

ARCHIVIO ISTITUZIONALE DELLA RICERCA

Alma Mater Studiorum Università di Bologna Archivio istituzionale della ricerca

The burden of persuasion in structured argumentation

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Calegari R., Riveret R., Sartor G. (2021). The burden of persuasion in structured argumentation. New York : Association for Computing Machinery, Inc [10.1145/3462757.3466078].

Availability: This version is available at: https://hdl.handle.net/11585/836289 since: 2021-10-26

Published:

DOI: http://doi.org/10.1145/3462757.3466078

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (https://cris.unibo.it/). When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

Roberta Calegari, Régis Riveret, and Giovanni Sartor. 2021. The burden of persuasion in structured argumentation. In Proceedings of the Eighteenth International Conference on Artificial Intelligence and Law (ICAIL '21). Association for Computing Machinery, New York, NY, USA, 180–184.

The final published version is available online at: https://doi.org/10.1145/3462757.3466078

Rights / License:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<u>https://cris.unibo.it/</u>)

When citing, please refer to the published version.

The Burden of Persuasion in Structured Argumentation

Roberta Calegari Alma Mater Research Institute for Human-Centered Artificial Intelligence Bologna, Italy roberta.calegari@unibo.it

Régis Riveret Commonwealth Scientific and Industrial Research Organisation Brisbane, Australia regis.riveret@data61.csiro.au

Giovanni Sartor Alma Mater Research Institute for Human-Centered Artificial Intelligence Bologna, Italy giovanni.sartor@unibo.it

ABSTRACT

In this paper we provide an account of the burden of persuasion in the context of structured argumentation. A formal model for the burden of persuasion is defined, discussed, and used to capture the role of the burden of persuasion in adjudicating conflicts between conflicting arguments and in determining the dialectical status of arguments. We consider how our model can also capture adversarial burdens of proof, namely, those cases in which failure to establish an argument for a proposition burdened with persuasion entails establishing the complementary proposition.

KEYWORDS

Burden of persuasion, argumentation, legal reasoning

ACM Reference Format:

Roberta Calegari, Régis Riveret, and Giovanni Sartor. 2021. The Burden of Persuasion in Structured Argumentation. In Eighteenth International Conference for Artificial Intelligence and Law (ICAIL'21), June 21-25, 2021, São Paulo, Brazil. ACM, New York, NY, USA, 5 pages. https://doi.org/10. 1145/3462757.3466078

1 INTRODUCTION

The burden of proof is a central feature of many dialectical contexts. It is particularly relevant in those domains, such as legal disputations or political debates, in which controversial issues are discussed in order to adopt a decision, see [20] on burdens of proof in different dialogue types. Research in AI & law has devoted a number of contributions to the formal analysis of burdens of proof: models of defeasible legal reasoning have been criticised for not taking burdens of proof into account [11], the distinction between different standards of proof has been addressed [7], formal accounts of burdens of proof have been developed within models for formal argumentation [9, 17]. However, it seems to us that a comprehensive model of burdens of proof in legal reasoning is still missing.

In the legal domain, two types of burdens can be distinguished: the burden of production (also called burden of providing evidence, or 'evidential' burden), and the burden of persuasion [17]. This terminology is used in common law systems [21], but the distinction

is also recognised in civil law jurisdiction, possibly using a different terminology [10]. The focus of this paper is on the burden of persuasion. We will show how an allocation of the burden of persuasion may induce single outcomes in contexts in which the assessment of conflicting arguments would, without such an allocation, remain undecided. Our model combines Prakken and Sartor's [17] model with the insight from Carneades' [8], and takes into account the fact that the persuasiveness of an argument, in a dialectical context, is determined not only by the internal strength of the argument, as determined by the strength of the inference rules used for building the argument (according, for instance, to the last link criterion), but also by the applicable counterarguments. Our model originates from legal considerations and is applied to legal examples [4, 5]. However, the issue of the burden of proof has a significance that goes beyond the legal domain involving other domains - such as public discourse, risk management, etc. [3] - in which evidence and arguments are needed, and corresponding responsibilities are allocated, according to types of dialogues and dialectical or organisational roles [19, 20]. The novelty of this contribution consists of a new definition of defeat relations involving arguments burdened with persuasion, and a corresponding definition of the criteria for labelling such arguments.

BURDENS OF PERSUASION 2

Let us illustrate how the burden of persuasion works through two examples, one from criminal law and one from civil law.

Burden of persuasion in criminal law. In criminal law, the burden of production is distributed between prosecution and defence, while the burden of persuasion (in most legal systems) is always on prosecution. More exactly, in criminal law, the burden of production falls on the prosecution relative to the two constitutive elements of crime, namely, the criminal act (actus reus) and the required mental state (mens rea), be it intention/recklessness or negligence, while it falls to the defendant relative to justifications or exculpatory defences (e.g., self-defence, state of necessity, etc.). In other words, if both actus reus and mens rea are established, but no exculpatory evidence is provided, the decision should be criminal conviction. On the other hand, the burden of persuasion falls on the prosecution for all determinants of criminal responsibility, including not only for the constitutive elements of a crime but also for the absence of justifications of exculpatory defences.

Example 2.1. Let us consider a case in which a woman, Hellen, has shot and killed an intruder in her home. The applicable law consists of (a) the rule according to which intentional killing constitutes murder, and (b) the exception according to which there is no murder if the victim was killed in self-defence. Assume that it has

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICAIL'21, June 21-25, 2021, São Paulo, Brazil

been established with certainty that Hellen shot the intruder and that she did so intentionally. However, it remains uncertain whether the intruder was threatening Hellen with a gun, as claimed by the defence, or had turned back and was running away on having been discovered, as claimed by prosecution. The burden of persuasion is on prosecution, who needs to provide a convincing argument for murder. Since in this case it remains uncertain whether there was self-defence, prosecution has failed to provide such an argument. Therefore, the legally correct solution is that there should be no conviction: Hellen needs to be acquitted. □

Burden of persuasion in civil law. In civil law, burdens of production and burdens of persuasion may be allocated in different ways. The general principle is that the plaintiff only has the burden of proof (both of production and persuasion) relatively to the operative facts that ground its claim, while the defendant has the burden of proof relative to those exceptions which may prevent the operative facts from delivering their usual outcomes, such as justifications with regard to torts, or incapability and vices of consent in contracts. However, derogations from this principle may be established by the law, in order to take into account various factors, such as the presumed ability of each party to provide evidence in favour of his or her claim, the need to protect weaker parties against abuses, etc. In matters of civil liability, for example, it is usually the case that the plaintiff, who asks for compensation, has to prove both that the defendant caused the harm, and that this was done intentionally or negligently. However, in certain cases, the law establishes an inversion of the burden of proof for negligence. This means that in order to obtain compensation, the plaintiff only has to prove that s/he was harmed by the defendant. This will be sufficient to win the case unless the defendant provides a convincing argument that s/he was diligent (not negligent).

Example 2.2. Let us consider a case in which a doctor caused harm to a patient by misdiagnosing his case. Assume that there is no doubt that the doctor harmed the patient: she failed to diagnose cancer, which consequently spread and became incurable. However, it is uncertain whether or not the doctor followed the guidelines governing this case: it is unclear whether she prescribed all the tests that were required by the guidelines, or whether she failed to prescribe some tests that would have enabled cancer to be detected. Assume that, under the applicable law, doctors are liable for any harm suffered by their patients, but they can avoid liability if they show that they were diligent (not negligent) in treating the patient, i.e., that they exercised due care. Thus, rather than the patient having the burden of proving that doctors have been negligent (as it should be the case according to the general principles), doctors have the burden of providing their diligence. Let us assume that the law also says that doctors are considered to be diligent if they followed the medical guidelines that govern the case. In this case, given that the doctor has the burden of persuasion on her diligence, and that she failed to provide a convincing argument for it, the legally correct solution is that she should compensate the patient.□

These two examples share a common feature. In both, uncertainty remains concerning a decisive issue. However, this uncertainty does not preclude the law from prescribing a single legal outcome in each case. This outcome can be achieved by discarding the arguments that fail to meet the required burden of persuasion, i.e., the prosecution's argument for murder and the doctor's argument for her diligence, respectively.

3 ARGUMENTATION FRAMEWORK

We introduce a structured argumentation framework relying on a lightweight ASPIC⁺-like argumentation system [14]. For the sake of simplicity, we assume that arguments only consist of defeasible rules, to the exclusion of strict rules and of some constituents of a knowledge base—such as axioms, ordinary premises, assumptions, and issues that can be found in the complete model [14]. A framework based on defeasible rules is sufficient for our purposes and can be extended as needed with further structures.

3.1 Argumentation graphs

Let any literal be an atomic proposition or its negation. Literals are brought into relation through defeasible rules.

NOTATION 3.1. For any literal ϕ , its complement is denoted by $\overline{\phi}$, *i.e.*, if ϕ is a proposition p, then $\overline{\phi}$ is $\neg p$, while if ϕ is $\neg p$, then $\overline{\phi}$ is p.

Definition 3.1. A **defeasible rule** *r* is a construct of the form:

$$\rho: \quad \phi_1, ..., \phi_n, \sim \phi'_1, ..., \sim \phi'_m \Rightarrow \psi$$

with $0 \le n$ and $0 \le m$, and where

- ρ is the unique identifier for *r*, denoted by N(*r*);
- each $\phi_1, \ldots, \phi_n, \phi'_1, \ldots, \phi'_m, \psi$ is a literal;
- $\phi_1, \ldots, \phi_n, \sim \phi'_1, \ldots, \sim \phi'_m$ are denoted by Antecedent(r) and ψ by Consequent(r);
- ~ φ denotes the weak negation (negation by failure) of φ: φ is an exception that would block the application of the rule whose antecedent includes ~ φ.

The identifier of a rule can be understood as the name of the rule. It can be used as a literal to specify that the named rule is applicable, and its negation correspondingly to specify that the rule is inapplicable [13].

A superiority relation \succ is defined over rules: $s \succ r$ states that rule *s* prevails over rule *r*.

Definition 3.2. A superiority relation \succ over a set of rules *Rules* is an antireflexive and antisymmetric binary relation over *Rules*, i.e., $\succ \subseteq Rules \times Rules$.

A defeasible theory consists of a set of rules and a superiority relation over the rules.

Definition 3.3. A **defeasible theory** is a tuple $\langle Rules, \rangle$ where *Rules* is a set of rules, and \rangle is a superiority relation over *Rules*.

We can construct arguments by chaining rules from the defeasible theory, as specified in the following definition; cf. [6, 13, 18].

Definition 3.4. An **argument** *A* constructed from a defeasible theory $\langle Rules, \rangle$ is a finite construct of the form:

$$A: A_1, \ldots A_n \Rightarrow_r \phi$$

with $0 \le n$, and where

- *A* is the argument's unique identifier;
- A₁,..., A_n are arguments constructed from the defeasible theory (*Rules*, >);
- ϕ is the *conclusion* of the argument, denoted by Conc(*A*);

The Burden of Persuasion in Structured Argumentation

• r : Conc $(A_1), \ldots,$ Conc $(A_n) \Rightarrow \phi$ is the top rule of A, denoted by TopRule(A).

NOTATION 3.2. Given an argument $A : A_1, \ldots A_n \Rightarrow_r \phi$ as in definition 3.4, Sub(A) denotes the set of subarguments of A, i.e., Sub(A) = Sub(A_1) \cup \ldots \cup Sub(A_n) \cup \{A\}.

Different types of inconsistencies can appear between arguments, causing them to attack each other. In the ASPIC family of argumentation frameworks, attack is differentiated from defeat, with the latter taking preferences between arguments into account. Preferences over arguments are defined in the work reported here via a last-link ordering: an argument A is preferred over another argument B if the top rule of A is stronger than the top rule of B.

Definition 3.5. A **preference relation** > is a binary relation over a set of arguments \mathcal{A} , such that an argument *A* is preferred to argument *B*, denoted by A > B, iff TopRule(A) > TopRule(B).

Before specifying the notion of defeat between arguments, let us first identify burdens of persuasion, i.e., those literals the proof of which requires a convincing argument. We assume that such literals are consistent: it cannot be the case that there is a burden of persuasion both on ϕ and $\overline{\phi}$.

Definition 3.6 (Burdens of persuasion). Let BurdPers, the set of **burdens of persuasion**, be a set of literals such that if $\phi \in$ BurdPers then $\overline{\phi} \notin$ BurdPers. We say that an argument A is burdened with persuasion if Conc(A) \in BurdPers.

We now consider possible collisions between arguments, i.e., those cases in which an argument *A* challenges an argument *B*: (*a*) by contradicting the conclusion of a *B*' subargument (rebutting), or (*b*) by denying (the application of) the top rule of a *B*' subargument or by contradicting a weak negation in the body of the top rule of a *B*' subargument (undercutting). Note that our notion of rebutting corresponds to the notion of successful rebutting in [14].

Definition 3.7 (bp-rebut). Argument *A* **bp-rebuts** argument *B* iff $\exists B' \in Sub(B)$ such that $Conc(A) = \overline{Conc(B')}$ and

- (1) $Conc(A) \notin BurdPers, and B' \neq A, or$
- (2) $Conc(A) \in BurdPers and A > B'$.

According to Definition 3.7, for an unburdened argument A to rebut B by contradicting the latter's subargument B', it is sufficient that B' is non-superior to A. For a burdened argument A to rebut B by contradicting B', it is necessary that A is superior to B'. Thus, burdens of persuasion supplement priorities in deciding conflicts between arguments having opposed conclusions. They dictate the outcome of such conflicts when priorities do not already determine which argument is to prevail: when two arguments contradict one another, the one burdened with persuasion fails to bp-rebut the other, while the latter will succeed in bp-rebutting the first.

Undercutting is defined as usual, including both the case in which an the attacker excludes the application of the top rule of the attacked argument (by denying the rule's the name) and the case in which it contradicts a weakly negated literal in the body of that rule.

Definition 3.8 (bp-undercut). Argument A **undercuts** argument B iff $\exists B' \in Sub(B)$ such that: 1) Conc(A) = $\neg N(r)$ and TopRule(B') = r; or 2) Conc(A) = ϕ and $\sim \phi \in Antecedent(TopRule(B')).$

The notions of bp-rebuttings and undercuttings can then be used to define a defeat relation comprising bp-defeats and strict bp-defeats between arguments.

Definition 3.9 (bp-defeat). A **defeat** relation \rightarrow over a set of arguments \mathcal{A} is a binary relation over \mathcal{A} , i.e. $\rightarrow \subseteq \mathcal{A} \times \mathcal{A}$, such that $\forall A, B \in \mathcal{A}$, A defeats B, i.e. $A \rightsquigarrow B$, iff A bp-defeats B or A strictly-bp-defeats B:

- (1) A **bp-defeats** B iff A bp-rebuts B or A undercuts B
- (2) A **strictly-bp-defeats** B iff A bp-defeats B and B does not bp-defeats A.

Example 3.10 (Civil law example: rules and arguments). To exemplify the notions just introduced, let us formalise Example 2.2 through a set of rules. We assume that sufficient evidence is provided to support (in the absence of evidence to the contrary) the factual claims at issue (guidelines, \neg guidelines, harm), i.e., that the corresponding burdens of production are satisfied.

- $\begin{array}{lll} f1: & \Rightarrow \neg guidelines & r1: \neg guidelines \Rightarrow \neg dueDiligence \\ f2: & \Rightarrow guidelines & r2: guidelines \Rightarrow dueDiligence \\ f3: & \Rightarrow harm & r3: harm, \sim dueDiligence \Rightarrow liable \\ \end{array}$ We can then build the following arguments:
 - $\begin{array}{lll} A1: & \Rightarrow_{f1} \neg guidelines & A2: A1 \Rightarrow_{r1} \neg dueDiligence \\ B1: & \Rightarrow_{f2} guidelines & B2: B1 \Rightarrow_{r2} dueDiligence \\ C1: & \Rightarrow_{f3} harm & C2: C1 \Rightarrow_{r3} liable \end{array}$

If there were no burden of persuasion, the relations would be the following: arguments A1 and B1 defeat one another, B1 defeats A2, A1 defeats B2, A2 and B2 defeat one another, B2 strictly defeats C2. If on the contrary, there is burden of is on the doctors' diligence (*dueDiligence* \in BurdPers), then B2 fails to defeats A2, so that A2 strictly defeats B2.

Given a defeasible theory, arguments built from it and defeats between these arguments are gathered into an argumentation graph.

Definition 3.11. An **argumentation graph** constructed from a defeasible theory *T* is a tuple $\langle \mathcal{A}, \rightsquigarrow \rangle$, where \mathcal{A} is the set of all arguments constructed from *T*, and \rightsquigarrow is a defeat relation over \mathcal{A} .

NOTATION 3.3. Given an argumentation graph $G = \langle \mathcal{A}, \rightarrow \rangle$, we write \mathcal{A}_G , and \rightarrow_G to denote \mathcal{A} and \rightarrow respectively.

3.2 Labelling semantics

Let us now introduce the notion of $\{IN, OUT, UND\}$ -labellings of an argumentation graph, so that each argument in the graph is labelled IN, OUT or UND, depending on whether it is accepted, rejected, or undecided, respectively.

Definition 3.12. A {IN, OUT, UND}-labelling L of an argumentation graph G is a total function $L : \mathcal{A}_G \to \{IN, OUT, UND\}$.

NOTATION 3.4. Given a labelling L, we write IN(L) for $\{A \mid L(A) = IN\}$, OUT(L) for $\{A \mid L(A) = OUT\}$ and UND(L) for $\{A \mid L(A) = UND\}$.

There are various ways to specify {IN, OUT, UND}-labelling functions [1]. For example, they can be *complete* or *grounded*.

Definition 3.13. A complete {IN, OUT, UND}-labelling of an argumentation graph G is a {IN, OUT, UND}-labelling such that $\forall A \in \mathcal{A}_G$

(1) A is labelled IN iff all defeaters of A are labelled OUT, and

(2) A is labelled OUT iff A has a defeater labelled IN.

Definition 3.14. A grounded {IN, OUT, UND}-labelling of an argumentation graph *G* is a complete {IN, OUT, UND}-labelling *L* of *G* such that IN(L) is minimal.

Remark that any argument not labelled IN or OUT must be labelled UND, since any {IN, OUT, UND}-labelling is a total function.

While common specifications of {IN, OUT, UND}-labellings define reasonable positions [1], they do not cater for burdens of persuasion. We now specify the notion of bp-labelling, namely, a labelling which takes into account a set of burdens of persuasion.

Definition 3.15. A **bp-labelling** of an argumentation graph G, relative to a set of burdens of persuasion BurdPers, is a {IN, OUT, UND}-labelling L such that $\forall A \in \mathcal{A}_G$

(1) $A \in IN(L)$ iff $\forall B \in \mathcal{A}_G$ such that B bp-defeats $A : B \in OUT(L)$ (2) $A \in OUT(L)$ iff

- (a) $Conc(A) \in BurdPers and \exists B \in \mathcal{A}_G$ such that
- *B* bp-defeats *A*, and $B \in IN(L)$ or $B \in UND(L)$
- (b) $Conc(A) \notin BurdPers and \exists B \in \mathcal{A}_G$ such that • B bp-defeats A and $B \in IN(L)$

Burdens of persuasion affect conditions for rejection, as specified in condition 3.15 (2) (a): the rejection (the OUT labelling) of an argument burdened of persuasion may be determined by any defeating counterargument *B* that is accepted (IN) or also is uncertain (UND). However, as specified in condition 3.15 (2) (b), the rejection of an argument which is not burdened with persuasion requires a defeating counterargument *B* that is labelled IN.

The semantics just described does not always deliver a unique labelling. Multiple labellings may exist when arguments rebut each other, none of them being burdened with persuasion. If one of these arguments is labelled IN the other is labelled OUT, and vice versa. To address such a situation, we focus on IN-minimal labelling semantics, where for example both such arguments are labelled UND. Let us call such a labelling a *grounded bp-labelling*.

Definition 3.16. A bp-labelling L of an argumentation graph G is a **grounded bp-labelling** iff UND(L) is maximal.

PROPOSITION 3.17. Let L_1 be the grounded {IN, OUT, UND}-labelling of an argumentation graph G, and L_2 the grounded bp-labelling of G. If BurdPers = \emptyset then IN(L_1) = IN(L_2).

PROOF. It is easy to see that if condition 3.15(1) concerning arguments burdened with persuasion is removed from definition 3.15, we obtain the definition of grounded {IN, OUT, UND}-labellings.

Example 3.18 (Civil law example: graphs and bp-labelling). Let us consider again Example 2.2 and the corresponding rules and arguments built in Example 3.10. The argumentation graph and its grounded {IN, OUT, UND}-labelling are depicted in Figure 1 (top), in which all arguments are UND except arguments for undisputed facts. The result is not satisfactory according to the law, since it does not take into account the applicable burdens of persuasion. The doctor should have lost the case – i.e., be found liable – since she failed to discharge her burden of proving that she was diligent (non-negligent). The doctor's failure results from the fact that it remains uncertain whether she followed the guide-lines. To capture this aspect, we need to specify the burdens of

persuasion. Let us assume that (as under Italian law) we have BurdPers = {dueDiligence, liable}, i.e., the doctor has to provide a convincing argument that she was diligent, and the patient has to provide a convincing argument for the doctor's liability. As the burdened doctor's argument for dueDiligence is labelled OUT, her liability can be established even though it remains uncertain whether the guidelines were followed.



Figure 1: Grounded {IN, OUT, UND}-labelling of Example 2.2 in the absence of burdens of persuasion (top), and its bp-labelling with BurdPers = {dueDiligence, liable} (bottom).

This example shows how the model presented here allows us to deal with the *inversion of the burden of proof*, i.e., a situation in which one argument *A* is presented for a claim ϕ being burdened with persuasion, and *A* (or a subargument of it) is attacked by a counterargument *B*, of which the conclusion ψ is also burdened with persuasion. If no convincing argument for ψ can be found, then the attack fails, and the uncertainty on ψ does not affect the status of *A*. In the example, the argument for the doctor's due diligence fails to meet its burden of persuasion. Consequently, it fails to defeat the argument for the doctor's liability, which succeeds, meeting its burden of persuasion.

Example 3.19 (Criminal law example: rules, graphs and bp-labelling). Referring to Example 2.1, let us consider the following rules (for simplicity's sake, we do not specify pieces of evidence here, but we assume that all factual claims are supported by evidence):

f1: \Rightarrow killed	f2: \Rightarrow intention
f3: \Rightarrow threated	f4: $\Rightarrow \neg$ threated
r1: threated \Rightarrow selfDefence	r2: killed, intention \Rightarrow murder
r3: $selfDefence \Rightarrow \neg murder$	r4: \neg threated $\Rightarrow \neg$ selfDefence
with r3 > r2. We can build the following arguments:	
A	

A1: $\Rightarrow_{f1} killed$	B1 : \Rightarrow_{f3} threated
A2: \Rightarrow_{f2} intention	$B2: B1 \Rightarrow_{r1} selfDefence$
$A3: A1, A2 \Rightarrow_{r2} murder$	$B3:B2 \Longrightarrow_{r3} \neg murder$
C1 : $\Rightarrow_{f4} \neg threated$	
$C2: C1 \Rightarrow_{r4} \neg selfDefence$	

In the grounded $\{IN, OUT, UND\}$ -labelling of Figure 2 (top), all arguments are UND except for the undisputed facts. Thus, in the absence

The Burden of Persuasion in Structured Argumentation

of burdens of persuasion, we do not obtain the legally correct answer, namely, acquittal. To obtain acquittal we need to introduce burdens of persuasion. Prosecution has the burden of persuasion on *murder*: it therefore falls to the prosecution to persuade the judge that there was killing, that it was intentional, and that the killer did not act in self-defence.



Figure 2: Grounded {IN, OUT, UND}-labelling of Example 2.1 in the absence of burdens of persuasion (top), and bp-labelling with the burden of persuasion BurdPers = $\{murder\}$ (bottom).

The bp-labelling is depicted in Figure 2 (bottom). The prosecution failed to meet its burden of proving murder, i.e., its argument is not convincing, since it remains undetermined whether there was self-defence. Therefore, the argument supporting murder is labelled OUT, and the presumed killer is to be acquitted.

4 CONCLUSION

We have presented a formal model for the burden of persuasion. The model is based on the idea that arguments burdened with persuasion have to be rejected when there is uncertainty about them. We have shown how an allocation of the burden of persuasion may lead to a single outcome (IN arguments) in contexts in which the assessment of conflicting arguments would otherwise remain undecided. We have also shown how our model is able to address inversions of burdens of proof, namely, those cases in which the burden shifts from one party to the other. In such cases, there is the burden of persuasion over the conclusion of a multistep argument, and at the same time a burden of persuasion over the conclusion of an attacker against a subargument of that multistep argument. The model can be expanded in various ways, to capture further aspects of legal reasoning. For instance, it can also be supplemented with argumentation over burdens of persuasion [15], in a manner similar to the way in which argumentation systems can be expanded to include argumentation about priorities, see [12, 16]. More generally

we plan to study the properties of our semantics, and its connection with other semantics for argumentation [1, 2].

ACKNOWLEDGMENTS

R. Calegari and G. Sartor have been supported by the H2020 ERC Project "CompuLaw" (G.A. 833647).

REFERENCES

- Pietro Baroni, Martin Caminada, and Massimiliano Giacomin. 2011. An introduction to argumentation semantics. *The knowledge engineering review* 26, 4 (2011), 365–410. https://doi.org/10.1017/S0269888911000166
- [2] Pietro Baroni and Régis Riveret. 2019. Enhancing Statement Evaluation in Argumentation via Multi-labelling Systems. *Journal Artificial Intelligence Research* 66 (2019), 793–860. https://doi.org/10.1613/jair.1.11428
- [3] Roberta Calegari, Andrea Omicini, and Giovanni Sartor. 2020. Argumentation and Logic Programming for Explainable and Ethical AI. In XALit 2020 – Italian Workshop on Explainable Artificial Intelligence 2020 (CEUR Workshop Proceedings, Vol. 2742). Sun SITE Central Europe, RWTH Aachen University, Italy, 55–68.
- [4] Roberta Calegari and Giovanni Sartor. 2020. Burden of Persuasion in Argumentation. In Proceedings 36th International Conference on Logic Programming (Technical Communications), ICLP 2020 (Electronic Proceedings in Theoretical Computer Science, Vol. 325). OPA, Rende (CS), Italy, 151–163. https: //doi.org/10.4204/EPTCS.325.21
- [5] Roberta Calegari and Giovanni Sartor. 2020. A Model for the Burden of Persuasion in Argumentation. In Legal Knowledge and Information Systems. JURIX 2020: The Thirty-third Annual Conference (Frontiers in Artificial Intelligence and Applications, Vol. 334), Serena Villata, Jakub Harašta, and Petr Křemen (Eds.). IOS, Brno, Czech Republic, 13–22. https://doi.org/10.3233/FAIA200845
- [6] Martin Caminada and Leila Amgoud. 2007. On the Evaluation of Argumentation Formalisms. Artificial Intelligence 171, 5–6 (2007), 286–310. https://doi.org/10. 1016/j.artint.2007.02.003
- [7] Arthur M. Farley and Kathleen Freeman. 1995. Burden of Proof in Legal Argumentation. In Proceedings of the 5th International Conference on Artificial Intelligence and Law. ACM, Maryland USA, 156–164. https://doi.org/10.1145/222092.222227
- [8] Thomas F. Gordon, Henry Prakken, and Douglas Walton. 2007. The Carneades model of argument and burden of proof. Artificial Intelligence 171, 10 (2007), 875-896. https://doi.org/10.1016/j.artint.2007.04.010
- [9] Thomas F. Gordon and Douglas N. Walton. 2009. Proof Burdens and Standards. In Argumentation in Artificial Intelligence. Springer, Boston, MA, 239–258. https://doi.org/10.1007/978-0-387-98197-0_12
- [10] Ulrike Hahn and Mike Oaksford. 2007. The Burden of Proof and Its Role in Argumentation. Argumentation 21 (2007), 36–61. https://doi.org/10.1007/s10503-007-9022-6
- [11] Ronald E. Leenes. 2001. Burden of Proof in Dialogue Games and Dutch Civil Procedure. In Proceedings of the 8th International Conference on Artificial Intelligence and Law. ACM, Missouri USA, 109–18. https://doi.org/10.1145/383535.383549
- [12] Sanjay Modgil and Henry Prakken. 2010. Reasoning about Preferences in Structured Extended Argumentation Frameworks. In Proceedings of COMMA 2010, Computational Models of Argumentation. IOS, Italy, 347–58. https://doi.org/10. 3233/978-1-60750-619-5-347
- [13] Sanjay Modgil and Henry Prakken. 2014. The ASPIC⁺ framework for structured argumentation: a tutorial. Argument & Computation 5, 1 (2014), 31–62. https: //doi.org/10.1080/19462166.2013.869766
- [14] Henry Prakken. 2010. An Abstract Framework for Argumentation with Structured Arguments. Argument and Computation 1 (2010), 93–124. https://doi.org/ 10.1080/19462160903564592
- [15] Henry Prakken, Chris Reed, and Douglas N. Walton. 2005. Dialogues about the Burden of Proof. In Proceedings of the 10th International Conference on Artificial Intelligence and Law. ACM, Bologna, Italy, 115–124. https://doi.org/10.1145/ 1165485.1165503
- [16] Henry Prakken and Giovanni Sartor. 1996. Rules about Rules: Assessing Conflicting Arguments in Legal Reasoning. Artificial Intelligence and Law 4 (1996), 331–68. https://doi.org/10.1007/BF00118496
- [17] Henry Prakken and Giovanni Sartor. 2010. A Logical Analysis of Burdens of Proof. Legal Evidence and Proof: Statistics, Stories, Logic 1 (2010), 223–253.
- [18] Gerard Vreeswijk. 1997. Abstract Argumentation Systems. Artificial Intelligence 90, 1-2 (1997), 225-279. https://doi.org/10.1016/S0004-3702(96)00041-0
- [19] Douglas Walton. 1996. Arguments from Ignorance. Pennsylvania State University Press, Pennsylvania. https://doi.org/10.1007/978-3-319-15013-0_3
- [20] Douglas Walton. 2014. Burden of proof, presumption and argumentation. Cambridge University Press, USA. https://doi.org/10.1017/CBO9781107110311
- [21] C.R. Williams. 2003. Burdens and standards in civil litigation. Sydney Law Review 25 (2003), 165–188.