden of production and tactical burden of proof, and these notions have been described by different scholars in different ways. They have also been linked in various ways with notions like presumptions, standards of proof, and shifts and distributions of burdens of proof. What adds to the complexity is that different legal systems describe and treat the burden of proof in different ways. For instance, in common law jurisdictions the just-mentioned distinction between three kinds of burden of proof is explicitly made while in civil law systems it usually remains implicit.

This chapter aims to clarify matters concerning burden of proof from a logical point of view. We take a logical point of view since, although some differences in notions and treatments might reflect legitimate *legal* differences between jurisdictions, we think that to a large extent the burden of proof is an aspect of *rational* thinking and therefore subject to a logical analysis. In particular, we claim that the burden of proof can be adequately analysed in terms of logical systems for defeasible argumentation, i.e., logics for fallible (but not fallacious!) reasoning. The grounds for this claim are fourfold. First, since legal proof almost always has an element of uncertainty, we cannot impose a deductive form onto real legal evidential reasoning. Second, while this reason still leaves open the use of other approaches, such as story-based or statistical approaches, we think that the notion of argumentation and related notions such as counter-argument, rebuttal and dispute, are very natural to legal thinking. Third (and a special case of the second reason), logics for defeasible argumentation are arguably suitable as a formal underpinning of much work of the influential New Evidence scholars, such as Anderson, Tillers, Twining and Schum (e.g. Anderson et al. 2005, who revived and modernized Wigmore's famous charting method for making sense of legal-evidential problems. (See Prakken 2004 for a defence of the thesis that argument-based logics can be a formal underpinning of this work.) Finally, logics for defeasible argumentation have a firm theoretical basis both in philosophy (especially in argumentation theory) and in logic (especially in its applications in artificial intelligence (AI)). In short, logics for defeasible argumentation are

theoretically well-founded and mature analytical tools that fit well with legal thinking in general and with legal-evidential reasoning in particular.

We start in section 2 with an overview of how the various kinds of burden of proof and related notions have been described and related in the jurisprudential literature. Then we describe the logical background of our analysis in section 3, that is, the idea of logics for defeasible argumentation and how they can be embedded in models of legal procedure. We have deliberately made this section of a tutorial nature, since a secondary aim of this chapter is to introduce these logics to the legal-jurisprudential community. Section 4 forms the heart of the chapter: it contains our formal account of the various notions of burden of proof in terms of argumentation logics. In this section we also discuss to what extent shifts and distributions of the three kinds of burden of proof can be logically modelled, and how our logical model accounts for different proof standards. In section 5 we briefly discuss some other notions related to burden of proof that are sometimes distinguished in the law. We conclude in section 6 by discussing to what extent our analysis can be adapted to other approaches, such as statistical ones.

2. Doctrinal Discussions on the Burdens of Proof in Civil Law and Common Law

In this section we briefly discuss accounts of burden of proof in doctrinal analysis, both in civil law and in common law jurisdictions. In common law systems generally a clear distinction is made between the burden of production and the burden of persuasion, although different characterizations and denominations are used for this two kinds of burden (see Williams 2003). What we call the *burden of production* is characterized by Capper and Cross (1999, 113) as 'the obligation to show, if called upon to do so, that there is sufficient evidence to raise an issue as to the existence or non-existence of a fact in issue'. Strong (1992, 425) describes the burden of production as 'the liability to an adverse ruling (generally a finding or directed verdict) if evidence on the issue has not been produced'. This burden is also called the 'evidential burden' (Capper and Cross 1990, 113), or the 'duty to produce evidence' (Wigmore 1940, § 2487), or the duty of passing the judge (Keane 1994, 55), or the burden of adducing evidence (Zuckerman 2006, para. 21.35). What we call the *burden of persuasion* (Zuckerman 2006, para. 21.33) is characterized by Capper and Cross (1990, 113) as the 'obligation of a party to meet the requirement of a rule of law that a fact in issue must be proved or disproved', and by Strong (1992, 426) as meaning that if the party having that burden has failed to satisfy it, the issue is to be decided against that party. This burden is also called the 'legal burden' (Denning 1945, Capper and Cross 1990, 113), the 'risk of non-persuasion' (Wigmore 1940, § 2487) or 'probative burden' (*DPP* v. *Morgan* [1976] AC 182 at 209, Lord Hailsham).

The *proof standards* for these two burdens are quite different. For the burden of persuasion the fact-finder must be convinced that the statement holds 'beyond

reasonable doubt' (in criminal cases) or 'on the balance of probabilities' (in civil cases; in such cases the phrase 'more probable than not' is also used). For the burden of production the proof standard is much lower. Sometimes it is said that just a 'scintilla of evidence' is needed, sometimes that the evidence is such that 'reasonable minds can disagree' on the issue, or even that there is evidence 'upon which a jury can properly proceed to find a verdict for the party producing it, upon whom the onus of proof is imposed' (as required in *Improvement Co.* v. *Munson*, 14 Wall. 442, 81 US 448 [1872]).

The distinction between a burden of production and a burden of persuasion is more significant in common law jurisdictions, since in these systems the discharge of the burden of production is a precondition for moving to the trial phase, where the factual issue is decided by the jury according to the burden of persuasion. Accordingly, these burdens are verified at different moments. According to Wigmore (1962, Volume IX at 283, cited in Williams 2003) 'The risk of non-persuasion operates when the case has come into the hands of the jury, while the duty of producing evidence implies a liability to a ruling by the judge disposing of the issue without leaving the question open to the jury's deliberations.' Strong (1992, 426) says:

> The burden of persuasion becomes a crucial factor only if the parties have sustained their burdens of producing evidence and only when all of the evidence has been introduced. It does not shift from party to party during the course of the trial simply because it need not be allocated until it is time for a decision. When the time for a decision comes, the jury, if there is one, must be instructed how to decide the issue if their minds are left in doubt. The jury must be told that if the party having the burden of persuasion has failed to satisfy that burden, the issue is to be decided against that party. If there is no jury and the judge is in doubt, the issue must be decided against the party having the burden of persuasion.

However, the distinction between the two burdens is also recognized in some civil law jurisdictions. For instance, the German legal doctrine distinguishes between a subjective burden of proof (*subjektive Beweislast*, also called burden of providing a proof, *Beweisführungslast*) and objective burden of proof (*objektive Beweislast*). As observed by Hahn and Oaksford (2007), the first corresponds, more or less, to the burden of production and the second to the burden of persuasion (see, for instance, Rosenberg et al. 1993, §§ 112–24).

The relation between the burdens of persuasion and production depends on whether the case is a criminal one or a civil one. For civil cases they usually go together since both are usually determined by the 'operative facts' for a legal claim, i.e., the facts that legally are ordinarily sufficient reasons for the claim. The law often designates the operative facts with rule–exception structures. For example, the operative facts for the existence of a contract generally are that there was an offer which was accepted but this rule can have many exceptions, such as that one party deceived the other party or that the party making or accepting the

offer was insane when doing so. Now, in civil cases, the general rule is that the party who makes a legal claim has both the burden of production and the burden of persuasion for the operative facts of the claim, while the other party has the two burdens for any exception.

For instance, if the plaintiff claims that a contract between him and defendant exists then he must produce evidence that he made an offer which the defendant accepted to fulfil his burden of production, and in the final stage the fact-finder must regard it as more probable than not that this offer and acceptance were made, otherwise the plaintiff loses. Suppose the plaintiff succeeds in both tasks and that the defendant claims she was insane when she accepted his offer. Then if the defendant has not produced evidence for her insanity, the plaintiff wins since the judge must rule as a matter of law that she was not insane. However, if she did produce evidence for her insanity, then she only wins if the fact-finder regards it as more probable than not that she was insane, otherwise the plaintiff still wins, even if the evidence on insanity is balanced.[1]

In criminal cases the burdens of production and persuasion on an issue can be on different parties, since the principle according to which one cannot be convicted unless one's guilt is proved also covers the non-existence of exceptions preventing such guilt. More precisely, this principle implies that the accused has to be acquitted when there remains reasonable doubt concerning the existence of such an exception (for instance, self-defence in a murder case), so the prosecution also has the burden of persuasion for the non-existence of such exceptions. In other words, in criminal cases the prosecution has the burden of persuasion not only for the legal operative facts for a claim (say, for murder, that there was a killing and that it was done with intent) but also for the non-existence of exceptions (such as that the killing was not done in self-defence). However, for the burden of production this is different: the prosecution has this burden only for the legal operative facts (in our murder example, 'killing' and 'intent'); for the exceptions the burden of production is on the defence. As Spencer and Spencer (2007, Chapter 2) say, in the British legal system:

> the prosecution cannot be expected to put up evidence to anticipate every specific defence the accused may present; thus in order to plead self-defence the accused will have to provide some evidence to enable the court to consider the matter.

So in our murder case example the defence must produce evidence that he acted in self-defence but once he has produced such evidence, the prosecution has the burden of persuasion that there was no self-defence. Similarly, in the Italian legal

1 It should be noted that our observations in this chapter on allocating the proof burdens only hold as a general rule; legal systems leave some freedom to courts to make exceptions to them in special cases, for instance, on the basis of fairness. The allocation of proof burdens can even be the subject of dispute; see Prakken and Sartor (2007) for examples and a logical formalization of such disputes.

system, the accused has the burden of producing evidence sufficient to create such a doubt in respect of the existence of a cause of justification, while the prosecutor then has the burden of persuading the court that the cause of justification does not exist (see Tonini 2007, 311, who grounds this conclusion on article 530 of the Italian code of criminal procedure, specifying that the judge has to acquit the accused person in a case where there is doubt in respect of the existence of a cause of justification). In sum, in criminal proceedings the two burdens go together only for operative facts; for exceptions they are separated.

As Williams (2003) observes, common law doctrine usually does not clearly distinguish the burden of production from the so-called 'tactical burden', which he characterizes as the situation when, if the party does not produce evidence or further evidence he or she runs the risk of ultimately losing in respect of that issue. The same criticism is raised by other authors, such as Keane (1994), according to whom, by providing evidence a party does not shift the burden of production onto the other party but only shifts a tactical burden, since once the burden of production is fulfilled, the issue is determined regardless of the burden of production. By contrast, the tactical burden of proof is not allocated by law but induced by the defeasible nature of the reasoning and the estimated quality of the evidence and arguments produced so far. In civil law countries, to the best of our knowledge, the distinction between burden of production and tactical burden is not usually explicitly considered. Nevertheless, since this notion is induced by the logic of the reasoning process instead of being assigned by law, it is also relevant for these systems. In the words of Williams (2003), this burden is a matter of tactical evaluation in that a party must assess the risk of losing in respect of an issue if no further evidence concerning that issue is produced.

Suppose in our murder example that the prosecution has provided evidence for 'killing' and 'intent', after which the defence produced evidence for 'self-defence'. The prosecution must now assess the risk of losing if the current stage were the final stage. If this risk is real then the prosecution had better provide counter-evidence on 'self-defence'. In other words, the prosecution now not only has the burden of persuasion but also a tactical burden with respect to self-defence. Clearly, a tactical burden can shift between the parties any number of times during a proceeding, depending on who would be likely to win if no more evidence were provided. In our example, if the prosecution provides counter-evidence against self-defence, then the defence must estimate the likelihood of losing if it does not provide further evidence supporting self-defence. If this likelihood is real, then the tactical burden has shifted to the defence. By contrast, the burden of production never shifts since once fulfilled it is disregarded in the rest of the proceeding. With Williams (2003), we believe that those who argue that this burden can shift confuse it with the tactical burden.

The tactical burden is also relevant in civil cases, for instance, when an exception is not to a legal rule but to a commonsense generalization. Suppose in our contract example that the defendant disputes the plaintiff's claim that the defendant accepted his offer, and that the plaintiff provides two witness testimonies

in support of his claim. The commonsense generalization used by the plaintiff here is that if two witnesses say the same, they are usually telling the truth. If the defendant provides no counter-evidence to the witnesses' credibility, then she runs the risk of losing in respect of this issue, since the fact-finder is likely to accept this generalization. However, her burden to provide such counter-evidence is not a burden of persuasion but only a tactical burden: since the plaintiff has the burden of persuasion for his claim that the defendant accepted his offer, the defendant's task is to cast sufficient doubt on whether she accepted the offer; she does not have to persuade the fact-finder that she did not accept the offer. Thus in civil cases the nature of an exception is important: if it is an exception to a legal rule, then it carries the burdens of production and persuasion, while if it is an exception to a commonsense generalization, it only carries a tactical burden, the strength of which depends on whether the generalization is used to fulfil a burden of persuasion or to prevent such fulfilment.

Both the murder and the contract example show that the tactical burden has no single fixed proof standard. A tactical burden can be said to be fulfilled if its intended effect is made likely, and this effect is different depending on whether a party has the burden of persuasion or not. The party that has it must convince the fact-finder (to the relevant degree) that the statement on which it rests holds while the other party only needs to make the fact-finder doubt (to the relevant degree) whether the statement holds.

Summarizing, we distil the following characterizations from the above discussion. The burden of persuasion specifies which party has to prove a statement to a specified degree (its proof standard) with the penalty of losing in respect of the issue. Whether this burden is met is determined in the final stage of a proceeding, after all evidence is provided. That a burden of persuasion for a statement is fulfilled means that a rational fact-finder is, to the required degree, convinced that the statement is true; so if the burden is not met, this means that such a fact-finder is not convinced to that degree that the statement is true; he need not be convinced that it is false. The burden of production specifies which party has to offer evidence on an issue at different points in a proceeding. If such evidence does not meet the (low) proof standard for this burden, the issue is decided as a matter of law against the burdened party, while otherwise the issue is decided in the final stage by the trier of fact according to the burden of persuasion. Both these burdens are assigned as a matter of law. By contrast, the *tactical burden of proof* is a matter of tactical evaluation in that a party must assess the risk of ultimately losing in respect of an issue if no further evidence concerning that issue is produced.

Our task in the remainder of this chapter is to make this characterization more precise. The most important issue is how a rational fact-finder confronted with conflicting evidence and arguments on a claim can decide whether the claim has been proven. Once we know this, we can define how such decisions on various claims affect the overall outcome of a case, which in turn allows us to give a precise characterization of the tactical burden of proof. However, to answer these questions, we first need to find suitable logical tools.

3. Logical Background: Defeasible Argumentation

In this section we sketch the logical background of our analysis, logics for defeasible argumentation as they have been developed in research on artificial intelligence and applied to legal reasoning by ourselves and others.

3.1. Introductory Remarks

Introductory textbooks to logic often portray logically valid inference as 'foolproof' reasoning: an argument is valid if the truth of its premises guarantees the truth of its conclusion. However, we all construct arguments from time to time that are not foolproof in this sense but that merely make their conclusion plausible when their premises are true. For example, if we are told that John and Mary are married and that John lives in Amsterdam, we conclude that Mary will live in Amsterdam as well since we know that usually married people live where their spouses live. Sometimes such arguments are overturned by counter-arguments. For example, if we are told that Mary is living in Rome to work at the foreign offices of her company for two years, we have to retract our previous conclusion that she lives in Amsterdam. However, as long as such counter-arguments are not available, we are happy to live with the conclusions of our fallible arguments. The question is: are we then reasoning fallaciously or is there still logic in our reasoning?

The answer to this question has been given in three decades of research in artificial intelligence on so-called logics for defeasible reasoning (cf. Prakken and Vreeswijk 2002), partly inspired by earlier developments in philosophy (e.g. Toulmin 1958; Rescher 1977) and argumentation theory (e.g. Walton 1996). At first sight it might be thought that patterns of defeasible reasoning are a matter of applying probability theory. However, many such patterns cannot be analysed in a probabilistic way. In the legal domain this is particularly clear: while reasoning about the facts can (at least in principle) still be regarded as probabilistic, reasoning about normative issues is clearly of a different nature. Moreover, even in matters of evidence, reliable numbers are usually not available so that the reasoning has to be qualitative.

In this section we sketch an account of defeasible reasoning that respects that arguments can be defeasible for various reasons. In short, the account is that reasoning consists of constructing arguments, of attacking these arguments with counter-arguments, and of adjudicating between conflicting arguments on grounds that are appropriate to the conflict at hand. Just as in deductive reasoning, arguments must instantiate inference schemes (now called 'argument schemes') but only some of these schemes capture foolproof reasoning: in our account, deductive logic turns out to be the special case of argument schemes that can only be attacked on their premises.

We shall, in this section, deliberately use a tutorial style, since our primary aim is to explain these ideas to legal theorists with an introductory knowledge of logic but with perhaps no knowledge of the modern developments of the last 30 years.

We think it is important that these developments become widely known among legal theorists, since attacks on the usefulness of formal logic for the law often wrongly presuppose that formal logic equates to deductive logic. In this section we present the ideas of a research community rather than just our own ideas, but in order not to overload the text with references we will limit them to a few key and overview publications.

3.2. Logic of Defeasible Argumentation

As stated previously, we assume that any argument instantiates some argument scheme. (More precisely, in general, arguments chain instantiations of argument schemes into trees, since the conclusion of one argument can be a premise of another.) Argument schemes are inference rules: they have a set of premises and a conclusion. What are the 'valid' argument schemes of defeasible reasoning? Much can be said on this and we will do so later on, but at least the deductively valid inference schemes of standard logic will be among them. Let us examine how deductive arguments can be the subject of attack.

According to the Dutch civil code, persons who are not minors have the capacity to perform legal acts (this means, for instance, that they can engage in contracts or sell their property). Suppose also that some person is not a minor. Then these premises instantiate the deductive scheme of *modus ponens*. In formulas of propositional logic:

> *Argument A:*
> Person & ¬Minor → Has-Legal-Capacity
> Person
> ¬Minor
> Therefore, Has-legal-capacity

(Here '&' stands for 'and', '¬' for 'not' and '→' for 'if … then'.)

Do we have to accept the conclusion of this foolproof argument? Of course not: any first lesson in logic includes the advice: if you don't like the conclusion of a deductive argument, challenge its premises. (In fact, this is the only way to attack a deductive argument since if we accept its premises then its deductive nature forces us to accept its conclusion.) Suppose someone claims that the person is in fact a minor since he is younger than 18. Then the following deductive argument against the premise '¬Minor' can be constructed, which also instantiates *modus ponens*.

> *Argument B:*
> Person & Younger-than-18 → Minor
> Person
> Younger-than-18
> Therefore, Minor.

Now we must choose whether to accept the premise '¬Minor' of argument A or whether to give it up and accept the counter-argument B. Here it is important to note that many cases of premise attack are cases where the premise is assumed in the absence of the contrary. For instance, in our example, being a minor is legally recognized as an exception to the legal rule that persons have the capacity to perform legal acts, so when applying this rule it is reasonable to assume that a person is not a minor as long as there is no evidence to the contrary. Now since argument B provides such evidence to the contrary, we must give up the premise of A and accept the counter-argument.

This leads to a first refinement of deductive logic. It turns out that some arguments have two kinds of premises: ordinary ones and *assumptions*, i.e., premises we are prepared to give up as soon as we have evidence that they are false.

However, not all counter-arguments are attacks on an assumption. Consider again our example. The law of Dutch civil procedure also says that persons younger than 18 who are married are not minors. This gives rise to a deductive argument that attacks the first premise of argument C.

Argument C:
Person & Younger-than-18 and Married → ¬Minor
Person
Younger-than-18
Married
Therefore, ¬Minor

It is important to see that, although superficially argument C attacks argument B's conclusion, C in fact attacks a premise of B, namely its rule premise 'Person & Younger-than-18 → Minor'. This can be seen as follows. If all premises of both arguments are accepted, then a contradiction can be derived, namely, 'Minor & ¬Minor'. To restore consistency, one of these premises has to be false. Since the second and third premise of argument B are also premises of argument C, accepting all premises of C means having to give up the first premise of B. In conclusion, argument C can be extended with:

Argument C continued:
Therefore ¬(Person & Younger-than-18 → Minor)

Moreover, argument B can be continued in a similar way. If all premises of B are accepted, at least one premise of C has to be false. Now the choice is between the rule premise 'Person & Younger-than-18 and Married → ¬Minor' and the factual premise 'Married'. Let us for the sake of illustration assume that the latter is beyond dispute: then argument B can be continued as follows:

Argument B continued:
Moreover, Married
Therefore ¬(Person & Younger-than-18 and Married → ¬Minor)

We now see that the conflict between arguments A and B is not a case of assumption attack, since the conflict is in fact between the two rule premises of these arguments. In such cases some comparative standard has to be applied to see which of these premises has to be preferred. In general, many such standards could be used. When the conflict is caused by conflicting statutory rules (as in our example), we might be able to resolve it on the basis of the hierarchical ordering of the respective regulations (for example, 'federal law precedes state law'), we might prefer the most specific rule on the basis of the principle *lex specialis derogat legi generali* (in our example we could prefer argument C over argument B on this ground) or we might be able to apply some specific statutory conflict rule (for instance, Dutch contract law gives precedence to rules concerning labour contracts over rules concerning other types of contracts). When instead the conflict arises because sources of evidence conflict (such as conflicting witness statements) we might be able to resolve the conflict on the basis of their relative trustworthiness. And when interpretations of a legal concept conflict, we might resort to the underlying purposes or values that are at stake. (Lawyers may even argue about what are the appropriate standards for comparing arguments: this can also be modelled in our account but for simplicity we refer the reader to the literature, e.g. the overview chapter, Prakken and Sartor 2002).

This leads to an important notion of defeasible argumentation, namely a *defeat* relation between conflicting arguments. Whatever conflict resolution method is appropriate, logically speaking we always end up in one of two situations: either the conflict can or cannot be resolved. In the first case we say that both arguments *defeat* each other and in the latter case we say that the preferred argument defeats the other and not vice versa (or that the first argument *strictly* defeats the other). So 'X strictly defeats Y' means 'X and Y are in conflict and we have sufficient reason to prefer X over Y' while 'X and Y defeat each other' means 'X and Y are in conflict and we have no sufficient reason to prefer one over the other'. It should be noted that this 'binary' nature of the outcome of the comparison does not preclude the use of comparative standards which are a matter of degree: even with such standards it must still be decided whether a certain difference in degree is sufficient to accept one argument and reject the other. (As we will explain below in section 4, this is the key to a proper modelling in our approach of differences in proof standards in different legal contexts.)

To summarize, arguments can at least be constructed with deductive argument schemes, their premises are either ordinary ones or assumptions, and arguments can be attacked by arguments that negate one of their premises. If such an attack is on an assumption, the attacker strictly defeats its target, while if the premise attack is on an ordinary premise, some suitable comparative standard has to be used to

see whether one of the arguments strictly defeats the other or whether they both defeat each other.

However, this is not all we can say: it turns out that the binary defeat relation between arguments is not enough to determine which arguments we can accept and which ones we must reject. Suppose that in our example argument C is indeed preferred over argument B on the basis of the *lex specialis* principle. Then we have: B strictly defeats A but C in turn strictly defeats B! Clearly in this case we are justified in accepting A and rejecting B even though B strictly defeats A, since A is 'reinstated' by argument C. However, this is not all: while in this simple case the outcome is intuitive, we can easily imagine more complex examples where our intuitions fail. For instance, another argument D could be constructed such that C and D defeat each other, then an argument E could be constructed that defeats D but is defeated by A, and so on: which arguments can now be accepted and which should be rejected? Here we cannot rely on intuitions but need a *calculus*. Its input will be all the arguments that can be constructed on the basis of a given pool of information, while its output will be an assessment of the *dialectical status* of these arguments in terms of three classes (three and not two since some conflicts cannot be resolved). Intuitively, the *justified* arguments are those that (directly or indirectly) survive all conflicts with their attackers and so can be accepted, the *overruled* arguments are those that are attacked by a justified argument and so must be rejected; and the *defensible* arguments are those that are involved in conflicts that cannot be resolved. Furthermore, a statement is justified if it has a justified argument, it is overruled if all arguments for it are overruled, and it is defensible if it has a defensible argument but no justified arguments. In terms more familiar to lawyers, if a claim is justified, then a rational adjudicator is convinced that the claim is true; if it is overruled, such an adjudicator is convinced that the claim is false; while if it is defensible, they are convinced neither that it is true nor that it is false.

This then is a main component of a logic for defeasible argumentation: a calculus for determining the dialectical status of arguments and their conclusions. What does this calculus look like? Currently there is no single universally accepted one and there is an ongoing debate in AI on what is a good calculus. However, we need not go into the details of this debate, since there is a surprisingly simple and intuitive calculus that suffices for most applications. The idea is to regard an attempt to prove that an argument is justified as a *debate* between a proponent and an opponent of the argument. Since the idea of the game is to test whether, on the basis of a *given* set of statements, a justified argument for a statement of interest can be constructed, both players must construct their arguments on the basis of such a given set of statements. (So unlike in actual legal procedures, the players are not allowed to add new statements to those that are available. See further, section 3.3.) The proponent starts with the argument that he wants to prove justified and then the turn shifts to the opponent, who must provide all its defeating counter-arguments. It does not matter whether they strictly defeat their target or not, since the opponent's task is to interfere with the proponent's attempt to prove his argument justified.

For each of these defeating arguments the proponent must then construct one strict defeater (it has to be a strict defeater since the proponent must prove his argument justified). This process is repeated for as long as it takes: at each of her turns, the opponent constructs mutual and strict defeaters of the proponent's previous arguments, while at each of his turns, the proponent constructs a strict defeater for each of the opponent's previous arguments, and so on. The idea is that our initial argument is justified if the proponent can eventually make the opponent run out of moves in every one of the opponent's lines of attack.

This process can be visualized as follows (the difference in colours will be explained below).

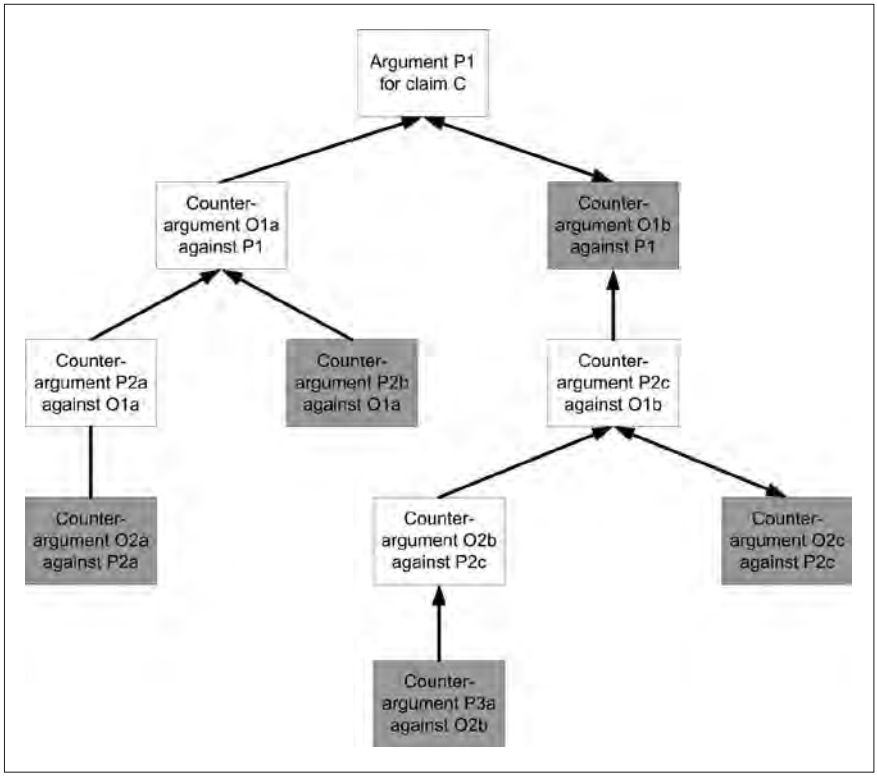

Figure 9.1 A dialectical tree

Note that if an argument is justified this does not mean that the proponent will in fact win the game: he could make the wrong choice at some point. All that it means is that the proponent will win if he plays optimally. In terms of game theory, an argument is justified if the proponent has a winning strategy in a game that starts

with the argument. In fact, there is a simple way to verify whether the proponent has a winning strategy. The idea is to label all arguments in the tree as *in* or *out* according to the following definition:

1. An argument is *in* if all its counter-arguments are *out*.
2. An argument is *out* if it has a counter-argument that is *in*.

In the figures *in* is coloured as grey and *out* as white. It is easy to see that because of (1) all leaves of the tree are trivially *in*, since they have no counter-arguments. Then we can work our way upwards to determine the colour of all the other arguments, ultimately arriving at the colour of the initial argument. If it is grey, i.e., *in*, then we know that the proponent has a winning strategy for it, namely by choosing a grey argument at each point where he has to choose. If, on the other hand, the initial argument is white, i.e., *out*, then it is the opponent who has a winning strategy, which can be found in the same way. So in the above figure the opponent has a winning strategy, which she can follow by choosing argument O1b in her first turn.

Suppose now that at the next stage of the dispute the proponent can construct a strictly defeating counter-argument against O2c. Then the situation is as follows:

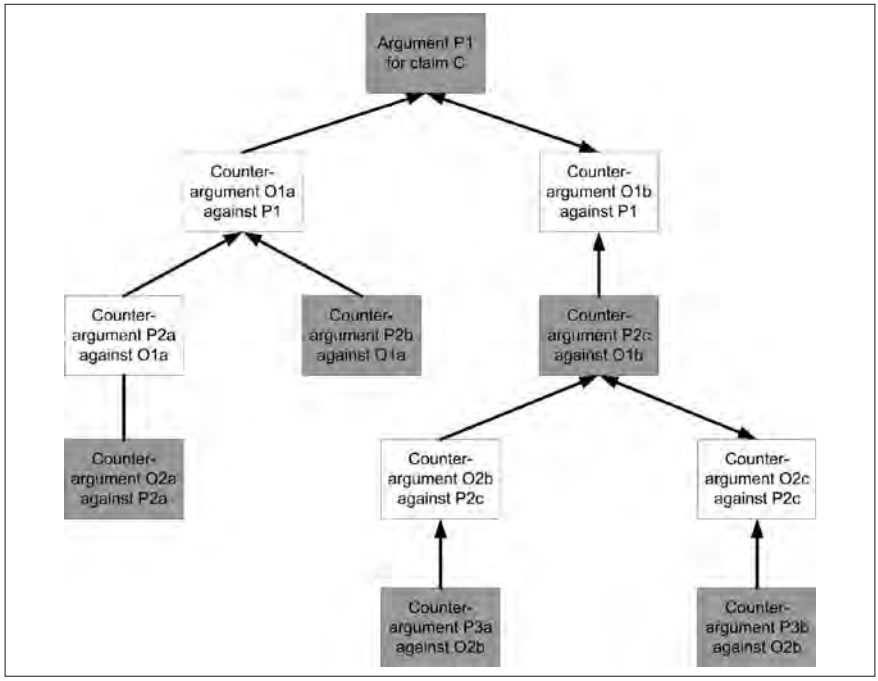

Figure 9.2 An extended dialectical tree

Now argument P1 is *in* on the basis of the new information state so this time it is the proponent who has a winning strategy. He can follow it by choosing P2b instead of P2a when confronted by O1a. This illustrates that when a dispute moves to a new information state, the dialectical status of arguments may change.

Finally, it should be noted that each argument appearing as a box in these trees has an internal structure. In the simplest case it just has a set of premises and a conclusion, but when the argument combines several inferences, it has the structure of an inference tree as is familiar from standard logic. For example, argument B above could be extended with an argument that since the person is a minor, he does not have the capacity to perform legal acts. In tree form, with the conclusion at the top and the premises at the bottom:

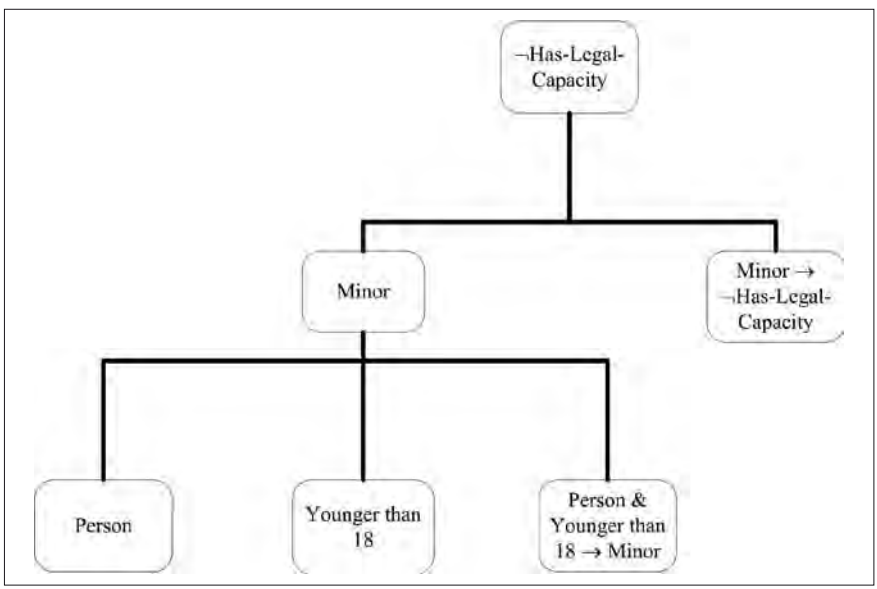

Figure 9.3 An argument

It is important not to confuse these two tree structures. Figure 9.3 displays a single argument: the nodes are statements and the links are inferences. In Figures 9.1 and 9.2 the nodes are complete arguments of which the internal structure is left implicit and the links are defeat relations between arguments. Each node in Figures 9.1 and 9.2 itself implicitly has a tree form as in Figure 9.3.

3.3. Representing Rules and Exceptions

Let us return to the representation of rules and exceptions. The statutory rule used in argument B that persons younger than 18 are minors turned out to have an exception in the case where the person is married. In fact, any statutory rule is subject to exceptions and some of them cannot even be foreseen when drafting the rule. A well-known example is the *Riggs* v. *Palmer* case in American inheritance law (discussed by Dworkin [1977]), in which a grandson had killed his grandfather and then claimed his share of the inheritance. The court made an exception to inheritance law based on the principle that no person shall profit from their own wrongdoing. Moreover, not just statutory rules but also interpretation rules can have exceptions. To use a simplified version of an example of Gardner (1987) on interpretation rules in American contract law, one such rule was 'a statement "I accept" is an acceptance', but an exception was 'a statement "I accept" followed by terms that do not match the terms of the offer is not an acceptance'. Finally, the generalizations often used in evidential reasoning are also subject to exceptions. Consider such generalizations as 'fleeing from the crime scene indicates consciousness of guilt'; an exception would, for example, be that the person may be an illegal immigrant wanting to avoid the police.

Rules that are subject to exceptions are, in AI, often called 'default rules' or in short 'defaults'. Now a convenient way to logically express the default nature of the rules used in legal reasoning is to use general assumptions, which say no more than that there is no exception to the rule. Such general assumptions have their counterpart in legal natural language with expressions such as 'unless there is evidence to the contrary' or 'unless the law provides otherwise'. So, for instance, we could write

R1: Person & ¬Exception-to-R1 → Has-Legal-Capacity
R2: Minor → Exception-to-R1
R3: Declared-insane → Exception-to-R1
R4: … (and so on)

(Note that we now need to give names to rules, to indicate which rule is to be blocked by an exception.) In the same way, we can give exceptional rules a general assumption, to indicate that they, too, may have exceptions:

R5: Minor & ¬Exception-to-R5 → ¬Has-Legal-Capacity
R6: Married → Exception-to-R5
R7: Representative-consents → Exception-to-R5
R8: … (and so on)

We could even adopt a further convention to leave the general assumption clauses implicit in the notation, and this is what we will do in the remainder of this chapter.

3.4. Presumptive Argument Schemes

So far we have only considered deductive argument schemes and we have modelled the defeasibility of arguments as the possibility of premise attack, distinguishing two kinds of premises, ordinary ones and assumptions. It has been suggested that this is all we need: if we distinguish ordinary premises and assumptions and adopt a suitable calculus for adjudicating conflicts between arguments, then the only argument schemes we need are those of deductive logic (see e.g. Bayón 2001). However, if we have a closer look at arguments as they are constructed in practice, we see that the assumptions they make often are not just specific statements but conform to certain reasoning patterns. For instance, evidential arguments are often based on stereotypical evidential sources, such as expert or witness testimony, observation or memory. Other evidential arguments apply the scheme of abduction: if we know that A causes B and we observe B, then in the absence of evidence of other possible causes we may presumptively conclude that A is what caused B. Arguments based on such patterns speak about states of affairs in general: unlike specific generalizations like 'summer in Holland is usually cool' or 'fleeing from a crime scene typically indicates consciousness of guilt' they express general ways of inferring conclusions. For these reasons it is natural to regard such patterns not as patterns for premises of an argument but as a new kind of argument scheme, namely, *presumptive*, or *defeasible* argument schemes. (See e.g. Walton (1996) for a general collection of presumptive argument schemes and Prakken (2005) for an overview of schemes for legal reasoning.)

For instance, the scheme for witness statements could be written as follows:

> *Argument scheme from witness testimony:*
> Person W says that P is true
> Person W was in a position to observe P
> Therefore (presumably), P is true

Another common scheme of legal evidential reasoning is that of temporal persistence.

> *Temporal persistence:*
> P holds at time T1
> Time T2 is later than time T1
> Therefore (presumably), P holds at time T2

A backwards variant of this scheme has 'Time *T2* is earlier than time *T1*' as its second premise. Anderson et al. (2005, Chapter 2) discuss an example in which a murder took place in a house at 4.45pm and a man was seen entering the house at 4.30pm and leaving it at 5.00pm. This gives rise to a forward and a backward application of the temporal persistence scheme, both supporting the presumptive conclusion that the man was in the house at the time of the murder. Temporal

persistence is also often used for proving the existence of a legal right. For instance, ownership of a good is usually proven by proving that it was bought and delivered; the other party must then prove that later events terminated the right of ownership.

The use of presumptive argument schemes gives rise to two new ways of attacking an argument. This is because even if all premises of a presumptive argument are true, its conclusion may still be false since its premises make its conclusion only plausible. The first new form of attack arises when another scheme with a contradictory conclusion is used. For example, two witnesses may give conflicting testimonies, or a witness testimony may be contradicted by direct observation or an abductive argument. Such a conclusion-to-conclusion attack is usually called a *rebutting* attack. A rebutting counter-argument may attack the final conclusion of its target but it may also attack an intermediate conclusion. For example, argument C above attacks the argument in Figure 9.3 by rebutting its intermediate conclusion 'Minor'.

The second new form of attack is based on the idea that a presumptive argument scheme has typical exceptional circumstances in which it does not apply. For example, a witness testimony is typically criticized on the witness's truthfulness or the functioning of his memory or senses. And an application of the practical syllogism may be criticized by pointing at better ways than *A* to realize the same consequences or at negative consequences brought about by realizing *A*. In general, then, each argument scheme comes with a set of *critical questions* which, when answered negatively, give rise to defeating counter-arguments, called *undercutters*. For example, the witness testimony scheme could be given the following critical questions (based on the work of David Schum; see e.g. Anderson et al. 2005):

> *Critical questions to the argument scheme from witness testimony:*
> CC1: Is the witness truthful?
> CC2: Did the senses of the witness function properly?
> CC3: Does the memory of the witness function properly?

And the main critical question of the temporal persistence scheme is whether there is reason to believe that *P* does not hold at a time *T3* between *T1* and *T2*. Undercutting counter-arguments do not attack a premise or the conclusion of their target but instead deny that the scheme on which it is based can be applied to the case at hand. Obviously, such denial does not make sense for deductive argument schemes, since such schemes *guarantee* that their conclusion is true if their premises are true. In sum then, while deductive arguments can only be attacked on their premises, presumptive arguments can also be attacked on their conclusion and on their inference steps. (We will formalize undercutters as arguments for a conclusion '¬name', where 'name' is a placeholder for the name of the undercut argument scheme.) Note that as with rebutting attack, an undercutting attack can be launched both at the final inference of its target and at an intermediate one.

How many presumptive argument schemes are there? Here the classical logician will be disappointed. One of the main successes of modern formal logic has been that an infinite number of valid deductive inferences can be captured by a finite and even very small number of schemes. However, things are different for defeasible inference: many different classifications of presumptive argument schemes have been proposed, and the debate as to what should count as an argument scheme is still ongoing. Moreover, while some schemes, such as abduction and the practical syllogism, can arguably be used in any domain, other schemes may be domain-dependent. For instance, Anderson (2007) points out that in legal contexts the witness and expert schemes have different critical questions than in ordinary commonsense reasoning. For a detailed discussion of argument schemes relevant for legal evidential reasoning see Bex et al. (2003).

Some readers might wonder whether legal reasoning about evidence is argument-based in the way sketched above at all. In particular, at first sight our approach would not seem to be able to model the ubiquitous phenomenon of accrual, or aggregation of various pieces of evidence pointing in the same direction. For example, if several witness testimonies support the same claim, then the argument scheme from witness testimony gives rise to three different arguments that in no way can be combined, while yet intuitively a party's position seems stronger the more witnesses it can produce who support the same claim. However, in Prakken (2005) it is shown how an argument scheme for argument accrual can be modelled in a logic for defeasible argumentation, and how the resulting formalism can be applied to reasoning about evidence. Because of space limitations, the reader is referred to that publication for the details.

Finally, to return to our question at the beginning of this section, we have seen that there is indeed logic in defeasible argumentation: the form of arguments must fit a recognized argument scheme (of course, still considering the debate on what should be recognized as such), and the dialectical status of an argument can be determined in a systematic dialectical testing procedure. On the other hand, what cannot be provided by such a logic are the standards for comparing conflicting arguments: these are contingent upon input information, just like the information from which arguments can be constructed.

3.5. Embedding of Argumentation in Procedural Settings

The previous subsections assumed that arguments are constructed and their dialectical status is determined on the basis of a given pool of information. However, at least two notions of burden of proof (the burden of production and the tactical burden) assume a dynamic setting in which new information can be introduced at various points in the proceedings. We must therefore explain how our logical account of defeasible argumentation can be embedded in a dynamic procedural setting.

In this chapter we assume that the exchange of arguments in a dispute is regulated by some legal procedure. However, we abstract from the details of such a

procedure and simply assume that a dispute consists of a sequence of stages which are characterized by different pools of input information and where the parties (including the adjudicator) can move from one stage to another by formulating new claims and arguments. The information pool of a stage then consists of all claims and the premises of all arguments stated up to that stage. (We also abstract from the fact that many procedures allow the parties to dispute, concede and retract claims and premises.) The outcome of a dispute is determined by applying the calculus for the dialectical status of arguments to the final stage. In fact, when verifying a tactical proof burden we shall also apply the argument game to intermediate stages, to verify what would be the outcome of the dispute if an intermediate stage were the final stage. As for the standards for comparing conflicting arguments, we assume that they are used in the final stage by the adjudicator. This means that if we apply the dialectical calculus to an intermediate stage, we have to guess which standards will be used in the final stage.

In terms of dialectical trees this can be formulated as follows (see also Modgil and Prakken, 2008). At each stage the parties add arguments to the dialectical tree defined above but with one difference: since at intermediate stages the comparative standards are unknown, the parties must move all counter-arguments that they can construct to any argument of the previous turn: at the final stage the adjudicator (after possibly having added her own arguments) applies the comparative standards to the final dialectical tree. In doing so, she may have to prune the tree: if an attacker moved by the proponent does not strictly defeat its target, then the entire sub-tree starting with this argument is pruned from the tree. The same holds for the opponent's arguments that do not defeat their targets. In the resulting pruned dialectical tree some of the opponent's arguments may have more than one strict defeater. It should then be checked whether a unique choice between these strict defeaters can be made in such a way that all branches of the tree end with a move by the proponent: if (and only if) this is possible, the proponent has a winning strategy for his initial argument so that it is justified.

This process is visualized in Figure 9.4.

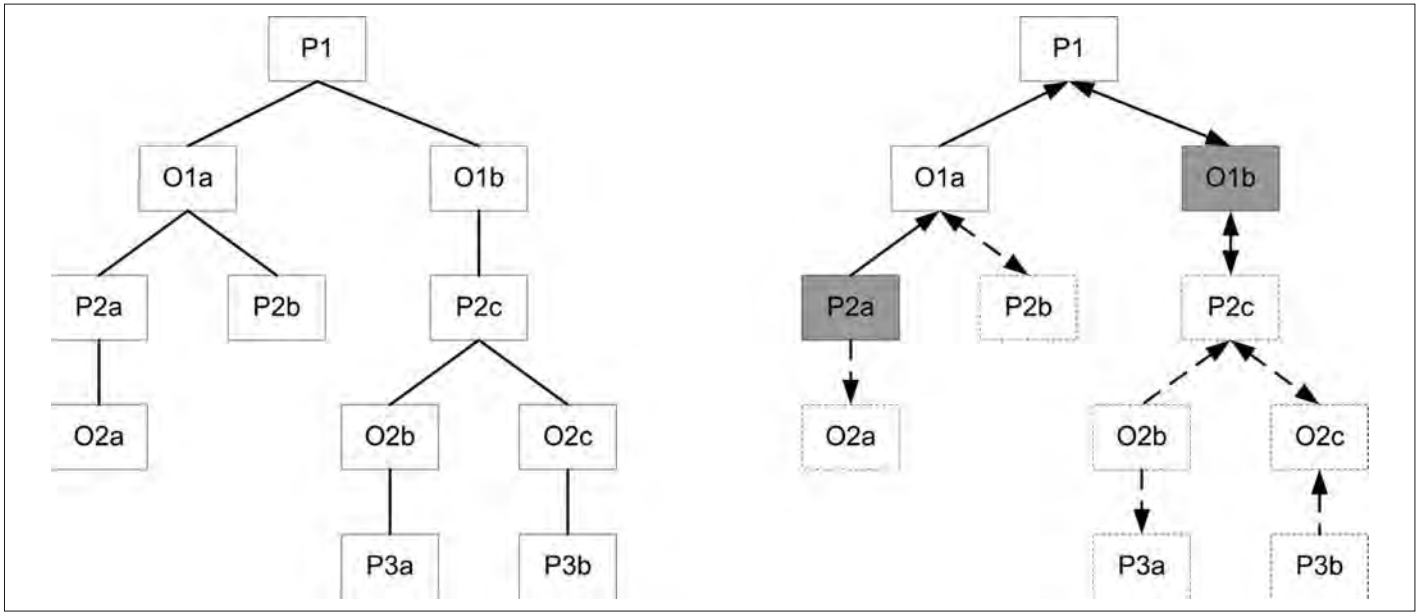

Figure 9.4 Adjudicating conflicts between arguments

In the dialectical tree on the left-hand side the links between arguments have no arrow, to indicate that we only know that they are in conflict with each other. On the right-hand side the adjudicator's decisions are displayed as defeat links. The parts of the tree that have been chopped off are displayed with dashed lines: O2a is chopped off since it is strictly defeated by P2a; then P2b is chopped off since it only weakly defeats O1a, and the entire sub-tree below O1b is chopped off since P2c only weakly defeats O1b. The remaining tree is labelled as above and we find that P1 is not justified.

4. A Logical Account of the Burdens of Proof[2]

We first recall the definitions of section 2. The *burden of persuasion* specifies which party has to prove a statement to a specified degree (the standard of proof) with the penalty of losing in respect of the issue. Whether this burden is met is determined in the final stage of a proceeding, after all the evidence is provided. The *burden of production* specifies which party has to offer evidence on an issue at different points in a proceeding. If the burden of production is not met, the issue will be decided as a matter of law against the burdened party, while if it is met, the issue will be decided in the final stage according to the burden of persuasion. Both these burdens are assigned as a matter of law. By contrast, the *tactical burden of proof* is a matter of tactical evaluation in that a party must assess the risk of ultimately losing in respect of an issue if no further evidence concerning that issue is produced.

How can these notions be analysed against the logical background of section 3? As remarked above, we need two things: a logic of defeasible argumentation and its embedding in a procedure for dispute. We will first formally characterize the three proof burdens and then illustrate them with two legal examples. In the discussion of these examples it will become clear that the procedure given in section 3 for determining the dialectical status of arguments must be refined: in particular since the version in section 3 does not allow for distributions of the burden of persuasion between the sides in a dispute.

4.1. Logical Definition

The *burden of persuasion* for a claim can, in terms of section 3, be defined as the task of making sure that in the final stage of the proceeding there exists a justified argument for the claim. *Proof standards* for the burden of persuasion can be formalized by a careful definition of the defeat relation between arguments, in particular when rebutting counter-arguments are compared: a stronger rebutting argument should strictly defeat a weaker argument only if the degree to which it is

2 The ideas described in this section were earlier expressed in more condensed form in Prakken and Sartor 2006.

stronger satisfies the applicable proof standard; otherwise both arguments defeat each other. For example, if the standard is 'on the balance of probabilities', the fact-finder can already say that A strictly defeats B if A is just a little bit stronger than B, while if the standard is 'beyond reasonable doubt', the fact-finder can say this only if, when faced with only A and B, he would certainly accept A's conclusion. Recall that in section 2 we said that if a burden of persuasion for a statement is fulfilled, a rational fact-finder is (to the required degree) convinced that the statement is true, while if the burden is not met, such a fact-finder is not convinced that the statement is true. Recall also that in section 3 we informally said that if a claim has a justified argument, then a rational adjudicator is convinced that the claim is true, while if it has defensible but no justified arguments they are neither convinced that it is true nor that it is false. We have now seen how the legal notions of section 2 and the logical notions from section 3 are related: having the burden of persuasion for a claim amounts to the task of having a justified argument for it in the final state of the proceeding.

The *burden of production* is harder to define in logical terms. Its logical aspect is that the adjudicator must at the appropriate stage of a proceeding examine whether an evidential argument has been produced for the claim on which the burden of production rests. What goes beyond logic is the demand that this argument must have sufficient internal strength in that 'reasonable minds' can disagree about whether its conclusion would hold if only its premises were known. In section 4.2 we will show how a negative decision on this issue can be expressed as an argument.

Finally, once the burden of persuasion has been assigned by law, the *tactical burden of proof* is automatically induced by the argument game for testing an argument's dialectical status, as applied to any given stage of the proceeding: at a given stage a party has a tactical burden of proof with regard to an issue if the evidence and the arguments thus far provided lead to assessing that issue in a way that goes against that party (and so would likely be concluded by the triers of fact if no new elements were provided to them before the end of the proceedings). In fact, the 'strength' of the tactical burden depends on the allocation of the burden of persuasion: the party who has the burden of persuasion for an initial claim is proponent in the argument game and therefore has to strictly defeat the other parties' arguments, while the other party, being the opponent in the argument game, only has to weakly defeat the proponent's argument. While the dialectical asymmetry of the argument game thus accounts for the fact that for one party the tactical burden is stronger than for the other, its embedding in a dynamic setting accounts for the possibility that the tactical burden shifts between the parties: once a party finds herself in a situation where, according to her assessment, she would likely lose if nothing else is known, this means that the tactical burden has shifted to her.

4.2. A Criminal Case

Consider again our example in section 2 from Dutch law about murder, with a general rule that killing with intent is punishable as being murder, and a separate rule expressing an exception in the case of self-defence. (We now use a double arrow instead of a single one to express that we represent defeasible rules; recall that according to our notational convention each such rule Ri has an implicit assumption ¬Exception-to-Ri.)

R1: Killing and Intent ⇒ Murder
R2: Self-defence ⇒ Exception-to-R1

As said in section 2, the law thus expresses that the prosecution has the burdens of production and persuasion of 'Killing' and 'Intent' while the defence has the burden of production for 'Self-defence' and the prosecution has the burden of persuasion for '¬Self-defence'.

Consider now a murder case and assume that the prosecution can satisfy his burden of persuasion with respect to 'Killing' and 'Intent' with evidential arguments (which we leave implicit for the sake of brevity). Then the tactical burden shifts to the defence, since if she provides no other evidence the adjudicator will likely convict her. The burdens of production then imply that the accused can only escape conviction by providing some minimally credible evidence of an exception to R1, such as that the killing was done in self-defence. For instance, the defence could provide a witness who says that the victim threatened the accused with a knife. (Below we will only list rules and facts; the arguments constructed with them are visualized in Figure 9.5.)

F1: Witness W1 says 'knife'
R3: Knife ⇒ Threat-to-life
R4: Killing and Threat-to-life ⇒ Self-defence

(Note that the argument for 'Self-defence' uses the argument scheme from witness testimony to conclude 'Knife' from fact F1.) Suppose that at this point in trial the judge has to assess whether the defence has satisfied her burden of production with respect to 'Self-defence'. As explained above, this amounts to deciding whether reasonable minds can disagree on whether there was self-defence if only the premises of this argument are known to hold. Let us first assume that the judge rules that this is not the case, so that the defence has not satisfied her burden of production. How can such a ruling be expressed in our logical analysis of section 3? A detailed answer depends on the precise grounds for the ruling but in any case it can be logically expressed as a strictly defeating counter-argument of the defence's argument for self-defence. The proper procedural setting will then disallow counter-arguments to the ruling so that the defence's argument will

certainly be overruled in the final stage of the dispute, and the defence loses on the issue of self-defence as a matter of law.

Let us now instead assume that the defence satisfied her burden of production for 'Self-defence'. In Anglo-American systems this means that the issue of self-defence must be addressed by the fact-finder in the final stage of the dispute. Moreover, if the current stage were the final stage, there would be a chance that the defence would win. To avoid the risk of losing, the prosecution should therefore provide additional evidence to take away the reasons for doubt raised by the defence. In other words, the prosecution now at least has a tactical burden to provide evidence against 'Self-defence'. Moreover, the prosecution also has the burden of persuasion against this claim. This is automatically captured by our logical analysis since the prosecution, being the proponent in the argument game, has to strictly defeat the defence's argument for 'Self-defence'.

Let us assume that the prosecution attempts to fulfill his burden of persuasion with a witness who declares that the accused had enough time to run away.

F2: Witness W2 says 'time-to-run-away'
R4: Time-to-run-away $\Rightarrow$ $\neg$Threat-to-life

Let us also assume that the evidence is of the kind that is usually sufficient to persuade the fact-finder, i.e., it is likely that the fact-finder will say that the argument using R4 for '$\neg$Threat-to-life' strictly defeats the argument using R3 for 'Threat-to-life'. Then the proposition 'murder' is justified again in the new stage of the dispute (again relative to the party's assessment of the likely decisions of the fact-finder). This shifts the tactical burden to the defence to provide counter-evidence that probably makes 'self-defence' at least defensible in the resulting stage. For example, the defence could provide evidence that witness W2 is a friend of the victim, which makes her unreliable.

F3: Witness W2 is-friend-of-victim
R5: Witness W is-friend-of-victim $\Rightarrow$ $\neg$Witness-Testimony-Scheme

Here Rule R5 expresses an undercutter of the presumptive argument scheme from witness testimony. By undercutting the prosecution's argument for '$\neg$Threat-to-life', the new argument reinstates the defence's argument for `Threat-to-life' and therefore makes the prosecution's main claim overruled again.

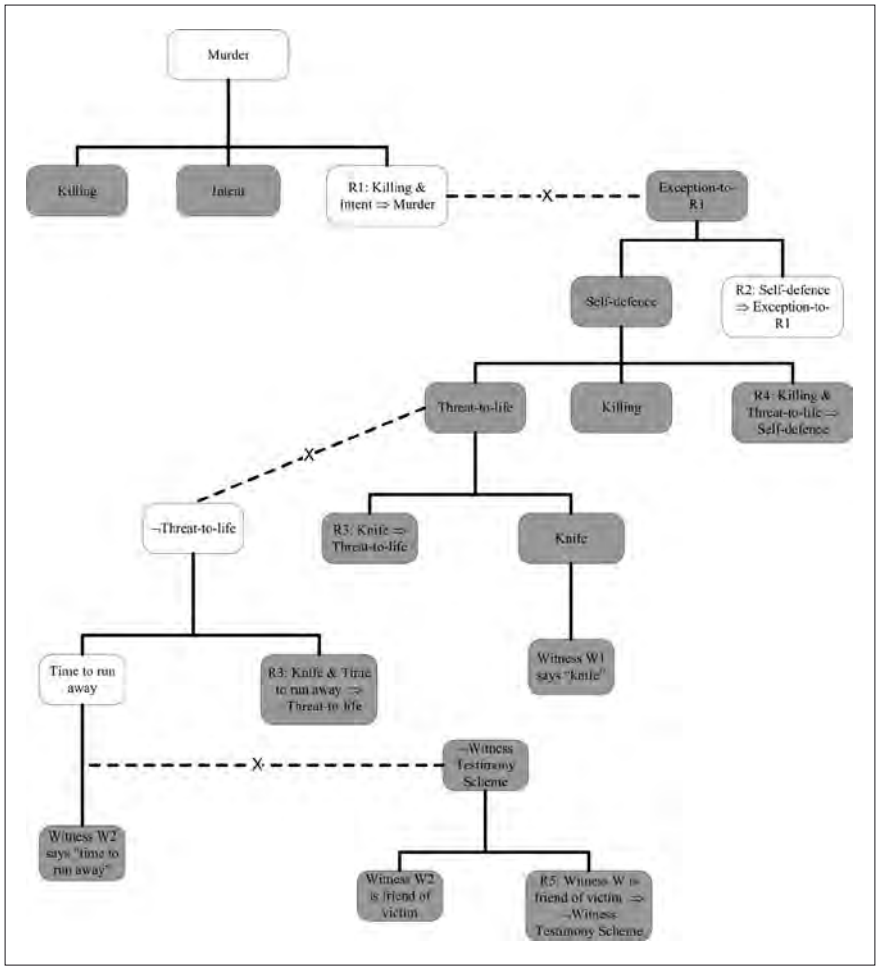

Figure 9.5 The arguments in the criminal case

All arguments constructed so far are visualized in Figure 9.5. The dotted lines with a cross indicate attack relations between arguments. The colouring of the premises and conclusions reflects their dialectical status if all attacks are regarded as successful. A grey statement is justified and a white statement is overruled.

This example has illustrated when and how the three kinds of burden rest on the parties during a dispute. The burdens of production and persuasion are fixed: they cannot shift from one party to the other. The burden of production on an issue is fulfilled as soon as the burdened party provides the required evidence on that issue and after that it is no longer relevant. The burden of persuasion, once created, remains on a party until the end of the discussion of the case, namely, until the

point when a party is precluded from giving any further input to the fact-finders. By contrast, the tactical burden on an issue is not fixed; it can shift between the parties any number of times during the discussion of the case, depending on who would probably win on that issue if no more evidence were provided.

4.3. A Civil Case

In our criminal example the burden of production was divided between the parties but the burden of persuasion was on one side only. Is this always the case? Our second example from section 2 illustrates that at least in civil cases the burden of persuasion can also be distributed between the parties.

Consider again the legal rule that a contract is created by an offer and acceptance and a statutory exception in case the offeree was insane when accepting the offer.

R1: Offer and Acceptance $\Rightarrow$ Contract
R2: Insane $\Rightarrow$ Exception-to-R1

As we said in section 2, a plaintiff who wants to argue that a contract was created has the burdens of production and persuasion for 'Offer' and 'Acceptance' while the defendant has the burdens of production and persuasion for any statutory exception, in this case, for 'Insane'. So it does not suffice for the defendant to cast doubt in this issue: she must (to the relevant degree) convince the adjudicator that she was insane. In terms of section 3, the defendant must, in the final stage of the proceeding, have a justified argument for the claim that she was insane.

This, however, creates a problem for the logic of section 3. Recall that a claim's dialectical status can be logically tested in an argument game between a proponent and an opponent. The problem is that this game fixes the dialectical asymmetry throughout the game: all of the proponent's counter-arguments must be strictly defeating (i.e., doubt-removing), while the opponent's counter-arguments can be weakly defeating (i.e., doubt-raising). So the opponent's counter-arguments always succeed if they cast doubt. However, our example shows that doubt-raising arguments do not suffice if the opponent has the burden of persuasion: in that case doubt-removing arguments are needed.

To meet this demand, in Prakken (2001) the argument game of section 3 was modified to allow for the two players in a dialogue (plaintiff and defendant) to have different dialectical roles (proponent or opponent) for different propositions. So, for instance, the plaintiff in our example can be the proponent with respect to 'Offer' and 'Acceptance' while he can be the opponent with respect to 'Insane'. Accordingly, the new argument game assumes as input not just a set of arguments ordered by a binary defeat relation but also an allocation of proof burdens for statements to the plaintiff and the defendant, expressing who has the burden of persuasion for each proposition. In fact, an interesting observation can be made here about the relation between logic and law. Usually, an existing logic is simply applied to legal reasoning, but here we have a case where an essential feature

of legal reasoning requires a change of an existing logic to make it suitable for modelling that feature.

5. Other Relevant Notions

We now briefly present three additional kinds of burden that are sometimes distinguished, namely, the burden of contesting, the burden of claiming and the burden of argument. In our formal approach they can be accounted for by the embedding of the logic in a procedure for dispute.

A *burden of contesting* exists when a statement unfavourable to a party will be assumed to hold if that party does not contest that proposition. This burden is usually conditional on the fact that the other party has claimed the statement at issue. In other words, not contesting a proposition claimed by the other party implicitly counts as conceding it. For instance, in both Dutch and Italian law of civil procedure, a factual proposition claimed by one party and not contested by the other party must be accepted by the judge even if no evidence for it was provided by the claiming party.

In our formal account this can be modelled by assuming that in the final stage of a proceeding such non-contested statements belong to the available pool of information as a special kind of premise that cannot be attacked by counter-arguments. Thus we have three kinds of premises: ordinary premises, assumptions and 'certain' premises.

The opposite of the burden of contesting is the *burden of claiming*. Such a burden exists for a statement if it cannot be accepted by the fact-finder because it was not claimed by the party for which it is favourable. For instance, in Dutch civil procedure (where this burden is called *stelplicht*) the plaintiff has the burden of claiming (at the first possible stage) all operative facts of his main claim while the defendant has the burden of claiming (at the first possible stage) any exception she wishes to argue for. Note that this is not the same as a burden of production; neither of these burdens entails the other: first, the burden of production, unlike the burden of claiming, implies the burden to give an argument; and second, it is conceivable that a party satisfies a burden of production for an exception by producing evidence for it without explicitly claiming that the exception holds (in our murder case example the judge might be able to infer the possibility of self-defence from the suspect's statements even if he has not explicitly claimed that he acted in self-defence).

In our formal account this burden can be modelled by disregarding, in the final stage of a proceeding, any argument for a conclusion on which a burden of claiming rests if the conclusion was not claimed by the interested party at the required stage of the proceeding.

Finally, the burden of production can be generalized to a *burden of argument*, which, unlike the burden of production, can also apply to non-factual statements. Such a burden exists for a statement if it cannot be accepted by the adjudicator

because no argument for it is provided. For example, Dutch civil procedure has such a burden for the plaintiff's main claims and the defendant's main counterclaims, and calls it *substantiëringsplicht*.

In our approach this can be modelled by making sure that statements on which a burden of argument rests cannot be in the pool of information from which arguments are constructed in the final stage of a proceeding. If they are needed to establish some conclusion, they must first themselves be derived as an intermediate conclusion from the available pool of information.

6. Conclusion

In this chapter we have given a logical analysis of several notions concerning the burden of proof. It has turned out that logics for defeasible argumentation, when embedded in a dynamical setting, provide the means to logically characterize the difference between several kinds of proof burdens. Our main contributions have been a precise distinction between two burdens that are sometimes confused, namely the burden of production and the tactical burden, and the insight that the burden of persuasion can be verified by applying an argumentation logic to the body of information available at the final stage of a proceeding.

It remains to be discussed to what extent our analysis applies to other approaches than argument-based ones. As long as any reasoning formalism is used that allows for a fallible notion of proof, much of our analysis still applies. All that is needed is that the formalism accepts as input a description of an evidential problem and produces as output a fallible assessment whether a certain claim has been proven. These assumptions are clearly satisfied by non-monotonic logics that are not argument-based, but the same holds for formalisms such as probability theory or Thagard's connectionist theory of explanatory coherence (Thagard 2004), at least if their output (a posterior probability or a numerical measure of coherence) is combined with a numerical proof standard for proof of the relevant statement. The key is the embedding of such formalisms in a dynamic setting as described in section 3.5, which allows the burden of persuasion to be defined as the burden of making sure that in the final stage of the proceedings the numerical value of the statement exceeds the proof standard. Clearly, such a notion of proof is fallible in that a statement provable in this way may no longer be provable on the basis of extended input information. In consequence, the notion of tactical burden also applies to these formalisms: at each stage in a proceeding a party should assess the likely outcome if that stage were the final stage, and if this outcome is unfavourable to that party, it should introduce new information that changes it.

On the other hand, what cannot be easily modelled in non-logical approaches are the logical relations that sometimes hold between different statements to be proven. See, for instance, our example in section 4.3, where the burdens on the plaintiff to prove offer and acceptance and on the defendant to prove that she was insane were assigned because of the (defeasible) logic governing the legal rules R1

and R2, of which these facts, respectively, are the conditions. So if probabilistic or connectionist methods are used to prove conditions of legal rules, there is still the theoretical issue of combining these methods with the logic governing these rules. This problem does not arise if the evidential part of the reasoning is modelled as argumentation (at least not if the set of argument schemes is sufficiently rich to combine evidential and non-evidential forms of reasoning). The same can be said about the relation with presumptions, since they can have non-statistical justifications. For example, some presumptions are based on considerations of fairness, such as having better access to information. For these reasons it is not immediately clear how they can be modelled in a purely probabilistic or connectionist approach; an argument-based logic, by contrast, can naturally model reasoning with presumptions, since presumptions usually have a rule-like structure, saying that if certain facts are proven, certain other facts may be taken to hold in the absence of counter-evidence (cf. Prakken and Sartor 2006).

Having said this, a full account of the relationship between presumptions and the three kinds of burden of proof is by no means trivial. However, such an account goes beyond the scope of the present chapter and must await another occasion.

Concluding our investigations, we have, of course, not proven that our argument-based account of burdens of proof and presumptions is fully adequate, but at least we hope that we have put a tactical burden of proof on those who want to argue otherwise.

References

Anderson, T. (2007), 'Visualization tools and argument schemes: A question of standpoint', *Law, Probability and Risk* 6, 97–107.

Anderson, T., Schum, D. and Twining, W. (2005), *Analysis of Evidence*, 2nd Edition (Cambridge: Cambridge University Press).

Bayón, J.C. (2001), 'Why is Legal Reasoning Defeasible?', in Soeteman (ed.).

Besnard, Ph., Doutre, S. and Hunter, A. (eds) (2008), *Computational Models of Argument: Proceedings of COMMA 2008* (Amsterdam: IOS Press).

Bex, F.J., Prakken, H., Reed, C. and Walton, D.N. (2003), 'Towards a formal account of reasoning about evidence: Argumentation schemes and generalisations', *Artificial Intelligence and Law* 11, 125–65.

Capper, C. and Cross, R. (1990), *Cross on Evidence* (London: Butterworths).

Denning, A. (1945), 'Presumptions and burdens', *The Law Quarterly Review* 61, 379–83.

Dworkin, R.M. (1977), *Taking Rights Seriously* (Cambridge, MA: Harvard University Press).

Engers, T.M. van (ed.) (2006), *Legal Knowledge and Information Systems. JURIX 2006: The Nineteenth Annual Conference* (Amsterdam: IOS Press).

Gabbay, D. and Guenthner, F. (eds) (2002), *Handbook of Philosophical Logic*, 2nd Edition, vol. 4 (Dordrecht, Boston and London: Kluwer Academic Publishers).

Gardner, A. von der Lieth (1987), *An Artificial Intelligence Approach to Legal Reasoning* (Cambridge, MA: MIT Press).

Hahn, U. and Oaksford, M. (2007), 'The burden of proof and its role in argumentation', *Argumentation* 21, 36–61.

Kakas, A. and Sadri, F. (eds) (2002), *Computational Logic: Logic Programming and Beyond. Essays In Honour of Robert A. Kowalski*, Part II. Lecture Notes in Computer Science 2048 (Berlin: Springer Verlag).

Keane, A. (1994), *The Modern Law of Evidence* (London: Butterworths).

Modgil, S. and Prakken, H. (2008), 'Applying Preferences to Dialogue Graphs', in Besnard et al. (eds).

Prakken, H. (2001), 'Modelling defeasibility in law: Logic or procedure?', *Fundamenta Informaticae* 48, 253–71.

—— (2004), 'Analysing reasoning about evidence with formal models of argumentation', *Law, Probability and Risk* 3:1, 33–50.

—— (2005), 'AI and law, logic and argument schemes', *Argumentation* 19:3, 303–20.

Prakken, H. and Sartor, G. (2002), 'The Role of Logic in Computational Models of Legal Argument: A Critical Survey', in Kakas and Sadri (eds).

Prakken, H. and Sartor, G. (2006), 'Presumptions and Burdens of Proof', in Van Engers (ed.).

Prakken, H. and Sartor, G. (2007), 'Formalising arguments about the burden of persuasion', in *Proceedings of the Eleventh International Conference on Artificial Intelligence and Law* (New York: ACM Press) pp. 97–106.

Prakken, H. and Vreeswijk, G.A.W. (2002), 'Logics for Defeasible Argumentation', in Gabbay and Guenther (eds).

Rescher, N. (1977), Dialectic: A Controversy-Oriented Approach to the Theory of Knowledge (Albany: State University of New York Press).

Rosenberg, L., Schwab, H. and Peter, G. (1993), *Zivilprocessrecht* (Munich: Beck).

Soeteman, A. (ed.) (2001), Pluralism and Law (Dordrecht: Kluwer Academic Publishers).

Spencer, M. and Spencer, J. (2007), *Q and A: Evidence 2007–2008* (Oxford: Oxford University Press).

Strong, J.W. (1992), *McCormick on Evidence*, 4th Edition (St. Paul, MN: West Publishing Company).

Thagard, P. (2004), 'Causal inference in legal decision making: Explanatory coherence vs. Bayesian networks', *Applied Artificial Intelligence* 18, 231–249.

Tonini, P. (2007), 'Onere della prova nel processo penale', in *Enciclopedia Giuridica del Sole 24 Ore* (Milan: Il Sole 24 Ore), pp. 306–12.

Toulmin, S.E. (1958), *The Uses of Argument* (Cambridge: Cambridge University Press).

Walton, D.N. (1996), *Argumentation Schemes for Presumptive Reasoning* (Mahwah, NJ: Erlbaum).

Wigmore, J.H. (1940), *A Treatise on the Anglo-American System of Evidence in Trials at Common Law*, 3rd Edition (Boston, MA: Little, Brown & Company).

—— (1962), *Evidence in Trials at Common Law*, vol. IX (New York: Little, Brown & Company).

Williams, C.R. (2003), 'Burdens and standards in civil litigation', *Sydney Law Review* 25, 165–88.

Zuckerman, A. (2006), *Zuckerman on Civil Procedure* (London: Sweet & Maxwell).