Legal Reasons for Floating Conclusions

Giuseppe PISANO [a,1], Henry PRAKKEN [b] and Giovanni SARTOR [a]
[a] *Law Faculty - CIRSFID, University of Bologna, Italy*
[b] *Department of Information and Computing Sciences, Utrecht University, The Netherlands*

Abstract. We formally characterise legal reasons in argumentation-based models of (legal) reasoning, with particular attention to floating conclusions, i.e., conclusions supported by incompatible arguments. Our account captures the intuition that such a conclusion may be justified on the basis of incompatible alternative explanations, each of which independently provides a reason for it.

To this end, we generalise our *ASPIC*$^+$-based model of cause-in-fact in legal cases by developing notions of justification for arguments and propositions under grounded and preferred semantics, and by defining our notion of a reason in terms of these justifications.

Keywords. Legal argument, reasons, causation, floating conclusions

1. Reasons and Floating Conclusion in the Law

While many works exist on reasons in practical and legal philosophy [1,2,3], there is no general formal account of what counts as a reason within legal argumentation, i.e., what factual or legal premises contribute to determining a legal conclusion, providing effective support to that conclusion. Our analysis aims to cover this gap by providing a notion of reasons based on the one of cause given in [4]. A key feature in our approach is the consideration of cases in which undecided conflicts exist between arguments which, directly or indirectly, support the same legal conclusion. In the AI and Logic community this phenomenon is treated under the heading of "floating conclusions" [5,6,7,8]. Floating conclusions are common in law, since lawyers often frame arguments based on incompatible uncertain premises, affirming that the conclusion they desire follows whichever premise is adopted, so that a choice between them is not needed.

In common law, alternative pleading allows a party to allege two or more claims which are inconsistent with each other. For example, a plaintiff in a car accident does not know whether the evidence will lead the jury to conclude that the defendant ran into the plaintiff intentionally or by mistake. Thus, the plaintiff is allowed to plead the inconsistent claims that defendant committed the intentional (and not negligent) tort of battery and that the defendant was negligent.

It is also usually assumed that, when facing incompatible factual pleadings, a judge does not necessarily need to identify which alternative occurred. Such an identification is not needed if the evidence proves that one or the other must have occurred and both pro-

[1] Corresponding Author: Giuseppe Pisano, g.pisano@unibo.it.

duce exactly the same operative judgment, or that one provides for a lesser outcome that is somehow included in the outcome of the other. In US law a reference is provided by Ashwander v. Tennessee Valley Auth., 297 U.S. 288 (1936), where the Court argued that to obtain equitable relief in a case possibly involving fraudulent intent or legal duress, the plaintiff did not need to show one of them; it was sufficient to show breach of trust or duty.

Cases in which judges state a conclusion equally supported by alternative possible facts can also be found in UK law. For instance, in Lancashire County Council v B ([2000] 2 AC 147), the House of Lords held that the injuring of a child by a carer could be established even if there was uncertainty on which one of his three carers had caused the injury, because under each alternative the harm was attributable to the care received by the child.

Civil law courts reason similarly: when alternative factual claims entail the same legal outcome and one of them must have occurred the court need not determine which (Cass. civ., 23 September 2021, n. 25856). The Dutch Supreme Court applied the same reasoning in a murder case (Breda Six, ECLI:NL:HR:2017:3189).

A similar approach may be found in interpretation, where, following a legally minimalist approach (see [9]), a court may refrain from choosing between competing interpretations when the choice is unnecessary for deciding the case. For instance, in Privinvest Shipbuilding Sal (Holding) and Others v Filipe Jacinto Nyusi [2024] EWCA Civ 184, the judges argued that an interpretive issue did not need to be addressed, as it would have made no difference to the decision (paragraph 58).

Another example of this reasoning is Judge Breyer's approach to Morse v. Frederick, 551 U.S. 393 (2007), a case concerning a student's First Amendment rights and a school principal's qualified immunity. The student had been punished for displaying a flag with the words "Bong hits 4 Jesus" (where "Bong hits" could be understood as an invitation to smoke marijuana). Judge Breyer argued that there was no need for the court to decide (as it did) whether under US law the school has the right to restrain the student's freedom of expressing himself in this way, since the principal would be shielded from liability even if the school did not have such a right, based on the doctrine that one cannot be made liable for violating a constitutional right when there is uncertainty on the interpretation of that right.

It must be noted that judges tend to accept a floating conclusion only when they are satisfied that at least one of the alternative factual or interpretive claims holds, or at least where no defensible arguments have been proposed which challenge all such claims (without relying on them). This fits Horty's observation that conflicts between arguments may lead to doubts on the merit of all of them [6].

Our contribution is thus twofold. First, we define justifications for arguments and propositions in terms of minimal (strongly) admissible sets [10,4]. Since floating conclusions are not supported under grounded semantics, we work with preferred semantics, where our justifications are close to the minimal sufficient explanations of Borg & Bex [11] for sceptically accepted formulas, including floating ones. Our account mainly differs in using strong admissibility when possible, which excludes defenders that play no active role. Second, and mainly, we define legal reasons as propositions whose contrary, added as an axiom, would undo a justification, generalising our account of cause-in-fact [4]. Reasons thus include premises, intermediate conclusions and incompatible alternatives, as well as absent exceptions, which are not captured by existing argument-based

explanations, since these only contain (elements of) arguments for the conclusion or their defenders [11,12].

In Section 2 we provide the formal background for our analysis, referring to both abstract and structured argumentation. Section 3 provides the notions of weak and strong explanation of an argument, on the basis of which we develop concepts of justification for arguments and propositions. Section 4 introduces our definition of legal reason and in Section 5 we illustrate it with two legal examples.

2. Formal Preliminaries

We present our formal preliminaries, being the theory of abstract argumentation frameworks [13] and the *ASPIC*$^+$ framework for structured argumentation [14].

2.1. Abstract Argumentation Frameworks

An *abstract argumentation framework* [13] is a pair $AF = (\mathscr{A}_{AF}, \mathscr{D}_{AF})$, where $\mathscr{A}_{AF}$ is a set of arguments and $\mathscr{D}_{AF} \subseteq \mathscr{A}_{AF} \times \mathscr{A}_{AF}$ is a relation of defeat.[2] If A defeats B but B does not defeat A we say that A *strictly defeats* B. We write $A \in AF$ as shorthand for $A \in \mathscr{A}_{AF}$ and we will omit the subscripts if there is no danger of confusion. We will sometimes in text present an AF as $A \leftarrow B \leftrightarrow C$, to denote that $\mathscr{A} = \{A,B,C\}$ and $\mathscr{D} = \{(B,A),(B,C),(C,B)\}$.

Let $S \subseteq A$. Then S is *conflict-free* if no member of S defeats a member of S and S (directly) *defends* $A \in \mathscr{A}$ if for all $B \in \mathscr{A}$: if B defeats A, then some $C \in S$ defeats B. Moreover, S is *admissible* if it is conflict-free and defends all its members and S is *strongly admissible* if every $A \in S$ is defended by some $S' \subseteq S \setminus \{A\}$ which is strongly admissible [15].

To illustrate, let $AF = A \leftrightarrow B$. Then $\{A\}$ is admissible since it is conflict-free and A defends itself against B. Likewise for $\{B\}$. However, $\{A\}$ is not strongly admissible since A is not defended by $\{\}$ even though $\{\}$ is (trivially) strongly admissible. Likewise for $\{B\}$. So for admissibility an argument can successfully defend itself but for strong admissibility it must be defended by other arguments whose own defence does not depend on it.

Finally, $E \subseteq A$ is a *preferred extension* of AF if E is a $\subseteq$-maximal admissible set, and E is the *grounded extension* of AF if E is the subset-maximal strongly admissible subset of $\mathscr{A}$. The grounded extension is always unique while there can be multiple preferred extensions. We denote with $\mathscr{P}(AF)$ the set of preferred extensions of AF.

In Figure 1, $\{a,c,d\}$ is both the grounded and the (in this case unique) preferred extension. If we also let b defeat d, then the grounded extension is $\emptyset$, while there are two preferred extensions $\{a,c,d\}$ and $\{b\}$.

2.2. ASPIC$^+$

The *ASPIC*$^+$ framework [14] defines notions of argument, attack and defeat and thus generates abstract argumentation frameworks. For present purposes an informal sum-

[2] Dung used the term 'attack' but since we want to instantiate it with the *ASPIC*$^+$ defeat relation, we rename it to 'defeat'.

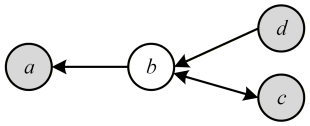


Figure 1. Illustrating preferred and grounded semantics.

mary suffices. *ASPIC*$^+$ relies on a logical language $\mathscr{L}$ (in this paper a set of propositional literals) over which a contrariness function is defined, which includes negation as follows: $\varphi \in \overline{\neg\varphi}$ and $\neg\varphi \in \overline{\varphi}$. As for notation, $-l$ denotes $\neg l$ if l is atomic and it denotes l' if l is of the form $\neg l'$. An *ASPIC*$^+$ theory consists of a set $\mathscr{R}$ of strict and defeasible inference rules over $\mathscr{L}$, and a knowledge base $\mathscr{K} \subseteq \mathscr{L}$ of unattackable *axioms* and attackable *ordinary premises*. Strict rules express deductive inferences, whereas defeasible rules only create a presumption that their conclusions hold.

Arguments are constructed by applying inference rules to premises in K and to conclusions obtained in earlier inference steps. Each inference step applies a strict rule $p_1, \ldots, p_n \rightarrow q$ or a defeasible rule $p_1, \ldots, p_n \Rightarrow q$. Here, each p_i is either a premise in K or the conclusion of a preceding inference step, and q is the conclusion of the current step. Arguments can thus be represented as trees (or, more generally, directed acyclic graphs) of inference steps. The leaves are the argument's premises, the root is its final conclusion, and the remaining nodes are intermediate conclusions. Subtrees of an argument are its subarguments.

Arguments can be *attacked* in three ways: on their ordinary premises (*undermining attack*), on their defeasible inference steps (*undercutting attack*) and on their defeasibly derived conclusions (*rebutting attack*). An undermining attacker concludes with a contrary of an ordinary premise of the attacked argument; an undercutting attacker concludes with the negation of the name of a defeasible inference rule used in the attacked argument; and a rebutting attacker concludes with a contrary of a defeasibly derived conclusion of the attacked argument. Attacks can be *direct* (on an argument itself) or *indirect* (on one of its proper subarguments). In general, *ASPIC*$^+$ distinguishes between attack and defeat relations between arguments, where the defeats are the attacks that succeed according to some preference relation on the set of all arguments. In this paper we assume for simplicity that there are no preferences, so all attacks succeed as defeats.

By providing arguments and defeat relations, *ASPIC*$^+$ generates abstract argumentation frameworks, and the acceptability status of these arguments can be evaluated by semantics for AFs. We denote by AF_{AT} the argumentation framework obtained from the argumentation theory AT.

Floating conclusions may arise within such frameworks.

Definition 1. [**Floating Conclusions**] Proposition φ is a floating conclusion relative to argumentation framework AF iff (a) every preferred extension of AF contains an argument for φ, but (b) no argument for φ is contained in every preferred extension of AF.

Note that this definition entails that no argument for a floating conclusion φ can be included in the grounded extension. If we assume that floating legal conclusions can be justified, we have to move from grounded to preferred semantics.

Example 1. Consider, for instance, the following knowledge base, where *Intention* and *Negligent* are assumed to be mutually contrarian (*Intention* $\in \overline{Negligent}$ and *Negligent* $\in$

$\overline{Intention}$):

$$\mathscr{K} = \{\ Harm,\ Reckless,\ Negligent\ \}$$
$$\mathscr{R} = \{\ r_1 : Harm, Intention \Rightarrow Liable, \quad r_2 : Harm, Negligent \Rightarrow Liable,$$
$$r_3 : Reckless \Rightarrow Intention\ \}$$

We can then build the following arguments:

$$A_1 = Harm \qquad A_2 = Reckless \qquad A_3 = Negligent$$
$$A_4 = A_2 \Rightarrow_{r_3} Intention \qquad A_5 = A_1, A_4 \Rightarrow_{r_1} Liable \qquad A_6 = A_1, A_3 \Rightarrow_{r_2} Liable$$

Note that A_3 and A_4 are in conflict, since *Intention* and *Negligent* are mutually contrarian. Consequently, A_3 attacks A_5 through its subargument A_4, while A_4 attacks A_6 through its subargument A_3. Accordingly, there are two alternative preferred extensions $P_1 = \{A_1, A_2, A_4, A_5\}$ and $P_2 = \{A_1, A_2, A_3, A_6\}$.

Both sets include an argument for *Liable*, namely A_5 and A_6, but neither of these arguments is included in the grounded extension because of the undecided conflict between *Intention* and *Negligent*.

3. The Explanations and Justifications of Arguments and Their Conclusions

Our analysis will be based on the idea of an explanation of an argument. We can distinguish two notions of an explanation within a given argumentation framework: a *strong explanation*, which explains the acceptance of an argument through a well-founded chain of defence, in which no argument relies on itself for its defence; and a *weak explanation*, which explains the acceptance of an argument through a set of arguments that collectively defend themselves, with arguments possibly contributing to their own defence. Weak explanations correspond to the compact explanations of [10], and strong explanations are their counterpart in terms of strong admissibility [15]. The weakness of weak explanations emerges in cases in which an argument A within the explanation both defeats and is defeated by an argument B (which is not strictly defeated by other arguments in the explanation), so that it remains unsettled which one of A or B should be endorsed. In terms of extensions, a weak explanation explains why an argument is in some given extension while leaving open that the argument is not in all extensions.

Definition 2. [**Strong and Weak Explanation**] Relative to the argumentation framework $AF = (\mathscr{A}_{AF}, \mathscr{D}_{AF})$

- S is a *strong explanation* for A iff S is a subset-minimal strongly admissible subset of $\mathscr{A}_{AF}$ containing A.
- S is a *weak explanation* for A iff S is a subset-minimal admissible subset of $\mathscr{A}_{AF}$ containing A.

Example 2. In Figure 1, $\{a,d\}$ is a strong (and also weak) explanation of a, while $\{a,c\}$ is only a weak explanation. The reason for the latter is that, although a is defended by $\{a,c\} \setminus \{a\} = \{c\}$, the set $\{c\}$ is not strongly admissible since c is not defended by $\{c\} \setminus \{c\} = \{\}$.

By the justification of argument A, we mean a set of arguments that is minimally sufficient to support the acceptance of A relative to an argumentation framework. The definition is intended to apply uniformly to both grounded and preferred semantics. In particular, when A is contained in a strong explanation, A belongs to the unique grounded extension and, consequently, to every preferred extension. In this case, the strong explanation suffices to justify A. If no strong explanation for A exists, justification is instead obtained by collecting weak explanations that, taken together, account for the acceptance of A across all preferred extensions. Since an argumentation framework may have infinitely many preferred extensions, this collection is not required to be finite.

Definition 3. [**Justification of an Argument**] Relative to the argumentation framework $AF = (\mathscr{A}_{AF}, \mathscr{D}_{AF})$, a set $\mathscr{S}$ is a justification for an argument A iff:

- $\mathscr{S} = \{S\}$, where S is a strong explanation for A, or
- no strong explanation for A exists, and $\mathscr{S} = \{S_1, \ldots S_n, \ldots\}$ is a subset-minimal set such that $S_1, \ldots S_n, \ldots$ are weak explanations for A, $S_1 \subseteq P_1, \ldots S_n \subseteq P_n, \ldots$ and $\{P_1, \ldots P_n, \ldots\} = \mathscr{P}(AF)$.

If, relative to argumentation framework AF, a justification $\mathscr{S}$ exists for an argument A, we say $\mathscr{S}$ is an AF-justification for A, and that A is AF-justified.

The notion of justification can subsequently be lifted from arguments to their conclusions. In particular, if a proposition φ is the conclusion of an argument A that is itself justified, explaining the justification of φ reduces directly to explaining the justification of A. This ensures that argument justification is recovered as a special case of proposition justification. If there is no single justified argument concluding for φ, the justification of φ may still be established when φ is supported in every preferred extension, although possibly by different arguments.

Definition 4. [**Justification of a Proposition**] Relative to the argumentation framework $AF = (\mathscr{A}_{AF}, \mathscr{D}_{AF})$, a justification of a proposition φ is

- an AF-justification $\mathscr{S}$ of an argument A for φ, if such a justification exists. Else:
- a subset-minimal set $\mathscr{S} = \{S_1, \ldots S_n, \ldots\}$ such that $S_1, \ldots S_n, \ldots$ are weak explanations for arguments $A_1, \ldots A_n, \ldots$ with $Conc(A_1) = \varphi, \ldots Conc(A_n) = \varphi, \ldots$ and such that $S_1 \subseteq P_1, \ldots S_n \subseteq P_n, \ldots$ where $\{P_1, \ldots P_n, \ldots\} = \mathscr{P}(AF)$.

If, relative to argumentation framework AF, a justification $\mathscr{S}$ exists for a proposition φ, we say that $\mathscr{S}$ is an AF-justification for φ and that φ is AF-justified.

Definition 4 accounts for cases of floating conclusions. In fact, through the second case of the definition, such floating propositions can be provided with a justification by collecting weak explanations for the different arguments supporting φ across the preferred extensions. Requiring $\mathscr{S}$ to be subset-minimal ensures that only the weak explanations necessary to account for the support of φ across all preferred extensions are retained.

Example 3. Consider again Example 1. The proposition *Liable* is a floating conclusion, since it is supported in every preferred extension, although by different arguments. In particular, A_5 concludes *Liable* in P_1, while A_6 concludes *Liable* in P_2, and neither argument is justified on its own.

$\{A_4, A_5\}$ is a weak explanation for A_5, and $\{A_3, A_6\}$ is a weak explanation for A_6. Since these are the only explanations, together they form the subset-minimal set $\{\{A_4, A_5\}, \{A_3, A_6\}\}$, thus constituting a justification for *Liable*.

4. From Justifications to Legal Reasons

In this section we shall move from the notion of a justification to the notion of a reason, namely, any proposition that contributes to explaining why a proposition is justified. Our model will characterise a reason for a literal ψ as any proposition φ such that, if its complement $(-\varphi)$ were true (if it were assumed as an undefeasible axiom), a justification for ψ would no longer hold. This covers both the case in which φ is included in (an argument in) the justification and the case in which $-\varphi$ would be included in a justified counterargument against that justification of ψ. This notion of a reason generalises the concept of a cause-in-fact we proposed in [4], which formalises the NESS concept of causality in fact [16].

Our notion of a reason is based on the idea of an axiom-expansion [4], which consists in inserting a set of axioms (indefeasible premises) in the argumentation theory. The axioms will defeat any incompatible argument, preventing their inclusion in any extension of the resulting argumentation framework.

Definition 5. [**Axiom-expansion**] The *axiom expansion* $AT + L$ *of an argumentation theory* $AT = (\mathscr{R}, \mathscr{K})$ with intervention L is defined as $AT + L = (\mathscr{R}, \mathscr{K} \cup L)$, where $L \subseteq \mathscr{L}$ is a set of axioms.

Definition 6. [**Legal Reason**] Literal φ is a *reason* for a literal ψ relative to argumentation framework AF_{AT} iff there exist

- an argument A for φ for which there is a strong or weak explanation relative to AF_{AT};
- an AF_{AT}-justification $\mathscr{S}$ for ψ;
- a set $L \subseteq \mathscr{L}$ containing an element $\chi \in \overline{\varphi}$,
- an argument B such that $L \subseteq Sub(B)$, B attacks an argument $C \in \bigcup \mathscr{S}$ and B is AF_{AT+L}-justified.

There are two possible ways in which an argument B may attack an argument $C \in \bigcup \mathscr{S}$: either B is a literal $-\varphi$ where φ is a conclusion of an argument in an explanation $S \in \mathscr{S}$ or B is an argument whose premises include all premises in L.

Example 4. Let us consider a simple example, dealing with liability for harm.

$$
\begin{aligned}
K = \{\ & Harm,\ MobilePhone,\ HighSpeed,\ \neg Necessity \quad \} \\
R = \{\ & r_1 : Harm, Intention \Rightarrow Liable, \quad r_2 : Necessity \Rightarrow \neg r_1, \\
& r_3 : Reckless \Rightarrow Intention, \quad r_4 : HighSpeed \Rightarrow Reckless, \\
& r_5 : MobilePhone \Rightarrow Reckless \quad \}
\end{aligned}
$$

With these rules and premises we can generate the following arguments:

$A_1 = Harm$ $A_2 = MobilePhone$ $A_3 = HighSpeed$
$A_4 = \neg Necessity$ $A_5 = A_3 \Rightarrow_{r_4} Reckless$ $A_6 = A_2 \Rightarrow_{r_5} Reckless$
$A_7 = A_5 \Rightarrow_{r_3} Intention$ $A_8 = A_6 \Rightarrow_{r_3} Intention$ $A_9 = A_1, A_7 \Rightarrow_{r_1} Liable$
$A_{10} = A_1, A_8 \Rightarrow_{r_1} Liable$

Using Definition 4, we obtain that *Liable* is a justified conclusion having two different justifications based on the two strong explanations for A_9 and A_{10}, namely $\{\{A_9\}\}$ and $\{\{A_{10}\}\}$. By applying Definition 6, we obtain the following reasons: *Harm*, *Intention*, *Reckless*, *HighSpeed*, *MobilePhone*, $\neg$*Necessity*. Each of these propositions is indeed a determinant of liability, in the sense that its negation would remove one justification for it. Importantly, this also captures alternative reasons. For instance, *MobilePhone* is a reason for liability even though without it there would still be liability. Similarly, $\neg$*Necessity* is also a reason for liability because expanding the theory with *Necessity* would remove (through r_2) at least one (and indeed both) of its justifications.

Example 5 (Back to Example 1)**.** Let us now go back to Example 1, where there was uncertainty on negligence or recklessness (as a kind of intention). In that framework *Liable* is a floating conclusion, as no argument for *Liable* has a strong explanation. However, *Liable* still is a justified conclusion, as a (different) argument for it exists in each preferred extension. Reasons for liability are *Harm*, *Intention*, *Reckless*, *Negligent*, which is what we expected. In particular, both *Negligent* and *Reckless* count as (alternative) reasons for liability.

5. Modelling Legal Examples

In this section we shall provide two legal examples of our approach, one dealing with an interpretation issue and the other with an evidence issue.

5.1. Interpretation

We refer to section 123 of the 1996 UK Employment Rights Act, where there was interpretive uncertainty on whether the term "loss" covered only monetary losses, or also injury to feelings (the issue was later settled in Dunnachie v Kingston-upon- Hull City Council [2004] UKHL 36). Two incompatible interpretations were available, Int_1 (only monetary harm is compensable) and Int_2 (also injury to feelings is compensable). Let us assume that an employee has suffered a monetary loss since their salary was not paid after the dismissal, and this is all they are asking.

$K = \{\ Int_1,\ Int_2,\ MissedSalary\ \}$
$R = \{\ r_1 : Int_1, MonetaryLoss \Rightarrow CompensableMonetaryLoss,$
$r_2 : Int_1, InjuryToFeelings \Rightarrow \neg CompensableInjuryToFeelings,$
$r_3 : Int_2, MonetaryLoss \Rightarrow CompensableMonetaryLoss,$
$r_4 : Int_2, InjuryToFeelings \Rightarrow CompensableInjuryToFeelings,$
$r_5 : MissedSalary \Rightarrow MonetaryLoss,$
$r_6 : Humiliation \Rightarrow InjuryToFeelings\ \}$

Given this knowledge base, it seems that *CompensableMonetaryLoss* should be certain (justifiable) since it follows from both incompatible interpretations. In fact, there is no need for the judge to make a choice between the two interpretations if the employee is only asking for monetary compensation. However, no argument for *CompensableMonetaryLoss* is included in the grounded extension, nor in every preferred extension (given the incompatibility between Int_1 and Int_2, which defeat each other). According to our Definitions 4 and 6, by contrast, *CompensableMonetaryLoss* is a justified conclusion, and each of the two interpretations (Int_1 and Int_2) is a reason for it (together with *MissedSalary*) .

Consider now the case of a different employee who is asking for compensation not only for their monetary loss but also for their humiliation (now K also includes $Humiliation$). According to our Definitions 4 and 6, still only *CompensableMonetaryLoss* is justified, since $CompensableInjuryToFeelings$ is not included in every preferred extension.

5.2. Dutch Supreme Court

We next analyse the Dutch Supreme Court judgement from 2017 in a murder case mentioned in the introduction. The case was about the murder of the mother of the owner of a Chinese restaurant at a time when the restaurant was closed. Six persons had been convicted in 1993 for the murder. The crucial evidence was that three of the convicted had confessed to the murder. In 2009 the three other convicted persons asked for a retrial because of new evidence. This retrial took place in 2013 at an appellate court.

One issue was that at the crime scene a blood drop had been found that turned out not to be of any of the six accused but of an unknown person of Asian origin. The main question was whether the blood drop was crime-related or not, for which the drying time of the blood drop appeared to be relevant. The defence argued that if it was crime-related, then this would point at another killer than one of the six accused so it would cast doubt on the reliability of the three confessions.

In Dutch criminal law, decisions of an appellate court on matters of fact can only be overturned by the Supreme Court if the appellate court's reasoning is incomprehensible, otherwise the Supreme Court has to accept the facts as deemed proven by the appellate court. The Supreme Court, after stressing this, essentially interpreted the court of appeal's reasoning as follows. The appellate court reasoned that two bodies of evidence E_1 and E_2 supported conflicting arguments about whether the blood drop was dried or not. The appellate court then reasoned that in both cases it is not certain that the blood drop was crime-related, so in both cases the blood drop does not cast doubt on the reliability of the confessions. The Supreme Court then ruled that "This judgement is – whatever may be said of the Court's considerations regarding the drying time of the droplet – not incomprehensible" so the court of appeal's decision could not be overturned on this issue.

This can be formally reconstructed with the following defeasible rules:

r_1: $E_1 \Rightarrow$ *blood drop had dried*
r_2: $E_2 \Rightarrow \neg$ *blood drop had dried*
r_3: *blood drop had dried* $\Rightarrow \neg$ *certain that crime-related*
r_4: $\neg$ *blood drop had dried* $\Rightarrow \neg$ *certain that crime-related*
r_5: $\neg$ *certain that crime-related* $\Rightarrow \neg$ *doubt on reliability of confessions*

Given the ordinary premises E_1 and E_2, two arguments can be constructed for the conclusion $\neg$ *doubt on reliability of confessions*. Argument A applies r_1, r_3, r_5 and argument B applies r_2, r_4, r_5. A is rebutted on its subargument using r_1 by $B's$ subargument applying r_2 while B is likewise rebutted on its subargument using r_2 by $A's$ subargument applying r_1. If both rebutting arguments are regarded as equally strong, then all arguments in the example except the premise arguments E_1 and E_2 are defensible in preferred semantics: one preferred extension contains A and all its subarguments while the other extension instead contains B and all its subarguments. Yet both extensions contain an argument for the conclusion $\neg$ *doubt on reliability of confessions*, which therefore is justified.

Definition 6 correctly identifies all reasons for $\neg$ *doubt on reliability of confessions*. Indeed, we can build a justification containing the two possible weak explanations for this conclusion, one containing A and the argument for *blood drop had dried*, and the other containing B and the argument for $\neg$ *blood drop had dried*. For each conclusion of an argument or subargument occurring in these explanations, expanding the theory with its negation yields a justified attacker of a member of at least one of the explanations. Accordingly, the reasons for $\neg$ *doubt on reliability of confessions* are E_1, E_2, *blood drop had dried*, $\neg$ *blood drop had dried* and $\neg$ *certain that crime-related*.

6. Conclusion

A particular challenge for formally characterising reasons in legal argumentation arises when incompatible premises support the same conclusion, a form of reasoning recognised in legal practice when courts accept a common legal consequence without resolving which of the competing alternatives holds.

In this paper, starting from our previous *ASPIC*$^+$-based account of cause-in-fact [4], we generalised the underlying notion of causation to provide an account of reasons that accommodates such cases. By introducing notions of justification for floating conclusions, our account can identify alternative reasons for a conclusion even when no single argument for that conclusion is accepted in every preferred extension, as we illustrated through two legal examples.

Future work will investigate the application of the proposed account to a broader range of legal cases and more complex patterns of legal reasoning. Moreover, we intend to extend the existing implementation in [17] to compute the notions of justification and legal reason introduced in this paper. In the worst case, computing justifications and reasons would require exponential time with known algorithms since justification inherits the complexity of sceptical acceptance under preferred semantics, and identifying reasons may require exponentially many axiom-expansions. We expect, however, that legal argumentation frameworks of the size considered here remain tractable in practice.

Acknowledgements

Funded by the Italian Ministry of University and Research (MUR) under the Italian Fund for Science (FIS 3) – Call FIS 3 (DD No. 1802 of 21/11/2024), project 'GenAI4Law – Generative AI for the Law' (Project Code FIS-2024-07323) – CUP J53C25002490001 – Principal Investigator Prof. Giovanni Sartor.

References

[1] Raz J. Reasons for Action, Decisions and Norms. In: Raz J, editor. Practical Reasoning. Oxford University Press; 1978. p. 128-43.

[2] Horty J. Reasons as Defaults. Oxford University Press; 2012.

[3] Star D. The Oxford Handbook of Reasons and Normativity. Oxford University Press; 2018.

[4] Pisano G, Prakken H, Sartor G, Liepina R. Modelling cause-in-fact in legal cases through defeasible argumentation. In: Proceedings of the 20th International Conference on Artificial Intelligence and Law. New York: ACM Press; 2025. p. 268-77.

[5] Makinson D, Schlechta K. Floating conclusions and zombie paths: two deep difficulties in the "directly sceptical" approach to defeasible inheritance nets. Artificial Intelligence. 1991;48:199-209.

[6] Horty J. Skepticism and Floating Conclusions. Artificial Intelligence. 2002;135:55-72.

[7] Prakken H. Intuitions and the modelling of defeasible reasoning: some case studies. In: Proceedings of the Ninth International Workshop on Nonmonotonic Reasoning. Toulouse, France; 2002. p. 91-9.

[8] Schuster D, Broersen J, Prakken H. On floating conclusions. In: Proceedings of the 16th International Conference on Deontic Logic and Normative Systems. London: College Publications; 2023. p. 199-215.

[9] Sunstein CR. Legal Reasoning and Political Conflict. Oxford University Press; 1996.

[10] Fan X, Toni F. On computing explanations in argumentation. In: Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI 2015); 2015. p. 1496-502.

[11] Borg A, Bex FJ. Minimality, necessity and sufficiency for argumentation and explanation. International Journal of Approximate Reasoning. 2024;168:109143.

[12] Čyras K, Rago A, Albini E, Baroni P, Toni F. Argumentative XAI: A Survey. In: Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI-21); 2021. p. 4392-9.

[13] Dung PM. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming, and n–person games. Artificial Intelligence. 1995;77:321-57.

[14] Modgil S, Prakken H. The ASPIC+ framework for structured argumentation: a tutorial. Argument and Computation. 2014;5:31-62.

[15] Caminada M. Strong admissibility revisited. In: Parsons S, Oren N, Reed C, Cerutti F, editors. Computational Models of Argument. Proceedings of COMMA 2014. Amsterdam etc: IOS Press; 2014. p. 197-208.

[16] Wright RW. The NESS account of natural causation: a response to criticisms. In: Goldberg R, editor. Perspectives on Causation. Hart Publishing; 2011. p. 285-322.

[17] Pisano G, Prakken H, Sartor G. From cause-in-fact to legally relevant cause through argumentation. In: Proceedings of the 21st International Conference on Artificial Intelligence and Law. New York: ACM Press; 2026. p. forthcoming.