

# A Survey of Uncertainties and Their Consequences in Probabilistic Legal Argumentation

Matthias Grabmair and Kevin D. Ashley

**Abstract** In this chapter, we present a survey of different sorts of uncertainties lawyers reason with and connect them to the issue of how probabilistic models of argumentation can facilitate litigation planning. We briefly survey Bayesian Networks as a representation for argumentation in the context of a realistic example. After introducing the Carneades argument model and its probabilistic semantics, we propose an extension to the Carneades Bayesian Network model to support probability distributions over argument weights, a feature we believe is desirable. Finally, we scout possible future approaches to facilitate reasoning with argument weights.

## 1 Introduction

Formal models of argument can be considered as qualitative alternatives to traditional quantitative graphical probability methods for reasoning about uncertainty, as has been observed by Parsons et al. (1997). We understand a *formal model of argument* as a description of a dialectical process in formal mathematical terms. It typically contains entities such as arguments and/or propositions in some language as well as (possibly typed) relations between those elements. *Probabilistic graphical models* is a summary term for graph structures encoding dependencies between a set of random variables.

In other words, discourse about what is likely to be true in an uncertain domain becomes a question of persuasion and proof rather than a subject of a supposedly objective probability. For a proposition of interest, the answer given by the system

---

M. Grabmair (✉) • K.D. Ashley  
University of Pittsburgh, Pittsburgh, PA, USA  
e-mail: [mag134@pitt.edu](mailto:mag134@pitt.edu)

as to its truthfulness is hence no longer a number but a qualitative label assigned to it according to some formal model of dialectic argumentation.

One task that combines argumentation and the need to deal with uncertainty is litigation planning, that is, planning whether to commence a lawsuit and what kinds of legal claims to make. The task of a lawyer in litigation planning is unique in that she is professionally engaged in argumentation to convince an audience of the truthfulness of certain assertions. For her, however, argumentation takes place in a planning context, and she needs decision support for the uncertainty she faces *in* argumentation. Her utility is the degree of persuasiveness of the arguments, and she has to plan her litigation strategy and research according to what will give her the maximum payoff in the task of convincing her audience.

The nature of this task, that is, systematizing uncertainty about the outcome of an argument, is a kind of practical reasoning. It is hence a promising step to bring quantitative models of uncertainty inside qualitative argument models with the goal of obtaining utility information for purposes of planning ones moves and allocating resources. In litigation, lawyers must view their case from multiple angles. They can argue about the facts playing in their favor, about the law rendering the established facts favorable for ones side, about the ethical or policy concerns given the facts, and, finally, about relating the case at bar to previous decisions in a way that helps their client. Motivated by the search for a promising systematization in which to invest further research and development resources, we survey how these four kinds of uncertainties interrelate using a realistic trade secret law case. We distill a set of criteria for the suitability of a probabilistic graphical representation to model legal discourse and present the Carneades argument model as well as its probabilistic interpretation using Bayesian Networks [BN] as our running example model. Finally, we explain and emphasize the role of argument weights in the process and suggest promising extensions to the model.

## 2 Survey of Relevant Uncertainties

In planning for litigation, lawyers take account of uncertainties both in the facts themselves and in how the law applies to the facts. They address uncertainty about the facts through evidence and evidential argument and uncertainty about the application of the law through dialectical argumentation.

In so doing, we believe that litigators need to address four main types of uncertainties:

- *Factual*: evidence and the plausibility of factual assertions
- *Normative*: dealing with the application of the law to facts
- *Moral*: dealing with an ethical assessment of the conflict
- *Empirical*: relating to prior outcomes in similar scenarios

## 2.1 *The Example Case*

We illustrate these in the context of a real lawsuit, ConnectU versus Zuckerberg,<sup>1</sup> in which plaintiffs C. and T. Winklevoss and D. Narendra sued defendant Mark Zuckerberg. The plaintiffs founded ConnectU while at Harvard University in 2002–2003; they allegedly developed a business plan for a novel institute-specific website and network for online communication and dating. According to the plan, ConnectU was to expand to other institutions and generate advertising revenue. To develop software for the website, the plaintiffs hired two students. They engaged Zuckerberg to complete the code and website in exchange for a monetary interest if the site proved successful. They gave him their existing source code, a description of the website’s business model and of the functionality and conceptual content, and the types of user information to be collected. They alleged that the defendant understood this information was secret and agreed to keep it confidential. The plaintiffs averred that they had stressed to the defendant that he needed to complete the code as soon as possible and that he assured them he was using his best efforts to complete the project. As late as January 8, 2004, he confirmed via email that he would complete and deliver the source code. On January 11, however, the defendant registered the domain name “thefacebook.com” and a few weeks later launched a directly competing website. By the time plaintiffs could launch ConnectU.com, Zuckerberg had achieved an insurmountable commercial advantage.

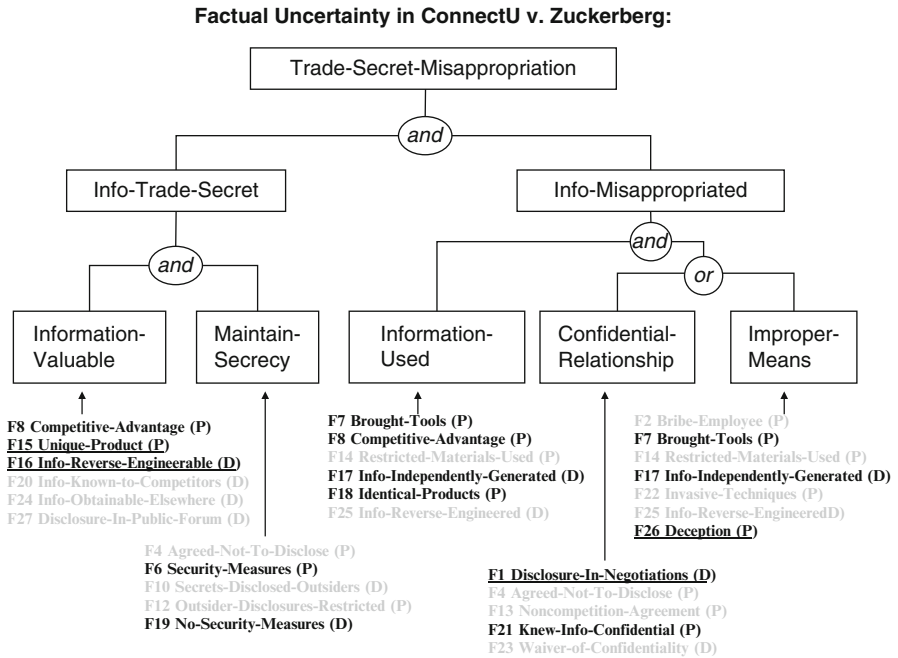
An initial decision is what kind of legal claim to bring such as trade secret misappropriation or breach of contract. In considering a claim of trade secret misappropriation, an attorney would be familiar with the relevant law and might even have in mind a model of the relevant legal rules and factors like the one employed in IBP (Ashley and Brüninghaus 2006) and shown in Fig. 1. Factors are stereotypical patterns of facts that strengthen a claim for one side or the other.

## 2.2 *Factual Uncertainty*

Regarding factual uncertainty, the attorney knows that showing the factual basis for some factors will be easier than others. In the diagram, easy to prove factors are underlined. Clearly, the information has commercial value; the social network model was a unique product at the time. And, the information may have given Zuckerberg a competitive advantage. On the other hand, to those in the know, the database is fairly straightforward and likely to be reverse-engineerable by a competent internet programmer in a reasonable amount of time. In addition, it is not clear whether the Winklevosses took security measures to maintain the secrecy of the information. They allege Zuckerberg knew the information was confidential,

---

<sup>1</sup> *ConnectU, Inc. v. Facebook, Inc. et al.*, Case-Number 1:2007cv10593, Massachusetts District Court, filed March 28, 2007.



**Fig. 1** IBP’s domain model of trade secret law with the applicable factors for plaintiff (P) and defendant (D) in ConnectU versus Zuckerberg. Inapplicable factors are grayed out. Easy to prove factors are underlined

but there was no agreement not to disclose, or they would have said. As to a breach of confidential relationship, it was the Winklevosses who gave the information to Zuckerberg, presumably in negotiations to get him to agree to be their programmer on this project. Another question mark concerns whether Zuckerberg used improper means to get the information. It seems clear that he never intended to actually perform the work the Winklevosses wanted him to perform; that was deceptive. Though he had access to the ConnectU source code (i.e., tools), he very may well have developed the Facebook code completely independently, not using the plaintiff’s information at all, and it may have been quite different from the ConnectU code.

### 2.3 Normative Uncertainty

Normative uncertainty refers to uncertainty in the application of legal concepts to facts based on dialectical argumentation involving two kinds of inferences: subsumption (i.e., the fulfillment of sets of conditions) or analogy (i.e., a kind of statutory interpretation based on similarity to some precedent or rule). For example,

a normative issue is whether ConnectU's knowledge or expertise constitutes a trade secret. Subsumption does not work too well in this example since Sec. 757 Comment b disavows "an exact definition of a trade secret,"<sup>2</sup> listing six relevant factors, instead. Analogy fills the gap, however. If a prior court has held it sufficient to establish a trade secret that the plaintiff's product was unique in the marketplace even though it could be reverse-engineered, that is a reason for finding a trade secret in ConnectU versus Zuckerberg. There is uncertainty, however, due to questions like whether the present case is similar enough to the precedent to warrant the same outcome. If one were fairly sure that the factors relating to an element (or legal concept) apply and uniformly support one side, then one could tentatively decide the element is satisfied. But here, since the factors are only fairly certain or doubtful and have opposing outcomes, a litigator will probably decide that the various elements are still open (indicated in bold in Fig. 1).

## 2.4 *Moral Uncertainty*

Moral uncertainty is involved in assessing if a decision is "just," "equitable," or "right" based on whether applicable values/principles are in agreement with or in conflict with the desired legal outcome of the case. For instance, in trade secret law, considerations of unfair competition, breach of confidentiality, or innovation all favor a decision for the plaintiff, while access to public information and protecting employees' rights to use general skills and knowledge favor the defendant. Most of these values are at stake in the ConnectU scenario and they conflict. If the allegedly confidential information is reverse-engineerable, that would undermine a trade secret claim or limit its value. On the other hand, if Zuckerberg obtained the information through deception, that would support the claim. Thus, there are questions about the nature of the trade-off of values at stake and which critical facts make the scale "tip over." It makes sense to consider the balance of values at the level of the individual elements as well as of the claim; the values at stake and the facts that are critical will differ at these levels.

## 2.5 *Empirical Uncertainty*

Finally, there is empirical uncertainty relating to prior outcomes in similar scenarios. In common law jurisdictions, precedents are recorded in detail, including information about the cases facts, outcome, and the courts normative and moral reasoning. In the United States and elsewhere, this case data may be retrieved manually and electronically with the result that attorneys can readily access relevantly similar cases. The precedents are binding in certain circumstances (*stare decisis*), but even if

---

<sup>2</sup> Restatement of Torts (First) Section 757, Comment b.

not binding, similar cases provide a basis for predicting what courts will do in similar circumstances. Given a problem's facts, a litigation planner can routinely ask for the outcomes that similar cases predict for claims, elements, and issues in the current case. The degree of certainty of these predictions (or, respectively, the persuasiveness of an argument based on prior decisions) depends on the degree of factual, normative, and moral analogy among the problem and the multiple, often competing precedents with conflicting outcomes.

## ***2.6 Interdependencies***

The four types of uncertainty are clearly not independent (e.g., normative and empirical, both involve questions of degree of similarity.) There interactions can be quite complex, as illustrated by the role of narrative. At a factual level, narrative and facts may mutually reinforce or be at odds. A narrative that is plausible in light of specific facts may enhance a jury's disposition to believe those facts; it also raises expectations about the existence of certain other facts. For instance, if an attorney tells a "need code" story (i.e., Zuckerberg really needed to see ConflictU's code to create Facebook), one would expect to find similarities among the resulting codes. If it is a "delaying tactic" story, the absence of similarities between codes supports a conclusion that Zuckerberg may simply have been deceiving the Winklevosses to get to the market first. At the normative and moral levels, the narrative affects trade-offs among conflicting values that may strengthen one claim and weaken another. Trade secret misappropriation does not really cover the Delaying Tactic story for lack of interference with a property interest in confidential information.

## **3 Desirable Attributes for a Probabilistic Argument Model to Assist Litigation Planning**

The prior work and literature on evidential reasoning using Bayesian methods as well as BNs is extensive (see prominently, e.g., (Kadane et al. 1996)). This chapter, however, focuses more narrowly on how probabilistic graphical models can enrich the utility of qualitative models of argument for litigation planning. This section will explain a list of desirable attributes that a probabilistic graphical argument model should provide in order to be useful for the task. We believe that there are four such attributes, namely, (1) the functionality to assess argument moves as to their utility, including (2) the resolution of conflicting influences on a given proposition, (3) easy knowledge engineering, as well as (4) the functionality to reason with and *about* argument weights.

### 3.1 *Assessment of Utilities*

As a starting point, the model should provide for a way to assess the effect of argument moves on the expected outcome of the overall argument. In other words, for every available argument move, one would like to know its expected payoff in probability to convince the audience. Quantitative probabilistic models allow for changes in the model's configuration/parameters to be quantified by calculating the relevant numbers before and after the change.

Similarly, Roth et al. (2007) have developed a game theoretical probabilistic approach to argumentation using defeasible logic. There, players introduce arguments as game moves to maximize/minimize the probability of a goal proposition being defeasibly provable from a set of premises, which in turn are acceptable with some base probability. Riveret et al.'s game theoretical model of argumentation additionally takes into account costs that certain argument moves may entail, thereby assessing an argumentation strategy as to its overall "utility" (Riveret et al. 2008). As Riveret et al. point out, the concept of a quantitative "utility" is complex and needs to take into account many factors (e.g., risk-averseness of a player) and does not necessarily need to be based on quantitative probability. We do not take position with regard to these questions in this survey chapter. Rather, we observe that quantitative probabilistic models enable an argument formalism to be used for utility calculations in litigation planning. We then suggest how this combination of models can also take account of certain phenomena in legal reasoning, for example, through probability distributions over argument weights (see Sect. 6).

### 3.2 *Easy Knowledge Engineering*

One of the difficulties in probabilistic models is the need to gather the necessary parameters, that is, probability values. This problem has been of particular importance in the context of BNs, where a number of approaches have been taken. For example, Díez and Druzdzel have surveyed the so-called "canonical" gates in BNs, such as noisy/leaky "and," "or," and "max" gates (see Díez et al. 2006 with further references) which allow for the uncertainty to be addressed in a clearly delimited part of the network with the remaining nodes behaving deterministically. This greatly reduces the number of parameters that need to be elicited. In furtherance, the BN3M model by Kraaijeveld and Druzdzel (Kraaijeveld et al. 2005) employs canonical gates as well as a general structural constraint on diagnostic BNs, namely, the division of the network into three layers of variables (faults, evidence, and context). A recent practical evaluation of a diagnostic network constructed under such constraints is reported in (Agosta et al. 2008; Ratnapinda et al. 2009).

We believe that a formal model of argument lends itself well to the generation of BNs from metadata (i.e., from some higher level control model using fewer or easier-to-create parameters). Specifically for the Carneades model, this route has previously been examined in (Grabmair et al. 2010). We will give an overview of it

in Section 5.3 and will build on it in Sections 6 and 7. This is conceptually related to the works on translation of existing formal models into probabilistic graphical models. For an example of such a translation, namely, of first-order logic to BNs, and a survey of related work on so-called probabilistic languages, see (Laskey 2008).

An example of a law-related BN approach that appears not to support easy knowledge engineering is the decision support system for owner's corporation (i.e., housing co-op) cases of Condlie et al. (2010). There, a BN was constructed using legally relevant issues/factors as its random variables. A legal expert was entrusted with the task of determining the parameters. The overall goal was to predict the outcome of a given case in light of prior cases in order to facilitate pre-litigation negotiation efforts. While this is a good example of a legal decision support system, it employs no argumentation-specific techniques or generative meta-model. From the example BN in the article, it appears that the domain model is essentially "hardwired" into a BN and would need major knowledge engineering efforts to maintain if relevant legislation/case law (or any other kind of applicable concerns) were to change and the complexity of the network would increase. We believe that the construction of decision support BNs under the structural constraints, and control of a qualitative argument model provides a more efficient and maintainable way to knowledge engineer such systems. By contrast, admittedly, our approach still requires validation.

### 3.3 *Conflict Resolution and Argument Weights*

The argument model needs to account for pro and con arguments as it resolves conflicts in supporting and undermining influences on a given proposition. Since conflicting arguments are very common in legal scenarios, it is critical for a system to be able to tell how competing influences outweigh or complement each other. In a quantitative model, this can be done straightforwardly, for example, by computing an overall probability based on some weighting scheme. In this way, a strong pro argument will outweigh a weak con argument with some probability. In qualitative models, to the best of our knowledge, such conflict resolution needs to be established in a model-specific way. For example, in a qualitative probabilistic network (Wellman 1990), one could introduce a qualitative mapping for sign addition where a strong positive influence (++) outweighs a weak negative influence (-), leading to an overall weaker positive influence (+) (Renooij et al. 1999). Also, Parsons (Parsons et al. 1997) has produced a qualitative model which allows for some form of defeasibility in combining multiple arguments into a single influence.

Dung's prominent work on argumentation frameworks (Dung 1995) only uses only one kind of attack relation and thereby abstracts away the precise workings which lead one argument to attack/defeat another (explained, e.g., in Baroni et al. 2009). There, arguments are equally strong, and resolving the argument means finding certain kinds of extensions in the argument graph. For a brief probabilistic interpretation and ranking of different extensions by likelihood, see Atkinson and



Bench-Capon (2007). Value-based argumentation frameworks (Bench-Capon 2003) expand upon this model and implement conflict resolution not through weights but through a hierarchy of values which prioritize conflicting arguments depending on which values they favor.

Whether conflict resolution is undertaken through a hierarchy of values or numerical argument weights (such as in Carneades (Gordon and Prakken 2007)) or in some way, two major conceptual challenges need to be addressed in the legal context. First, the notion of conflict resolution or degrees of persuasion needs to interface with the concepts of a standard and burden of proof. There is ample literature on the topic both in formal models of argument (e.g., in argumentation frameworks (Atkinson and Bench-Capon 2007) including a general survey and specifically in Carneades (Gordon et al. 2009)) and elsewhere, such as in the work of Hahn and Oaksford (2007), who (in our opinion, correctly) point out that proof standards may be relative to the overall utility at stake.

For example, in ConnectU versus Zuckerberg, a judge may be more likely to grant the plaintiffs' request for injunctive relief if it were claimed only for a limited duration such as the time it would reasonably take for a competent programmer to reverse engineer the information once the plaintiffs' program was deployed. Such a solution would possibly maximize the realization of the values protected, protecting property interest in the invention for a limited time while allowing access to public information afterward.

Second, an argument model should allow for meta-level reasoning about the conflict resolution mechanism, that is, arguments about other arguments (e.g., about their weight/strength). This is our main motivation for the extensions suggested in Sect. 6.

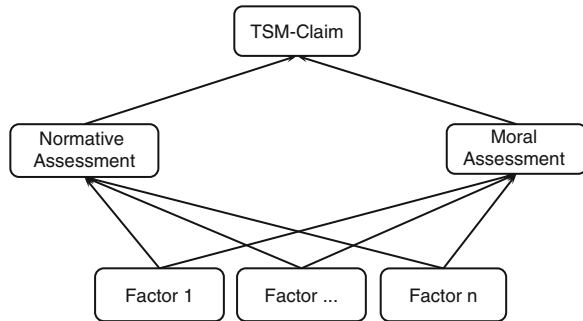
## 4 Sample Assessment of Graphical Models

This section will lead through a comparison of graphical models and their capabilities in the context of modeling the ConnectU versus Zuckerberg case. Connecting to our previous explanations of uncertainty, we systematize our legal problem by distilling it into the three major uncertainties: normative, moral, and factual. Empirical uncertainty is not explicitly represented, but rather affects both normative and moral uncertainty, where it flows into how decisions about precedents affect the argument. The overall outcome is influenced both by normative and moral assessments of the facts (Fig. 2). We will now expand this basic framework into a more detailed reflection of arguments.

### 4.1 A Graphical Structure of the Analysis

As we have seen in the IBP domain model sketched above, in order for a trade secret misappropriation claim to be successful, the information in question would

**Fig. 2** A basic graph for a legal decision problem



need to constitute a trade secret, and the information would need to have been misappropriated in some way. Assume that the primary legal issue before the court is just the former, that is, whether the source code and business documentation obtained by Zuckerberg constitute a trade secret.

A factor in the affirmative is that the code is part of a hitherto *unique product* (compare factors in Fig. 1). There is no service on the market yet with comparable prospects of success. The value underlying this factor is the *protection of property*, as the inventor of a unique product has an ownership interest in the information that is protected by the legal system.

On the other hand, a factor speaking against the trade secret claim could be the fact that the product is easily *reverse-engineerable*. After all, ConnectU (functionally similar to Facebook) is a straightforward online database-driven website application. After seeing it work, an internet programmer of average competence could recreate the functionality from scratch without major difficulty within a reasonable amount of time. The value underlying this conflicting factor is the protection of *public information*. The policy is that trade secret law shall not confer a monopoly over information that is publicly available and easily reproduced and hence in the public domain.

The two conflicting concerns influence both the normative and the moral assessments. In the normative part, *unique product* and *reverse-engineerable* essentially become two arguments for and against the legal conclusion that there is a trade secret. They, in turn, can be established from facts in an argumentative way. In the moral part, we have the conflicting interests behind these factors, namely, the desire to protect one's property and the goal to protect public information. However, these interests are not *per se* arguments in the debate of whether a moral assessment of the situation speaks for or against a trade secret misappropriation claim. Rather, they are *values* which are to be weighed against each other in a balancing act. Again, this determination is contingent on an assessment of the facts.

The overall structure just presented can be visualized as a directed graph model in Fig. 3. It shall be noted that this graph is an incomplete and coarse-grained representation of the concerns. Our intention is only to expand parts of the normative and moral layer in order to illustrate certain features of graphical probabilistic argument models in the subsequent sections.

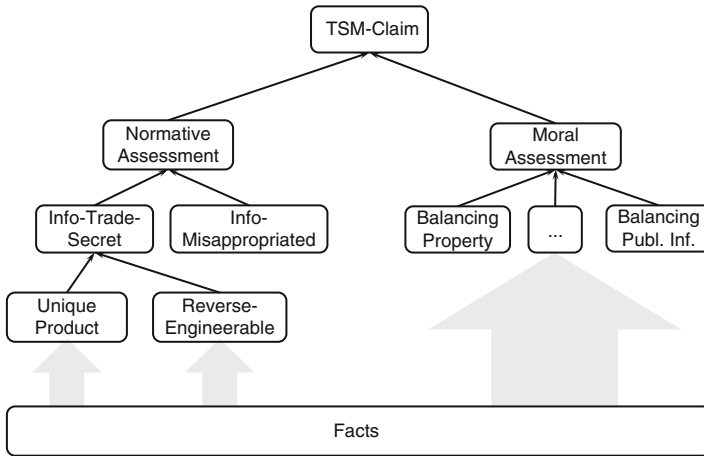


Fig. 3 The decision graph with moral and normative assessments expanded


### 4.2 Casting the Example into a Graphical Model

Figure 3 provides a basic graphical model of the decision problem in ConnectU versus Zuckerberg. On the normative side, there are two conflicting arguments (*unique product* and *reverse-engineerable*), each stemming from a certain factor and conflicting values, which underlie these factors. Overall, the two components are framed by the factual layer of the graph (shown as a single large node standing in for a set of individual facts nodes) feeding into the normative and moral parts, as well as the overarching conclusion, that is, the success of the trade secret misappropriation claim. A probabilistic interpretation of the graph implies that, for every arrow, our belief in the target proposition is influenced by our belief in the propositions that feed into it. Our belief in certain facts influences our belief in the status of the intermediate concepts involved. The way in which this influence takes place varies across the different kinds of uncertainty. For example, in factual uncertainty, a significant part of the reasoning may concern evidential questions of whether a certain fact is true (or not) and what its causal relations are to other uncertain facts. Here, and probably to a greater extent in normative uncertainty, multiple conflicting arguments are a typical occurrence and need to be resolved by some functionality of conflict resolution, as will be addressed further below.

### 4.3 Generic Bayesian Networks

This section will give a brief overview of Bayesian Networks (as introduced in Pearl et al. 1985) and a brief survey of how they can be used to model argumentation.

**Fig. 4** Pro and con arguments for trade secret issue



A Bayesian Network [BN] is a probabilistic graphical model based on a directed acyclic graph whose nodes represent random variables and whose edge structure encodes conditional dependencies between the variables. Each variable has a conditional probability function (typically represented as a table of conditional probabilities) defining the probability of the given node having a certain value conditioned on the values of its immediate parent variables. One can then compute a probability value of interest in the network by multiplying respective probabilities along a corresponding path. BNs are capable of both forward reasoning (i.e., determining the probability of effects given causes) and diagnostic reasoning (i.e., determining the probability of causes given effects). In the context of this chapter and at the current state of our research of applying BNs to argumentation, we only discuss the former.

From now on, we shall assume that the graph is completed in the interface of the facts layer to the assessment layer. One can take a straightforward approach and translate the graph in Fig. 3 into a BN. While the node and edge structure can be mapped over identically, one needs to define conditional probability tables. In our example, the main instances of interest are the conflicting arguments for *unique product* and *reverse-engineerability* (see Fig. 4).

Using a BN for representing argument graphs has a number of desirable features. First, it allows the calculation of probability values given its parameters, while qualitative models provide this functionality only to a very limited extent. In the context of argument models, this means that by changing the parameters or altering the graph in certain ways, one can calculate the increase or decrease in persuasiveness of a conclusion statement of interest. Second, probabilities are propagated through the network, thereby leading to a quantitative synthesis of pro and con influences on a given proposition. This is because, through the conditional probabilities, relative strength among competing arguments is expressed via the probability with which one set of arguments defeats another. This aspect of BNs as a representation technique addresses the first two of our desired attributes for an argument-based litigation support system (see Sect. 3).

An example of a system combining argument graphs and BNs is the work of Muecke and Stranieri (2007), who use a combination of a generic argument model and expert-created BNs to provide support for the resolution of property issues in divorce cases. The resulting system was intended to reflect the judicial reasoning to be anticipated if the case were to be litigated, also taking into account precedent cases in designing the argument graphs. However, our approach is not to be confused with work aiming to construct argument graphs from existing BNs (see, e.g., Keppens et al. 2012) for an approach to extracting an argument diagram from

BNs in order to scrutinize and validate evidential reasoning with the BN). Zukerman et al. (1998) developed a system which uses a combination of multiple knowledge bases and semantic networks to construct two BNs, one normative model and a user model. These are then used to gradually expand a goal proposition of interest into an argument graph (with premise propositions and inference links) until a desired level of probabilistic belief in the goal is established.

However, a difficulty is that BNs traditionally have proven to be difficult in terms of the required knowledge engineering. This is particularly true with regard to the elicitation of the necessary conditional probabilities (compare Diez et al. 2006). Also, although BNs reduce in parametric complexity compared to a full joint distribution, exact inference is still NP hard (see Cooper 1987), and individual conditional probability tables can become quite large very easily. We will now illustrate this rapid growth in the context of our example.

Let us assume, similar to Carneades (Gordon and Prakken 2007; Gordon et al. 2009), that statements can have three possible statuses: *accepted* (where the audience has accepted a statement  $p$  as true), *co-acceptable* (where the audience has accepted the statement's complement  $\bar{p}$  as true), and *questioned* (where the audience is undecided about whether to accept or reject a statement). If one applies this triad of possible statuses to the three statement sub-issues in Fig. 4, then we need to define  $3^3 = 27$  possible statuses (called "parameters") in the conditional probability table of *info-trade-secret*. For each of the possible three statuses of the conclusion *info-trade-secret*, we need to determine the probability of said status conditioned on every possible combination of the three statuses of each of *unique product* and *reverse-engineerable*.

A domain expert would need to provide all these probabilities and tackle the challenges this task entails, for instance:

- What is the probability of the conclusion being accepted if both the pro and con arguments are accepted by the audience and hence conflict? Which one is stronger and how much stronger?
- If neither argument is acceptable, is there a "default" probability with which the audience will accept the statement?
- If only the pro argument is accepted, what is the probability of the conclusion being accepted? Is this single accepted pro argument sufficient? Is there some threshold of persuasion to be overcome?
- Similarly, if only the con argument is accepted, is this sufficient to make the audience reject the conclusion? Does it make a difference whether the audience is undecided over a pro argument or whether it explicitly rejects it?

It is difficult for domain experts to come up with exact numbers for all these scenarios, irrespective of which argument model or which precise array of possible statuses is used. One may argue that the use of three statuses as opposed to a binary true/false scheme artificially inflates complexity to a power of three as opposed to a power of two (thereby leading to only  $2^3 = 8$  required probabilities). In scenarios where there are more than two arguments, however, the number of required probabilities outgrows the feasible scope very quickly. Also, once a table of

probabilities has been created, modifications like adding or removing arguments require significant effort.

In light of this difficulty, it is desirable to have a simpler mechanism for knowledge engineering when using BNs as a representation. One solution is to use an existing qualitative model of argumentation to generate BNs including their parameters. In such an approach, the superior “control” model (in our case, an argument model) steers the construction of the networks and the determination of its parameters. In argumentation, previous work on probabilistic semantics for the Carneades argument model (Grabmair et al. 2010) produced a formalism with which an instance of a Carneades argument graph can be translated into a BN. This method may resolve at least some of the issues enumerated above or, at least, provide a more usable representation.

## 5 Carneades

### 5.1 A Brief Introduction to the Carneades Model

This section gives a brief overview of the Carneades argument model as well as its BN-based probabilistic semantics and explains the benefits it provides over a generic BN model.

The Carneades argument model represents arguments as a directed, acyclic graph. The graph is bipartite, that is, both statements and arguments are represented as nodes in the graph. An argument for a proposition  $p$  (thereby also being an argument con  $\neg p$ ) is considered *applicable* if its premises hold, that is, if the statements on which it is based are *acceptable*. This acceptability is in turn computed recursively from the assumptions of the audience and the argument leading to a respective statement. For each statement, applicable pro and con arguments are assessed using formally defined *proof standards* (e.g., preponderance of the evidence, beyond a reasonable doubt), which can make use of numerical weights (in the zero to one range) assigned to the individual arguments by the audience.

In recent work on Carneades (Ballnat et al. 2010), its concept of argumentative derivability has been extended to an additional labeling method using the statuses of *in* and *out*, which we will use occasionally throughout this chapter. So far, we have spoken of a literal  $p$  being *acceptable*, which can alternatively be labeled as  $p$  being *in* and  $\neg p$  being *out*. In the opposite case, if  $\neg p$  can be accepted (i.e., is *in*),  $p$  is labeled *out*, whereas  $\neg p$  is labeled *in*. Finally, if  $p$  remains questioned, both  $p$  and  $\neg p$  are *out*. One can see that the new method is functionally equivalent to the original one yet entails a clearer separation of the used literals into positive and negated instances. This has been taken over into the BN-based probabilistic semantics for Carneades in (Grabmair et al. 2010).

**Fig. 5** A Carneades graph for the trade secret issue




Figure 5 recasts our trade secret example into a Carneades graph. Similar to the previous Fig. 4, three rectangular nodes represent the statements. The circular nodes stand for the arguments which use these statements as premises and are labeled with exemplary argument weights. Now, assume that the trade secret issue shall be decided based on a *preponderance of evidence* proof standards. Verbalized, this means that the side shall win who provides the strongest applicable argument. A statement is acceptable under the Carneades implementation of the preponderance of evidence standard if the weight of the strongest applicable pro argument is greater than the weight of the strongest applicable con argument (for a full formal definition, see Gordon et al. 2009). The inverse condition needs to be fulfilled for the statement to be rejected. If neither side can provide a trumping applicable argument, the statement remains questioned. The only constraint is that each weight is within the  $[0, 1]$  range. They represent a quantitative measure of the subjective persuasiveness assigned to the arguments by the audience, not probabilities, and hence do not need to sum to 1. Carneades makes the assumption that such numbers can be obtained and uses the set of real numbers as its scope, that is, there is no minimum level of granularity for the precision of the number.

Hence, in our example (Fig. 5), if both *unique product* and *reverse-engineerable* are acceptable or assumed true by the audience, the pro argument with a weight of 0.7 trumps the con argument with a weight of 0.5, and hence the statement *info-trade-secret* is acceptable.

## 5.2 Carneades Bayesian Networks

From the graphical display in Fig. 5, we can easily conceptualize a corresponding Carneades Bayesian Network [CBN] quantitatively emulating the same qualitative inferences. The structure would be identical, and we only would need to define probability functions/tables. These probabilities basically emulate the deterministic behavior of the Carneades model. Later, the model can be extended to support new phenomena of uncertainty, for example, probability distributions over assumptions. The following explanations are essentially a verbalized summary of the formalism given in (Grabmair et al. 2010), to which we refer the interested reader for greater detail.

First, we determine the probability parameters of the leaf nodes. Depending on whether our audience assumes *unique product* and/or *reverse-engineerability*

as accepted or rejected, or has no assumptions about them, the respective probability will be either 0 or 1 for the positive and negated versions of the respective statements. Second, the probability for each of the arguments being applicable will be equal to 1 if their incoming premise statement is acceptable and 0 otherwise. Third, the conditional probability table for *info-trade-secret* contains the probability of the statement being acceptable conditioned on the possible combinations of in/out statuses of the two arguments (i.e.,  $2 \times 2 = 4$  parameters). Here, the probability of the presence of a trade secret being acceptable for the audience is equal to 1 for the two cases where the pro argument based on *unique product* is applicable and 0 otherwise because the pro argument is required for acceptability and always trumps the con argument (based on *reverse-engineerable*) because  $0.7 > 0.5$ .

Notice that the just explained probabilities change once the Carneades argument graph changes. For example, if the weights of the arguments are changed such that the con argument strictly outweighs the pro argument, then the conditional probability table for *info-trade-secret* would need to be changed to assign a probability of 0 to the statement being *in* (or, more generally, the statement's acceptability) in case of conflicting applicable arguments.

This simple example illustrates how one can translate instances of a qualitative model of argumentation into a BN. Notice, however, that all probabilities we have defined so far are deterministic, that is, they are either equal to 1 or equal to 0. This makes the BN function exactly like the original Carneades argument graph. Once some part of the configuration changes, this change can be taken over into the BN by modifying the structure or conditional probability tables of the network. In this way, our qualitative model of argument becomes a control structure which takes care of constructing an entire BN. The required knowledge to be elicited from domain experts is hence limited to the information necessary to construct the Carneades argument graph, such as the assumptions made and the weights assigned by the audience.

Finally, the probabilistic semantics of Carneades allow for a conceptual distinction between the subjective weight of an argument (to the audience) and the probability of the argument being *applicable* given the probability of the argument's premises being fulfilled. For detail and related work, we refer the reader to (Grabmair et al. 2010).

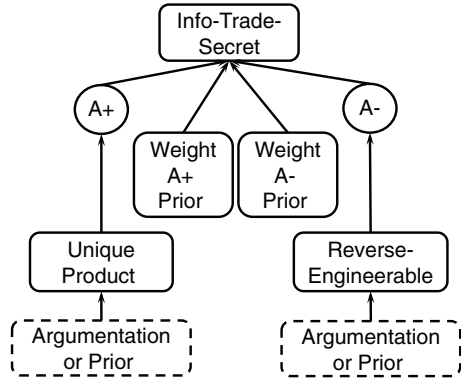
### 5.3 Carneades Bayesian Networks with Probabilistic Assumptions

Probabilistic semantics for a qualitative argument model provide little benefit if their parameters make the BN behave deterministically. A first extension introduces probabilistic assumptions and has been formally defined and illustrated in an example context in (Grabmair et al. 2010). Therefore, it shall only be explained briefly in this chapter.

Carneades allows for an audience to assume certain propositions as *in* or *out*. In realistic contexts, one might not be able to make such a determination about the



**Fig. 6** Example CBN with priors for arguments weights



audience. Instead, one may characterize an audience as being more or less likely to accept a certain statement than to reject it by some degree of confidence, or vice versa. One can capture this into a Carneades BN by allowing for prior probability distributions over the leaf node literal statements that the foundations of the argument and which may be assumptions. From there upward, the network propagates these values and provides for means to calculate the probability of success in the overall argumentative goal given the knowledge about the assumptions of the audience (see Fig. 6). The remaining parts of the network still are deterministically parametrized. This means that changes to the model graph (such as the introduction of new arguments or the alteration of weights) are still straightforward to reflect in the BN and can be assessed in terms of their impact on the probability of the success of the overall argument.

In a practical application, one can use available knowledge about the audience (e.g., the composition of a jury or a particular judge in a legal proceeding) to model these prior distributions.

### 5.4 Introduction to Argument Weights

Besides the assumptions about statements, Carneades associates with the audience the weights of the arguments employed in order to be able to resolve conflicting arguments using proof standards. Similar to probabilistic assumptions, one might not be able to determine that the audience will definitely prefer argument  $A+$  over argument  $A-$  (compare Fig. 6) by some margin. Instead, one may be able to say that the audience is more likely to assign high weight to an argument than a low one. Or one may say that the audience is so mixed that it is just as likely to give great weight to the argument as it is to disregard it. It is hence useful to model argument weights using prior probability distributions as well. Figure 6 displays our example network in such a structure. Notice that the statements *unique product* as well as *reverse-engineerable* can either be subject to further argumentation or assumed by the

audience to be *in*, *out*, or questioned according to some prior probability. For example, the audience can either assume the reverse-engineerability *as such* or it can be convinced to believe in it because of, for example, an IT expert testifying that an average programmer can reverse-engineer the database. In the latter case, one could expand *reverse-engineerable* into a subtree containing all arguments and statements related to this expert testimony.

## 6 Extension of Carneades to Support Probabilistic Argument Weights

We now formally define an extension/modification of the CBN model to support probability distributions over argument weights. We build on the formalism published in (Grabmair et al. 2010), which we have not reproduced in its entirety here. To make this chapter stand alone, we present slightly altered versions of three definitions used in (Grabmair et al. 2010) and accompany them with explanations about the symbols and constructs. These are the new argument weight random variable (representing the actual weights), the modified audience (with distributions over argument weights instead of static ones), and the enhanced probability functions (determining the probability of an argument having a certain weight and of its conclusion being acceptable). All other modifications are straightforward and have been omitted for reasons of brevity.

**Definition 1.** (argument weight variables). An *argument weight variable*  $w_a$  is a random variable with eleven possible values:  $\{0, 0.1, 0.2, \dots, 0.9, 1.0\}$ . It represents the persuasiveness of an argument  $a$ .

By contrast to the original Carneades models, we have segmented the range of possible weights using 0.1 increments. While the granularity of the segmentation is arbitrary, we consider a segmentation useful for purposes of this extension because (1) it is flexible yet allows to reasonably model a tie among weights of conflicting arguments (which would be harder if one were dealing with a continuous distribution) and (2) it allows the probabilities to be displayed in table form. This latter feature will become relevant in the next section, where we survey the possibility of yet another extension, subjecting the weight distributions to further argumentation. Next is our enhanced definition of an audience (compare to def. 12.3 in Gordon et al. 2009). For understanding,  $L$  stands for the propositional language in which the statements have been formulated.

**Definition 2.** (audience weight distributions). An *audience with weight distributions* is a structure  $\langle \phi, f \rangle$ , where  $\Phi \in L$  is a consistent set of literals assumed to be acceptable by the audience and  $f$  is a partial function mapping arguments to probability distributions over the possible values of argument weight variables. This distribution represents the relative weights assumed to be assigned

by the audience to the arguments. For a given set of weight values  $w_{a_1}, \dots, w_{a_n}$ , let  $f''(w_{a_1}, \dots, w_{a_n})$  be a partial function mapping the arguments  $a_1, \dots, a_n$  onto these values.

The main aspect of the extension is the modified probability functions given in the next definition. In the original definitions in (Grabmair et al. 2010), weights were not explicitly represented as random variables in the network. As we introduce them now, we alter the probability functions for arguments by making them conditioned on the values of their argument weight parent variables. We further assign probability functions for the weight variables themselves.

Carneades conceptualizes argumentation as an *argumentation process* divided into *stages*, which in turn consist of a list of arguments, each connecting premise statements to a conclusion statement. The central piece of the formalism is the *argument evaluation structure*, which contains the elements necessary to draw inferences, namely, the stage containing the arguments, the proof standards of the relevant statements, and, finally, an audience holding assumptions and assigning argument weights. Carneades features *in-out* derivability in (Ballnat et al. 2010) as explained above. When verbalized,  $(\Gamma, \phi) \vdash_{f''(w_{a_1}, \dots, w_{a_n}), g} p$  means that the arguments contained in the stage  $\Gamma$  and the set of assumptions by the audience  $\Phi$  argumentatively entail statement  $p$  (i.e.,  $p$  is *in*) given the argument weights determined by the partial function  $f''$  and the proof standards of the statements determined by the function  $g$ .

**Definition 3.** (modified probability functions). *If  $\mathcal{S} = \langle \Gamma, \mathcal{A}, g \rangle$  is an argument evaluation structure with a set of arguments  $\Gamma$  (called a “stage” in an “argumentation process”), a function  $g$  maps statements to their respective proof standard,  $\mathcal{A} = \langle \phi, f' \rangle$  is the current audience, the set of random variables  $V_{\mathcal{S}} = S \cup A \cup W$  consists of a set of statement variables  $S$ , a set of argument variables  $A$  as well as a set of weight variables  $W$ , and  $D_{\mathcal{S}}$  is the set of connecting edges, then  $\mathcal{P}_{\mathcal{S}} = \mathcal{P}_a \cup \mathcal{P}_s \cup \mathcal{P}_w$  is a set of probability functions defined as follows:*

*Definitions of  $\mathcal{P}_a = \{P_a | a \in A\}$  have been omitted for brevity reasons. See def. 13 in (Grabmair et al. 2010). It states that arguments are applicable with a probability of 1 if their premises hold.*


*If  $s = S$  is a statement variable representing literal  $p$  and  $a_1, \dots, a_n$  are its parent argument variables with corresponding weights variables  $w_{a_1}, \dots, w_{a_n}$ , then  $P_s$  is the probability function for the statement variable  $s$ :*

$$P_s(s = \text{in} | a_1, \dots, a_n, w_{a_1}, \dots, w_{a_n}) = \begin{cases} 1, & \text{if } (\Gamma, \Phi) \vdash_{f''(w_{a_1}, \dots, w_{a_n}), g} p, \\ 0, & \text{if } (\Gamma, \Phi) \not\vdash_{f''(w_{a_1}, \dots, w_{a_n}), g} p. \end{cases}$$

$$P_s(s = \text{out} | a_1, \dots, a_n, w_{a_1}, \dots, w_{a_n}) = 1 - P_a(s = \text{in} | a_1, \dots, a_n).$$

We further add probability functions for weight variables:

$$P_{w_a}(w_a = v) = f'(a, v).$$



**Fig. 7** Example probability distributions over weights (high peak, uniform, low shallow). The horizontal axes show argument weight values for an argument whose probabilities are displayed on the vertical axes

Finally, the total set of probability functions consists of one such function for each statement, argument, and weight:

$$\text{Let } \mathcal{P}_a = \{P_a | a \in A\}, \mathcal{P}_s = \{P_s | s \in S\} \text{ and } \mathcal{P}_W = \{P_{w_a} | a \in A\}.$$

In our example, we can use this function to automatically construct a conditional probability table for the statement *info-trade-secret*. Because there are 11 possible argument weights, its conditional side would contain  $11 \times 11 \times 2^2 = 484$  possible combinations of argument applicability and argument weights for both  $A+$  and  $A-$ . The literal is *in* with a probability of 1 in all combinations where (1)  $A+$  is applicable and  $A-$  is either not applicable or, in case, both arguments are applicable and (2) its weight  $w_{A-}$  is less than  $w_{A+}$ . This second case is new in this extension; in the basic model weights (and hence the outcomes of the aggregation of the arguments using the proof standard) are static. For example, using this extension, in Fig. 6, the probability of *info-trade-secret* being *in* would be determined by not only taking into account the probability of the audience accepting or rejecting reverse-engineerability or uniqueness of the product but also the possibility of either of the two arguments outweighing the other with some probability. Without the extension, one could only calculate the probability of the outcome given some static weights (e.g.,  $0.7 > 0.5$  as above).

For example, one can use a discretized bell-curve-like distribution over the 11 intervals (assuming it holds that  $\sum_{i=0}^{10} P(w = \frac{i}{10}) = 1$ ), depending on what one believes to be true about the audience). For example, in ConnectU versus Zuckerberg, we may be fairly certain that the uniqueness of the product is a persuasive argument and hence model a narrow bell-curve-like segmentation of probabilities with a high center (see Fig. 7 left). The reverse-engineerability, on the other hand, may only be low in weight, resulting in a shallow discrete curve-like distribution (see Fig. 7 middle). In a litigation planning context, these distributions correspond to the attorney's assumption about the audience. The narrower and higher the curve, the more certain she will be about exactly how persuasive an argument is. The more uncertain she is, the flatter the curve will become. If she cannot tell at all how strong her jury/judge will deem an argument, a uniform distribution can be employed (see Fig. 7 right).

The benefit of using probability distributions over argument weights as opposed to static weights is that the system can compute the probability of a statement being acceptable in the argument model by taking into account the probabilities with which conflicting arguments defeat each other under the different proof standards. For example, a strong argument may no longer deterministically defeat a weaker one, but only by some relatively high probability that can be determined by computing both weight prior distributions of the two arguments. Or an argument may fall below a minimum relevance threshold or not with some probability. These implicit computations can take place through Carneades' existing proof standard functions and the basic conditional probability functionality of BNs. Notice that this extension is set on top of the core formalism for CBNs in (Grabmair et al. 2010) and does not include prior probabilities over assumptions (compare Sect. 5.3). A comprehensive formalism including multiple extensions is intentionally left for future work.

## 7 Desiderata for Future Developments

In this section, we suggest two further extensions, namely, to influence argument weights by (a) making the argument weights subject to argumentation and, in order to connect to our previous explanations about moral uncertainty, (b) making the argument weights subject to influence from the results of the moral assessment of the case, that is, the results of the balancing of the principles involved. We will outline these extensions conceptually, justify why we consider them relevant, and hint at possible ways to implement them.

The common idea underlying the just suggested extensions as well as the following explanations is that argument weights are a crucial part of an argument model as they potentially enable the system to give definite answers in instances of conflicting arguments. Also, the weights stand for the realization that arguments differ in persuasiveness and are relative to a factual context and to abstract values. We do not take a position with regard to the philosophical and argument-theoretical implications of these notions, but only explore their ability to be modeled in this formal context.

### 7.1 *Weights Subject to Argumentation*

Arguments can be abstracted from and reasoned about. One can argue about how heavily a certain argument weighs in comparison to others and in light of the context in which it is employed. For example, in ConnectU versus Zuckerberg, if the plaintiff argumentatively establishes that the database is really difficult to reverse engineer and requires considerable skill, then the defendant's argument based on reverse-engineerability becomes weaker, despite still being applicable.

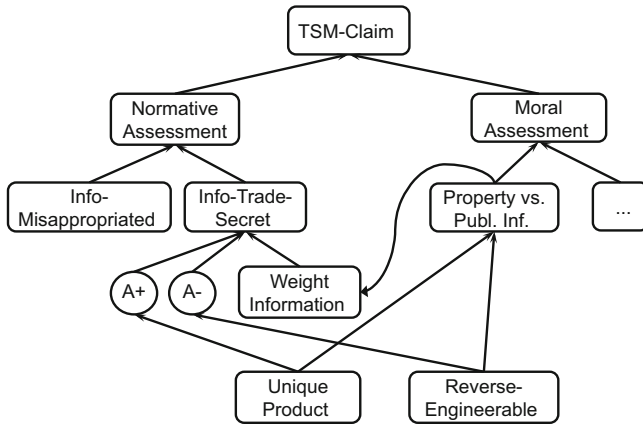
Once weights are represented explicitly in the argument graph, one can make them subject to argumentation. This presupposes the model's ability to express in its proposition information about the weights. What this information specifically means is dependent on how the argument model represents weights. In the Carneades-BN derivative, we have been describing in this chapter, this information about the weights takes the shape of normalized real numbers or, in the case of the extension, a prior probability distribution over a discretized range of weights.

As has been proposed above, a very strong argument could, for example, be modeled as a narrow bell-curve-like prior weight distribution with its peak centered at 0.8. On the other hand, a weak argument might be a similar bell curve but with its maximum probability located at only 0.3. Finally, an argument of average weight would be situated at 0.5. If one attempts to represent as a graph a meta-argument that a certain argument is either strong, weak, or average, the discourse essentially becomes an argument about a potentially unconstrained number of alternative choices as opposed to an argument about the acceptance of a certain statement as true. Carneades, with its current qualitative semantics, does not provide for such functionality.

In a generic BN representation, however, these choices can be represented as probability distributions conditioned on the status of the parent variables, where the parent variables here would represent, for example, the nodes in the fact layer or normative layer informing the weights. This essentially leads to an interdependence of argument formation and argument weights. We have just scouted the possibility of informing a choice of argument weights from fact patterns, which in turn can be argumentatively established. This can be considered as in part conceptually reminiscent of arguing from *dimensions* (Ashley et al. 1991). For example, in a trade secret misappropriation case, imagine that the manufacturer of a unique product has communicated production informations of the product to outsiders. This forms an argument against a trade secret claim because the manufacturer himself has broken the secrecy. However, it arguably makes a difference to whom he has communicated the secret (e.g., his wife, his neighbor, or a competitor's employee) and, in a second dimension, to how many people he has done so. This small example shows that the weight of an existing argument may be subject to argumentation itself based on the facts and the values at stake in the dispute. This leads over to our final suggested extension.

## 7.2 *Inform Weights from Values*

The values underlying the legal issues are another source of information for the question of which weights certain arguments shall carry. In *ConnectU versus Zuckerberg*, we have an argument based on the product's uniqueness furthering the protection of intellectual property interests, whereas an argument based on the product's reverse-engineerability goes against property interests in public information. This mechanism of arguments furthering values has been debated in



**Fig. 8** Conceptual argument graph with values informing weights of legal arguments

computational models of argument, perhaps most prominently in the context of the already mentioned value-based argumentation frameworks (Bench-Capon 2003), where a hierarchy of values resolves conflicts among arguments. In further work (Modgil 2006), this hierarchy can be made subject to argumentation as well, lifting the argument framework into a meta-representation. Feteris (2008) presents a pragma-dialectical model for how legal arguments for and against the application of rules can be weighted by underlying values and the factual context, connecting to prominent legal theories which we do not reference here for reasons of brevity.

One can adapt this notion in the context of probabilistic semantics for argument weights, by allowing the moral assessment of a case to influence the probability with which certain legal arguments will prevail over others. In other words, if a judge in *ConnectU* versus *Zuckerberg*, after considering the given facts, is of the opinion that property interests outweigh the interest in protecting public access (because, e.g., the evidence indicates that the social network idea is so unique and new that its farsighted inventor deserves a head start on its marketing), she might be more likely to accord greater weight to the legal arguments about the controversial issues which favor the property interest value (Fig. 8).

From a technical point of view, one can imagine that the decision about which principle is more important can determine a choice among different prior probabilities for the argument weights. A more satisfying solution, however, would be to not choose among predefined priors, but to have the priors of the weights be a qualitative outflow of the balancing decision between the two principles. If one value outweighs another by a larger or smaller margin, this margin should be reflected in the shape of the prior distribution of the argument weights. We are presently researching ways of modeling legal argumentation with values and have not yet conceived how to integrate it with a probabilistic argument model as described. Still, the idea of a value-based approach coheres with the intuitions of legal experts and has the potential of enriching the expressivity of probabilistic argument models.

## 8 Conclusions and Future Work

We have presented a survey of four kinds of uncertainties lawyers reason about and plan for in litigation and presented a realistic case example in trade secret law. Systematizing the uncertainties into a graphical model, we explained desirable attributes for probabilistic models of argument intended for litigation planning. Generic BNs as well as a probabilistic semantics for the Carneades model have been used as examples, and we have suggested extensions to the latter that allow more sophisticated reasoning with argument weights. We see great potential in combining quantitative models of uncertainty with qualitative models of argument and consider legal reasoning as an ideal testing ground for evaluating how qualitative models can be used most effectively to create quantitative models and how quantitative methods can provide steering heuristics and control systems for qualitative argumentation.

**Acknowledgments** We wish to thank the organizer Frank Zenker and the participants of the Workshop on Bayesian Argumentation at the University of Lund for the very inspiring discussion and all the feedback and questions we received in the context of our presentations and the workshop as a whole. We further thank Prof. Marek Druzdzel (Intelligent Systems Program, School of Information Science, University of Pittsburgh) and Collin Lynch (Intelligent Systems Program, University of Pittsburgh) for very helpful suggestions in the preparation of this chapter.

## References

- Agosta, J. M., Gardos, T. R., & Druzdzel, M. J. (2008). Query-based diagnostics. In M. Jaeger & T. D. Nielsen (Eds.), *Proceedings of the fourth European workshop on probabilistic graphical models (PGM-08), Hirtshals, Denmark, September 17–19* (pp. 1–8).
- Ashley, K. D. (1991). *Modeling legal arguments: Reasoning with cases and hypotheticals*. Cambridge, MA: MIT Press.
- Ashley, K. D., & Brüninghaus, S. (2006). Computer models for legal prediction. *Jurimetrics Journal*, 46, 309–352.
- Atkinson, K., & Bench-Capon, T. (2007). Argumentation and standards of proof. In *ICAIL 2007* (pp. 107–116). New York: ACM Press.
- Ballnat, S., & Gordon, T. F. (2010). Goal selection in argumentation processes: A formal model of abduction in argument evaluation structures. In *Computational models of argument proceedings of COMMA 2010* (pp. 51–62). Amsterdam: IOS Press.
- Baroni, P., & Giacomin, M. (2009). Semantics of Abstract Argument Systems. In I. Rahwan & G. Simari (Eds.), *Argumentation in Artificial Intelligence* (pp. 25–44). Berlin: Springer-Verlag.
- Bench-Capon, T. J. M. (2003). Persuasion in practical argument using value-based argumentation frameworks. *Journal of Logic and Computation*, 13(3), 429–448.
- Condliffe, P., Abrahams, B., & Zeleznikow, J. (2010). A legal decision support guide for owners corporation cases. In R. G. F. Winkels (Ed.), *Proceedings of Jurix 2010* (pp. 147–150). Amsterdam: IOS Press.
- Díez, F. J., & Druzdzel, M. J. (2007). Canonical probabilistic models for knowledge engineering (Technical Report CISIAD-06-01). UNED, Madrid (2006), Version 0.9, April 28 (2007), Available at <http://www.ia.uned.es/fjdiez/papers/canonical.pdf>



- Dung, P. M. (1995). On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2), 321–358.
- Feteris, E. T. (2008). Weighing and balancing in the justification of judicial decisions. *Informal Logic*, 28, 20–30.
- Cooper, G. F. (1987). *Probabilistic inference using belief networks is Np-hard* (Paper No. SMI-87-0195). Stanford: Knowledge Systems Laboratory, Stanford University.
- Gordon, T. F., & Walton, D. (2009). Proof burdens and standards. In I. Rahwan & G. Simari (Eds.), *Argumentation in artificial intelligence*. Berlin: Springer.
- Gordon, T. F., Prakken, H., & Walton, D. (2007). The Carneades model of argument and burden of proof. *Artificial Intelligence*, 171(10–11), 875–896.
- Grabmair, M., Gordon, T. F., & Walton, D. (2010). *Probabilistic semantics for the Carneades argument model using Bayesian networks: Proceedings of the third international conference on computational models of argument (COMMA)* (pp. 255–266). IOS Press: Amsterdam.
- Hahn, U., & Oaksford, M. (2007). The burden of proof and its role in argumentation. *Argumentation*, 21, 39–61.
- Kadane, J. B., & Schum, D. A. (1996). *A probabilistic analysis of the Sacco and Vanzetti evidence*. New York: Wiley.
- Keppens, J. (2012). Title: Argument diagram extraction from evidential Bayesian networks. *Artificial Intelligence and Law*, 20(2), 109–143.
- Kraaijeveld, P., & Druzdzel, M. (2005). GeNIeRate: An interactive generator of diagnostic Bayesian network models. In *16th international workshop on principles of diagnosis* (pp. 175–180).
- Laskey, K. B. (2008). MEBN: A language for first-order Bayesian knowledge bases. *Artificial Intelligence*, 172(2–3), 140–178.
- Modgil, S. (2006). Value based argumentation in hierarchical argumentation frameworks. In P. E. Dunne & T. Bench-Capon (Eds.), *Computational Models of Natural Argument, Proceedings of COMMA 2006* (pp. 297–308). IOS Press.
- Muecke, N., & Stranieri, A. (2007). An argument structure abstraction for Bayesian belief networks: Just outcomes in on-line dispute resolution. In *Proceedings of the fourth Asia-Pacific Conference on Conceptual Modeling* (pp. 35–40). Darlinghurst: Australian Computer Society, Inc.
- Parsons, S., Gabbay, D., Kruse, R., Nonnengart, A., & Ohlbach, H. (1997). Normative argumentation and qualitative probability. In *Qualitative and quantitative practical reasoning* (LNCS, Vol. 1244, pp. 466–480). Berlin/Heidelberg: Springer.
- Pearl, J. (1985, August 15–17). *Bayesian networks: A model of self-activated memory for evidential reasoning* (UCLA Technical Report CSD-850017). Proceedings of the 7th Conference of the Cognitive Science Society, University of California, Irvine, pp. 329–334.
- Ratnapinda, P., & Druzdzel, M. J. (2009). Passive construction of diagnostic decision models: An empirical evaluation. In *Proceedings of the international multicongress on computer science and information technology (IMCSIT-2009), Mragowo, Poland, October 12–14* (pp. 515–521). Piscataway: IEEE.
- Renooij, S., & van der Gaag, L. (1999). Enhancing QPNs for trade-off resolution. *Proceedings of the fifteenth conference annual conference on uncertainty in artificial intelligence (UAI-99)* (pp. 559–566). San Francisco: Morgan Kaufmann.
- Riveret, R., Prakken, H., Rotolo, A., & Sartor, G. (2008). Heuristics in argumentation: A game-theoretical investigation. In *Proceedings of COMMA 2008* (pp. 324–335). Amsterdam: IOS Press.
- Roth, B., Riveret, R., Rotolo, A., & Governatori, G. (2007). Strategic argumentation: A game theoretical investigation. In *Proceedings of the 11th international conference on artificial intelligence and law* (pp. 81–90). New York: ACM.
- Wellman, M. P. (1990). Fundamental concepts of qualitative probabilistic networks. *Artificial Intelligence*, 44, 257–303.
- Zukerman, I., Mcconachy, R., & Korb, K. (1998). Bayesian reasoning in an abductive mechanism for argument generation and analysis. In *Proceedings of the fifteenth national conference on artificial intelligence* (pp. 833–838). Menlo Park: AAAI.