Proceedings of the 23rd International Conference on Principles of Knowledge Representation and Reasoning
Main Track

Tenability and Weak Semantics: Modeling Non-uniform Defense

Uri Andrews [1] **, Luca San Mauro** [2] **, John Spoerl** [1]

[1]Department of Mathematics, University of Wisconsin–Madison
[2]Dipartimento di Ricerca e Innovazione Umanistica, University of Bari, Italy
andrews@math.wisc.edu, luca.sanmauro@gmail.com, john.spoerl@wisc.edu

Abstract

In Dung-style abstract argumentation, various semantics capture notions of acceptability of arguments. The admissibility semantics capture the notion that an argument can be consistently defended from any potential counterargument. Weak semantics often relax the demands of admissibility by restricting which counterarguments must be taken seriously (e.g., discounting self-defeating or otherwise incoherent attacks). Many prominent proposals for weak semantics remain extension-based in a stronger sense. While these semantics discount attacks from arguments which are considered unreasonable, they still require a uniform defense against all reasonable arguments, even if they are collectively inconsistent. This uniformity can be too demanding when defensibility is inherently strategic, and thus the appropriate reply depends on the opponent's line of attack.

We introduce *tenability*, a family of dialogue-based semantics that formalize when a designated argument (or a set of arguments) can be maintained in debate by a proponent against any conflict-free attack which the opponent may present. The approach is motivated by three natural benchmark patterns: self-defeating attack, floating assignment, and disjunctive reinstatement, on which tenability behaves differently from all weak semantics previously considered in the literature. We define three variants—static tenability, tenability, and strong tenability—via monotone commitment games over finite conflict-free moves, differing in the obligations imposed on the disputants. We establish the relative strength of these notions, prove implications and separations with previously studied weak semantics, and we analyze computational complexity on finite frameworks: deciding static tenability is Π_2^P-complete, while deciding tenability and strong tenability is $PSPACE$-complete.

1 Introduction

Dung's abstract argumentation frameworks (AFs) (Dung 1995) have become a core formalism for defeasible reasoning in knowledge representation (Bench-Capon and Dunne 2007; Rahwan and Simari 2009). An AF represents arguments as abstract entities related by an attack relation; a semantics then determines which sets of arguments (extensions) can be collectively maintained. A foundational role in this landscape is played by *admissibility*-based semantics: a set should be internally consistent (conflict-free) and able to withstand counterarguments by defending each of its members against *every* attacker. This uniform defense ideal underlies the standard Dung-style semantics (complete, grounded, preferred, stable) and their numerous refinements (Dung 1995; Baroni, Caminada, and Giacomin 2011).

At the same time, it is equally well known that strict admissibility may be overly conservative in the presence of pathological attack patterns, most notably those featuring self-defeating arguments and odd cycles. In such configurations, an attack may arguably fail to constitute a serious challenge precisely because its source undermines itself, yet admissibility treats it on par with ordinary attacks (Bodanza and Tohmé 2009; Baroni and Giacomin 2007). This and related concerns have led to a large and still expanding body of work proposing alternatives that weaken or reshape the admissibility requirement, yielding a landscape of semantics that is both broad and structurally diverse (Baroni and Giacomin 2007; Baroni and Giacomin 2009; Baroni, Caminada, and Giacomin 2011; Baroni et al. 2018; van der Torre and Vesic 2017). Classic examples include non-admissible or cycle-sensitive proposals (e.g. stage-style or semi-stable variants) (Verheij 1996; Caminada, Carnielli, and Dunne 2012) as well as resolution-based families designed to tame problematic cyclic structures while retaining key principles (Baroni, Dunne, and Giacomin 2011).

Within this broad spectrum, *weak semantics* aim to preserve the defense-based intuition while refining which attacks count or how they must be responded to. For orientation, we recall three influential lines that will serve as com-

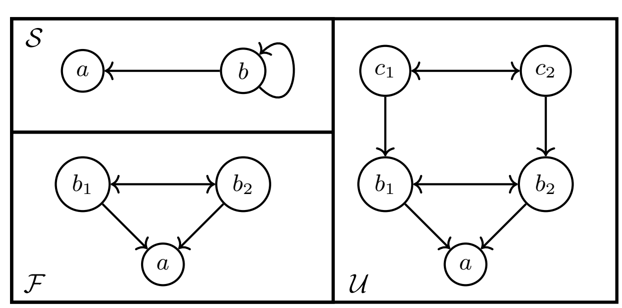

Figure 1: $\mathcal{S}$ - Self-Defeating Attack, $\mathcal{F}$ - Floating Assignment, $\mathcal{U}$ - Disjunctive Reinstatement

parison points throughout the paper.

- Cogency-based approaches start from a comparative idea: extensions are ordered by a binary *at least as cogent as* relation $X \succeq Y$, which captures that X defends itself against attacks from Y. In the standard cogency case, this admits an equivalent, more operational reading: selecting the undominated extensions amounts to retaining a Dung-style notion of defence while discounting self-defeating attacks. (Bodanza and Tohmé 2009; Bodanza, Tohmé, and Simari 2016; Blümel and Ulbricht 2022).
- Weak admissibility weakens defense *recursively* via reduct-style constructions that filter which attackers remain relevant once a candidate position is assumed (Baumann, Brewka, and Ulbricht 2020). This yields semantics that mitigate odd-cycle effects while staying close in spirit to defense-based acceptability, and has already prompted further extensions and complexity analyses (Dvořák, Ulbricht, and Woltran 2022; Dauphin, Rienstra, and van der Torre 2020).
- Weakly complete semantics revise Dung's complete labelings by relaxing the propagation of undecidedness: under suitable conditions, an argument may be accepted even if some attackers remain undecided, via the mechanisms of *undecidedness blocking* or *propagation* (Dondio and Longo 2021). This provides a principled way to model acceptance under unresolved counterevidence and connects naturally to ambiguity-blocking intuitions.

Despite their differences, the above approaches share a common structural commitment typical of extension and labeling-based semantics: acceptability is witnessed by a *single* stance that must be adequate against *all* attacks that the semantics deems relevant, *simultaneously*. This "uniform witness" pattern is also reflected in classical proof-theoretic and dialogical characterizations of argumentation semantics (Vreeswijk and Prakken 2000; Prakken 2005; Caminada and Wu 2009; McBurney and Parsons 2009). However, in many debates and evidential settings, defensibility is inherently *strategic*: the appropriate reply depends on which consistent line of attack an opponent commits to, and mutually incompatible challenges need not be met *jointly* by one fixed, consistent position. Requiring a single stance that neutralizes every consistent challenge at once can therefore be too demanding.

This paper develops *tenability*, a family of semantics designed to capture *non-uniform* defendability. The key conceptual shift is from *there exists one stance answering all relevant attacks* to *for every consistent attack, there exists an appropriate response*. We formalize this idea through monotone commitment protocols in which players build finite conflict-free positions by increasing moves, yielding three notions with increasing dialogical structure: *static tenability*, *strong tenability*, and *tenability*. These notions are motivated and contrasted on a small set of benchmark patterns in Section 2, formally defined in Section 3, and studied from a mathematical and computational perspective in Section 4. Section 5 positions tenability within the weak-semantics landscape and highlights the distinctive role of *non-uniform* defense, while Section 6 discusses limitations and directions for further work.

2 Motivation

Our starting point is a simple query: When is a designated argument a defensible in a dispute? In Dung-style abstract argumentation, defensibility is most often verified through a *uniform witness*: acceptance is certified by a single stance (an extension or a labeling) that must be adequate against all counterarguments deemed relevant by the semantics, simultaneously. This uniformity is also reflected in standard dialogical and game-theoretic characterisations of Dung-style semantics, where acceptance is tied to the existence of a (winning) defense that is robust across all opponent choices (Vreeswijk and Prakken 2000; Prakken 2005).

In this paper we isolate a complementary, strategy-oriented reading: *defensibility is the ability to answer any consistent line of attack*, possibly with different replies against incompatible challenges. This motivates *tenability* as a semantics of *non-uniform defense*.

To make the issue concrete, we focus on three small benchmark patterns shown in Figure 1: *self-defeating attack* $\mathcal{S}$, *floating assignment* $\mathcal{F}$, and *disjunctive reinstatement* $\mathcal{U}$. Across these benchmarks, we test whether a can be maintained against reasonable challenges. In an abstract AF, a minimal proxy for a reasonable line of challenge is legitimacy: an opponent should not be allowed to press, as a single line of contestation, a bundle of mutually incompatible arguments. Accordingly, we treat a legitimate attack as a conflict-free set of arguments advanced as one line of challenge.

Self-defeating attack ($\mathcal{S}$). In $\mathcal{S}$, the only attacker of a is self-defeating. Such patterns have long been recognised as problematic for classic admissibility and related semantics: if an attacker undermines itself (directly or via odd-cycle phenomena), it is unclear why it should count as a decisive challenge (Dung 1995; Bodanza and Tohmé 2009; Baumann, Brewka, and Ulbricht 2020; Dondio and Longo 2021). Under a consistency constraint, a self-defeating argument cannot constitute a reasonable attacking move. This aligns with the guiding intuition behind several weak-semantics proposals that discount self-defeating (or otherwise problematic) attacks when assessing defense (Bodanza and Tohmé 2009; Baumann, Brewka, and Ulbricht 2020; Dondio and Longo 2021).

Floating assignment ($\mathcal{F}$). The framework $\mathcal{F}$ exhibits a different phenomenon. Here a is attacked by two incompatible counterarguments b_1 and b_2 (they attack each other). While $\{b_1, b_2\}$ is not a legitimate line of attack, each singleton $\{b_1\}$ and $\{b_2\}$ *is* legitimate. Yet a has no satisfactory reply to either: any attempt to counterattack b_1 must employ b_2 as defender, and vice versa, but either choice immediately destroys consistency with a itself. So $\mathcal{F}$ captures the principle that a position is *untenable* whenever the opponent can mount a legitimate challenge that the proponent cannot answer even in isolation. This benchmark is widely discussed in connection with ambiguity and standards of proof in legal-style readings, where incompatible testimonies may or may not suffice for a verdict depending on the evidentiary threshold (see, e.g., (Horty 2001)). For tenability, the decisive point is structural: inability to answer a consistent singleton challenge already voids defensibility.

Disjunctive reinstatement ($\mathcal{U}$). The crucial benchmark for our purposes is $\mathcal{U}$. As in $\mathcal{F}$, the designated argument a faces two incompatible attackers b_1 and b_2. Unlike $\mathcal{F}$, each attacker can be met by independently consistent defenses: c_1 neutralises b_1 and c_2 neutralises b_2, and each of $\{a, c_1\}$ and $\{a, c_2\}$ is conflict-free. What fails is precisely uniform defense: there is no single conflict-free set extending $\{a\}$ that counterattacks *both* b_1 and b_2 at once, since c_1 and c_2 are themselves incompatible. This illustrates the gap between a uniform-witness reading of defense and a strategic reading. Under a strategic reading, once the opponent commits to a consistent line of challenge, the proponent may tailor her reply to that line. Thus a can be maintained even though it is not uniformly defended by any fixed conflict-free stance.

This gap also exposes an asymmetry that can arise when uniformity is imposed at the level of witnessing stances: the proponent is required to provide a single consistent "story" that addresses mutually incompatible accusations jointly, even when no rational opponent would be entitled to maintain those accusations together as one consistent line. We refer to this as the *Master versus Ignoramus* effect: the proponent is treated as a master who must answer every objection within one account, while the opponent enjoys the freedom of switching between incompatible objections without committing to their joint consistency.

To make the strategic reading more intuitive, consider the following toy narrative modeled by $\mathcal{U}$.

Miss Scarlett is accused of a crime that occurred in the ballroom. The following testimonies are available as evidence.

- a: *Miss Scarlett claims she was not at the scene at the time of the crime.*
- b_1: *Colonel Mustard claims he saw Miss Scarlett leaving the west exit of the library at the time of the crime.*
- b_2: *Mr. Green claims he saw Miss Scarlett leaving the east exit of the library at the time of the crime.*
- c_1: *Mrs. White testifies that Colonel Mustard and Professor Plum were with her in the lounge throughout the evening.*
- c_2: *Professor Plum testifies that Mr. Green and Mrs. White were with him in the study throughout the evening.*

If the opponent challenges a by committing to the line $\{b_1\}$, the proponent can respond with $\{c_1\}$; if the opponent instead commits to $\{b_2\}$, the proponent can respond with $\{c_2\}$. Crucially, in $\mathcal{U}$ the opponent is not entitled to press $\{b_1, b_2\}$ as a *single* consistent line of attack, since b_1 and b_2 are mutually incompatible.

Taken together, $\mathcal{S}$, $\mathcal{F}$, and $\mathcal{U}$ isolate three desiderata for a tenability-style notion of defensibility:

1. Inconsistent challenges should not by themselves undermine defensibility ($\mathcal{S}$).
2. Every consistent line of attack must be answerable on its own; if some consistent attack admits no consistent reply, the position is untenable ($\mathcal{F}$).
3. The defense may depend on the opponent's consistent line of challenge; incompatible attacks need not be rebutted jointly by a single consistent stance ($\mathcal{U}$).

These desiderata motivate our three tenability notions introduced next. *Static tenability* captures the one-shot strategic idea already visible in $\mathcal{U}$: for every consistent challenge B attacking a position A, there exists a conflict-free extension $C \supseteq A$ that counters every attack from B. *Strong tenability* and *tenability* lift this to dynamic disputes via monotone commitment protocols: the proponent must maintain and defend her accumulated commitments over time. Formal definitions are presented in Section 3.

3 Notions of Tenability

We now formalize *tenability* as a family of weak semantics grounded in the strategy-based reading developed in Section 2: defensibility is the ability to answer any consistent line of attack, possibly with different replies against incompatible challenges. Technically, we capture this via monotone commitment protocols over finite conflict-free moves.

We introduce three notions: *static tenability* (Section 3.1), *strong tenability* (Section 3.2), and *tenability* (Section 3.3). Tenability is our main target; the other two notions are introduced first because they isolate key components of non-uniform defense and help motivate the design choices in the final definition.

We fix some terminology regarding argumentation frameworks. An argumentation framework $\mathcal{F} = (F, \rightarrowtail)$ will always be identified with its set of arguments F, i.e., we write $a \in \mathcal{F}$ to mean $a \in F$.

Definition 3.1. Given an argumentation framework $\mathcal{F}$, argument $a \in \mathcal{F}$, and sets of arguments $A, B \subseteq \mathcal{F}$, we say:

- *a attacks* a set of arguments B ($a \rightarrowtail B$) if $(\exists b \in B)[a \rightarrowtail b]$;
- $A^+ = \{x \in \mathcal{F} \mid (\exists a \in A)[a \rightarrowtail x]\}$;
- *A is as cogent as B*, written $A \succeq B$, if A is conflict-free and every $b \in B$ which attacks A is counterattacked by some argument in A: $(\forall b \in B)[b \rightarrowtail A \implies b \in A^+]$;
- *B is more cogent than A* ($B \succ A$) if B is as cogent as A, but A is not as cogent as B.

3.1 Static Tenability

One possible way of understanding the debate over $a \in \mathcal{U}$ is to recognize that for every consistent attack on a, there is a consistent extension containing a which defends itself against said attack. In other words, a naïve defense strategy is one that has a reply for every set of incoming attacks. Thus, the minimal working definition to capture tenability is the following:

Definition 3.2. A set of arguments A is *statically tenable* if for every finite conflict-free B, there is a conflict-free $C \supseteq A$ which is as cogent as B.

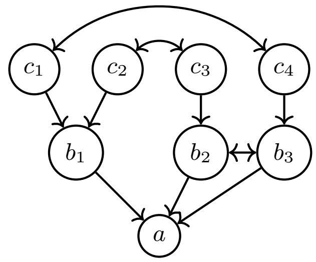

Figure 2: An argumentation framework $\mathcal{E}$ in which $\{a\}$ is statically tenable, but not strongly tenable.

As a preliminary test, we analyze the examples $\mathcal{S}$, $\mathcal{F}$, and $\mathcal{U}$ from Figure 1. Static tenability accepts $a \in \mathcal{S}$ because there is no conflict-free attack to defend against, and rejects $a \in \mathcal{F}$ because there is a conflict-free $B = \{b_1\}$ for which there is no defense. In the case of $\mathcal{U}$, there are 4 conflict-free sets which attack $\{a\}$. Namely, $\{b_1\}, \{b_1, c_2\}, \{b_2\}$, and $\{b_2, c_1\}$. The first two are handled by extending $\{a\}$ to $\{a, c_1\}$. The latter two are handled by extending $\{a\}$ to $\{a, c_2\}$. In any case, the proponent of $\{a\}$ has a strategy to counter any consistent attack on her beliefs.

This interpretation of tenability has some advantages and some drawbacks. The definition is concise and straightforward and represents a preliminary or heuristic check that an agent may apply to its beliefs. However, this definition does leave much to be desired. It merely represents the first exchange of a debate; it is a minimal requirement for an agent to defend its beliefs against incoming attacks. In a more realistic exchange, the attacker would continue its charge, bringing forth further arguments against the defending set. Consider for example the argumentation framework $\mathcal{E}$ (Figure 2). The set $\{a\}$ can be checked to be statically tenable. However, in an extended dispute, a has no adequate defense, hence should not be classified as tenable.

The extended dispute may go as follows: the opponent to a proposes b_1. In order to counterattack b_1, the proponent must extend its position by accepting c_1 or c_2. If she accepts c_1, then the opponent continues their attack with b_3. There is no way for the proponent to counter the attack b_3 without contradicting the positions she has already assumed during the debate. On the other hand, if she accepts c_2, then the opponent continues their attack with b_2. There is no consistent rebuttal to $\{b_1, b_2\}$ which extends $\{a, c_2\}$. The tables below are a representation of these debates as two player games between the proponent (*Pro*) and her opponent (*Opp*).

Pro	$\{a\}$	$\{a, c_1\}$
Opp	$\{b_1\}$	$\{b_1, b_3\}$

Pro	$\{a\}$	$\{a, c_2\}$
Opp	$\{b_1\}$	$\{b_1, b_2\}$

The important aspect we are highlighting is that the tenability of $\{a\}$ is conditional upon having a *winning strategy* in a certain dispute about $\{a\}$. Arguments and debates involve back-and-forth, so tenable arguments should be maintainable beyond the first round of attacks. We thus isolate precisely which disputes are relevant and give *dynamic* definitions for tenability.

3.2 Strong Tenability

Definition 3.3 (Strong Tenability)**.** Let $\mathcal{F}$ be an AF. A *strong tenability dispute* on $\mathcal{F}$ is a sequence of subsets $(X_0, X_1, ..., X_n)$ of $\mathcal{F}$, of which the even moves (X_{2i}) are Pro moves (for Proponent), and the odd moves (X_{2i+1}) are Opp moves (for Opponent), such that:

(1) each X_i is conflict-free;
(2) $X_i \subseteq X_{i+2}$ for all i;
(3) X_1 and each $(X_{i+2} \setminus X_i)$ are finite;[1]
(4) every Pro move, X_{2n}, is as cogent as the preceding Opp move X_{2n+1}, i.e., $X_{2n} \succeq X_{2n-1}$;
(5) every Opp move, X_{2n+1}, has an attacker of the preceding Pro move, X_{2n}, which is not counterattacked by X_{2n}, i.e., $X_{2n+1} \not\preceq X_{2n}$.

The *strong tenability game* is the alternating move game between Pro (or Proponent) and Opp (or Opponent) where the legal positions are the strong tenability disputes.

A strong tenability dispute $(X_0, X_1, ..., X_n)$ has *concluded* if there is no X_{n+1} such that $(X_0, X_1, ..., X_{n+1})$ is a strong tenability dispute. A concluded strong tenability dispute is *won* by the player making the last move X_n. A *strategy for* Pro is a function mapping each strong tenability dispute of even length to a legal move for Pro, i.e. mapping $(X_0, \ldots, X_{2n+1})$ to some set X_{2n+2} so that $(X_0, \ldots X_{2n+1}, X_{2n+2})$ is also a strong tenability dispute. We say that a strategy is a *winning strategy for* Pro if every concluded dispute where Pro's moves are determined by the strategy is won by Pro. A set $A \subseteq \mathcal{F}$ is *strongly tenable* if Pro has a winning strategy beginning with $X_0 = A$. Note that a winning strategy need not lead to a concluded dispute. Rather, if the strategy leads to an infinite sequence of moves, we consider this a victory for Pro, since she has successfully defended $X_0 = A$ against successive attacks.

In the case that Opp plays last in a concluded dispute, this means that Opp has played a set with an attacker of Pro's last move which is not consistently counterattacked by any collection of arguments I has played or can play. In the case that Pro plays last, Pro may only need to play a set of arguments 'as good as' any further counterarguments. In a debate this may look like the opponent realizing he has no strictly better counterarguments, so he resigns his assault upon A. Thus strong tenability semantics, following the lead of Dung, are also guided by the principle "the one who has the last word laughs best" (Dung 1995, p. 322). A further note is that backtracking is omitted as part of the dispute: A's strong tenability being defined in terms of a winning strategy means that if backtracking is necessary for a player, then that player loses the run of the dispute and a winning strategy would avoid such plays.

Compare the strong tenability dispute to the argument game for admissibility (see (Vreeswijk and Prakken 2000)) where there is more asymmetry in the roles of proponent and

[1]Note that the definition allows for X_0 to be infinite, while all subsequent attacks and defenses may only introduce finitely many new arguments. Formalizing strong tenability disputes and tenability disputes in this manner facilitates the future study of tenability semantics in infinite Argumentation Frameworks. See section 4.1.

opponent. The proponent must maintain consistency while the opponent need not ('Master versus Ignoramus'). It is exactly this asymmetry which is addressed through the family of tenability semantics. Indeed it underscores how admissibility requires a uniform response to challenges, whereas tenability does not. However, there is an asymmetry to our definition that persists: the opponent does not need to defend the arguments they put forth. Pro must defend all of the arguments it incorporates into the dispute, while *Opp* may leave loose ends. For this reason, we find it plausible (indeed likely) that a rational agent defending its position a will disregard the opponent's evidence if it is not also defended in the course of the dispute. Consider again the Strong Tenability dispute on $\mathcal{E}$, in the form of a dialogue:

> Pro: *I propose* a.
> *Opp: Well* b_1 *is true, so* a *is not.*
> Pro: *If* b_1 *is the position you hold, then I will defend* a *by adopting* c_1.
> *Opp: I now propose* b_3*, another reason* a *cannot hold.*
> Pro: *Now hold on, you haven't defended your own position* b_1. *I don't believe you are proposing your arguments in good faith. Thus I reject this threat upon* a.

The lesson we take from this is that a rational agent may not be forced to abandon her positions in the face of bad faith arguments. To be swayed from the original position, we should require the opponent to play by the same rules. If admissibility is captured by a game of 'Master vs Ignoramus', then the strong tenability dispute is one of *Master versus Novice* where the novice need not defend his positions, yet is expected to not take explicitly contradictory positions. We next define tenability in order to model a dispute of *Master versus Master*.

3.3 Tenability

In addition to the desiderata suggested by $\mathcal{F}, \mathcal{S}$, and $\mathcal{U}$, our investigation has yielded further refinements for an adequate notion of tenability:

(i) The proponent and opponent are allowed to progress the dispute by extending their current stance (accepting further arguments);

(ii) Both proponent and opponent must argue *in good faith*, that is, they must defend the positions they've proposed during the dispute;

(iii) An argument is tenable when the proponent has a (possibly non-uniform) rebuttal to all lines of attack from the opponent, i.e. has a winning strategy.

We now propose a refined definition of tenability incorporating these criteria.

Definition 3.4 (Tenability)**.** A *tenability dispute* on $\mathcal{F}$ is a sequence of subsets $(X_0, X_1, ..., X_n)$ of $\mathcal{F}$, of which the even moves (X_{2i}) are Pro moves, and the odd moves (X_{2i+1}) are *Opp* moves, such that:

(1) each X_i is conflict-free;
(2) $X_i \subseteq X_{i+2}$ for each i;
(3) X_1 and each $(X_{i+2} \setminus X_i)$ are finite;
(4) every move must be as cogent as the preceding move $(X_{i+1} \succeq X_i)$;
(5) every *Opp* move X_{2i+1} must be more cogent than the preceding Pro move X_{2i} $(X_{2i+1} \succ X_{2i})$.

The *tenability game* is the alternating move game between Pro (or Proponent) and *Opp* or (Opponent) where the legal positions are the tenability disputes. A tenability dispute $(X_0, X_1, ..., X_n)$ has *concluded* if there is no X_{n+1} such that $(X_0, X_1, ..., X_{n+1})$ is a tenability dispute. A concluded tenability dispute is won by the player making the last move X_n. We say that a strategy is a *winning strategy for* Pro if every concluded dispute where Pro's moves are determined by the strategy is won by Pro. A set $A \subseteq \mathcal{F}$ is *tenable* if Pro has a winning strategy beginning with $X_0 = A$.

The distinction between the tenability dispute and the strong tenability dispute is found only in the symmetry between Pro and *Opp* in the fifth requirement (5). One asymmetry still remains, in that *Opp*'s moves must contain an attacker of Pro's prior move which is not counterattacked, while Pro need only provide counterattacks. It is worth here justifying this asymmetry:

Because we are interested in modeling the plausibility of an argument, rather than its provability, we are seeking an inherently defensive notion. Hence Pro's plays need only provide an adequate defense of A, rather than proactively proving A. This same motivation also indicates why *Opp* must provide a argument which attacks and is not attacked by Pro's move: if *Opp* does not provide a strictly stronger play X_{2n+1}, then Pro's current play X_{2n} has already maintained itself against that charge. There is no good reason to adopt further beliefs to defend against X_{2n+1} as X_{2n} accounts for those doubts, hence X_{2n+1} is not a move which progresses the dispute. Furthermore, we have no requirement on the number of arguments that Pro and *Opp* may include in a dispute, besides that they may introduce only finitely many new arguments, which means that both Pro and *Opp* have the appropriate means at every stage of the dispute to both defend their prior positions and, in the case of *Opp*, put forward a new attack. The tenability dispute serves our intuitions about *Master vs. Master* debates.

4 Mathematical Analysis

The preceding discussion suggests that these notions of tenability are distinct, yet related. Indeed this is the case, as shown in the next proposition. Proofs are available in the extended version, to appear on arXiv.

Proposition 4.1. *Strong tenability implies tenability; tenability implies static tenability. Moreover, each implication is strict.*

The strictness of these implications are given by the framework in Figure 3, which is statically tenable, but not tenable, and the framework in Figure 2, which is tenable, but not strongly tenable.

One concern for tenability and other weak semantics is their compatibility with existing semantics. The following proposition shows that the tenability notions are indeed weakenings of admissibility.

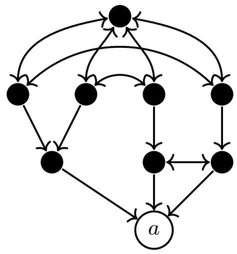

Figure 3: An argumentation framework in which $\{a\}$ is statically tenable, but not tenable.

Proposition 4.2. *If A is admissible, then A is (statically/-/strongly) tenable.*

We also show that (strong) tenability is compatible with the grounded extension. Further, in finite AFs, static tenability is also compatible with the grounded extension.

Theorem 4.3. *If A is (strongly/-) tenable and $\mathcal{G}$ is the grounded extension, then $A \cup \mathcal{G}$ is (strongly/-) tenable. If $\mathcal{F}$ is finite and A is statically tenable, then $A \cup \mathcal{G}$ is statically tenable.*

As an aside, there exists an infinite argumentation framework $\mathcal{F}$ and an argument a so that a is statically tenable, yet a is attacked by the grounded extension. See Theorem 4.8 and Figure 5.

The next proposition highlights a structural property that sharply distinguishes tenability from most familiar weak semantics: *downward closure* (or *conservativity*) with respect to set inclusion.

Proposition 4.4. *If A is (strongly/-/statically) tenable, and $X \subseteq A$, then X is (strongly/-/statically) tenable.*

At first sight this may look surprising, since many extension-based semantics are not closed under subsets: an extension is meant to be a self-standing position, and dropping arguments may destroy the internal support needed to defend the remaining ones. Tenability is deliberately different. Under our strategy-based reading, a set A is not certified because it is a single stance that "stands on its own"; rather, A is tenable because it can be *maintained* against consistent attacks by suitably extending it, possibly in different ways for different attacks. If A can be maintained, then any weaker commitment $X \subseteq A$ can be maintained as well: starting from X the proponent can simply follow the very same defense strategy, committing (when needed) to some of the additional arguments that were available from A. Intuitively, moving from A to X removes liabilities rather than resources: any consistent attacks on X were already attacks on A, and it cannot make defense harder, since the proponent is always allowed to add further arguments later. In this sense, conservativity is not an artefact but an intrinsic consequence of modeling acceptance as maintainability under monotone commitments.

A useful corollary is that credulous tenability of an argument a is equivalent to the tenability of the singleton $\{a\}$. Accordingly, tenability is best understood as a semantics of *maintainable commitments* rather than of *complete positions*.

One may wonder about the effect of restricting the number of arguments each player can introduce at each stage of the debate. In the case of strong tenability, this has absolutely no impact on which player has a winning strategy. Strong tenability is robust with respect to play size restrictions. This indicates that the winning motto of strong tenability is to play as little as possible, so as not to be forced into contradiction.

Theorem 4.5. *$A \subseteq \mathcal{F}$ is strongly tenable if and only if Pro has a winning strategy starting with A in the strong tenability game where each player may only introduce 1 new argument in each play.*

Tenability is not robust in this sense. Consider $\{a\}$ of the argumentation framework given in Figure 4. The winning strategy to refute a in the tenability game, playing $\{b_1, b_2\}$, is not available when we restrict each player to advancing a single argument at a time. If Opp attacks a with b_1, then he gets stuck defending b_1 from c_1. He would need to play c_2, but this is not a play *more cogent than* $\{a, c_1\}$, thus is not a legal move. Even if we were to allow Opp to play c_2, this would nullify the threat of b_2, highlighting that in the tenability game, b_1 and b_2 need to be played simultaneously for Opp to win.

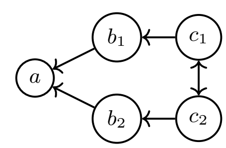

Figure 4: An AF demonstrating the need to play sets of arguments simultaneously.

We end this portion of the investigation by classifying the complexity of each semantics.

Theorem 4.6. *For finite $\mathcal{F}$:*

a. The problem of determining if $A \subseteq \mathcal{F}$ is statically tenable is Π_2^P-complete.

b. The problem of determining if $A \subseteq \mathcal{F}$ is tenable is $PSPACE$-complete.

c. The problem of determining if $A \subseteq \mathcal{F}$ is strongly tenable is $PSPACE$-complete.

In terms of complexity, this shows that tenability and strong tenability are on par with weak admissibility which is also PSPACE complete (Dvořák, Ulbricht, and Woltran 2022). Static tenability is on par with skeptical acceptance for the preferred semantics which is also Π_2^P (Dunne and Bench-Capon 2002). This is evidence for the relative practical applicability of these new semantics that their complexity for implementation is on par with existing semantics.

4.1 Tenability in the Infinite Setting

Though our main focus is on finite argumentation frameworks, we have made definitions for (strong/-/static) tenability which also work well in the infinite setting, and the previous results hold for infinite argumentation frameworks as well. We note that in each definition, we make the choice to allow the set of arguments A under discussion to be infinite, yet we demand that at each stage, each player may only play finitely many new elements. We note that alternate definitions could be made, but we choose the most concrete definition, hewing closely to the goal of modeling debate via dialog. In a debate, for A to be conclusively refuted, it must be refuted via a finite-length debate using finitely many arguments at each stage.

In set theory, games are considered between two players where alternating moves are used to construct some sequence $(X_0, X_1, \dots)$ and the victory condition for each player is determined by membership of the sequence $(X_0, X_1, \dots)$ in some set $\mathcal{S}$. A particular pitfall in this area is that, as a result of the axiom of choice, there can exist games so that neither player possesses a winning strategy (Jech 2003, Chapter 33), or even games that have winning strategies for different players, depending on the chosen model of set theory (Larson and Shelah 2008). This is distinctly not the case for the games we use to define (strong/-/static) tenability. Our notions are well-defined and no set theoretic problems arise:

Theorem 4.7. *If $A \subseteq \mathcal{F}$ is any argumentation framework, then either Pro has a strategy showing that A is (strongly/-/statically) tenable or Opp has a strategy showing that A is not (strongly/-/statically) tenable. Further, the (strong/-/static) tenability of $A \subseteq \mathcal{F}$ does not depend on the model of set theory used to evaluate its (strong/-/static) tenability.*

We also note that it is necessary to consider the tenability of infinite sets $A \subseteq \mathcal{F}$ in order to model consistently simultaneously defending the set A of beliefs. In particular, this is strictly stronger than consistently defending each finite subset. The argumentation framework given in Figure 5 demonstrates this. The set in question is a, together with the unattacked arguments.

Theorem 4.8. *There exists $A \subseteq \mathcal{F}$ so that A is not statically tenable, yet every finite subset of A is strongly tenable.*

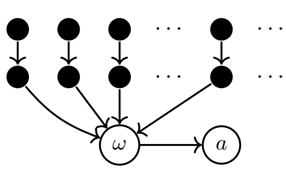

Figure 5: An argumentation framework with a set A, consisting of a along with all unattacked arguments, which is not statically tenable, yet every finite subset is statically tenable.

5 Comparison to Other Semantics

Tenability departs radically from labeling-based, extension-based, and recursion-based semantics. There is no way to label arguments `in`, `out`, `undec` that coincides with our intuitions about tenability. The core idea of a non-uniform defense strategy is not naturally captured by a single labeling. Further, the PSPACE-completeness indicates that tenability cannot be captured by the existence of a certain labeling with some simple properties. The dynamic mechanism of tenability thus separates it from many other weak semantics, but it is included with game-based weak semantics. Here, we will situate tenability among a few semantics that have played key roles in the recent literature.

The *cogency* semantics of Bodanza, Tohmé, and Simari (Bodanza, Tohmé, and Simari 2016) were introduced to directly address issues of self-defeating attack and acceptance in odd-length cycles. They do so by implementing the *as cogent as* relation which allows for comparison between sets of arguments: one set $X \subseteq \mathcal{F}$ is as cogent as another $Y \subseteq \mathcal{F}$, if X is admissible in the restriction of $\mathcal{F}$ to $X \cup Y$ (see Definition 3.1). The isolated comparison of sets of arguments via the cogency relation naturally corresponds to evaluating those sets on how they may be argued against each other in a dialogue. A set is *cogent* if it is undominated with respect to this relation. A weaker notion, *cyclic cogency* relaxes this further in anticipation of the issues presented by odd-length cycles.

Definition 5.1. Let $\mathcal{F}$ be an argumentation framework and $X \subseteq \mathcal{F}$. X is *cogent* if for every $Y \subseteq \mathcal{F}$, $Y \succeq X$ implies $X \succeq Y$. X is *cyclically cogent* if for every $D \subseteq \mathcal{F}$ at least one of the following holds: (1) $D \not\succeq X$; (2) $X \succeq D$; (3) there exists a chain $D = D_1 \prec D_2 \prec \cdots \prec D_n = X$.

The cogency relation plays a key role in defining tenability, however, (strong/-/static) tenability semantics require strictly less than cogent semantics; where the latter requires a fixed singular line of defense for a set of arguments, tenability semantics relaxes this to allow for the defense to change, dependent upon the attacks presented. We note that cogency and cyclic cogency are also captured by game-based semantics (Bodanza, Tohmé, and Simari 2016).

A different approach to defending against illegitimate arguments lies in the recursion-based weak admissibility semantics of Baumann, Brewka, and Ulbricht (Baumann, Brewka, and Ulbricht 2020). Informally, a set of arguments X is ruled acceptable whenever its attackers are rendered unacceptable after removing the arguments attacked by X.

Definition 5.2. For any argumentation framework $\mathcal{F}$ and $A \subseteq \mathcal{F}$, $\mathcal{F}^A$ is the restriction of $\mathcal{F}$ to $F \smallsetminus (A \cup A^+)$, called the *$A$-reduct* of $\mathcal{F}$. A set $A \subseteq \mathcal{F}$ is *weakly admissible* if it is conflict-free and for every attacker y of A, y belongs to no weakly admissible set of $\mathcal{F}^A$.

The last weak semantics we will consider is weakly complete semantics introduced by Dondio and Longo (Dondio and Longo 2021). Weakly complete semantics employs a different mechanism for evaluating acceptance, and capture acceptance in scenarios where undecided arguments may play lessened roles.

Definition 5.3. An extension is *weakly complete* if it is the set of `in`-labeled arguments in a *weakly complete labeling* of $\mathcal{F}$. Such a labeling is a map from $\mathcal{F}$ to $\{\texttt{in}, \texttt{out}, \texttt{undec}\}$ where for every $a \in \mathcal{F}$, each of the following holds: (1) if a is labeled `in`, then there are no attackers of a labeled `in`; (2) if a is labeled `out`, then it has at least one attacker labeled `in`; (3) if a is labeled `undec`, then it has at least one attacker labeled `undec` and it has no attackers labeled `in`.

Weakly complete semantics are by far the weakest semantics in this realm with respect to which arguments are credulously accepted. However, it is not the case that every, e.g., cogent extension is weakly complete, rather, every cogent extension is a subset of a weakly complete extension. We next introduce a formal method of capturing this relation between semantics and show all possible relationships between the semantics discussed in this paper.

Definition 5.4. Let σ and τ be semantics. We say that σ is an inclusion-based refinement of τ if every σ-extension is a subset of a τ-extension.

Observe that if $\sigma \subseteq \tau$, then σ is also an inclusion-based refinement of τ.

Theorem 5.5. *For finite argumentation frameworks, Figure 6 represents the implications and inclusions between weak semantics. Each solid arrow indicates implication. Each dashed arrow indicates inclusion-based refinement. If there is no directed path from one semantics to another, then neither implication holds.*

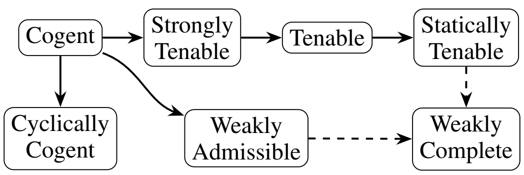

Figure 6: The relations among weak semantics in finite argumentation frameworks.

We remark that some of the refinements among semantics presented in Theorem 5.5 need not persist in infinite argumentation frameworks. In fact, semantics such as weak admissibility fail to be well-defined for arbitrary argumentation frameworks.[2] We leave the project of mapping out the relations among weak semantics in the arbitrary setting to future work.

6 Discussion

This paper introduced three tenability notions—*static tenability*, *strong tenability*, and *tenability*—all centered on the same conceptual shift: defensibility is not necessarily witnessed by a single uniform stance, but may consist in the ability to respond *non-uniformly* to any consistent line of attack. In this sense, tenability adds a new axis to the weak-semantics landscape: it preserves a defense-based intuition while making strategic non-uniformity explicit at the abstract level.

Because our definitions are purely abstract, tenability is not intended as a standalone model of practical reasoning in rich domains. In applications, the main modeling work lies upstream: extracting an AF from structured arguments, rules, evidence, and preferences. Well-established structured argumentation formalisms (and their instantiations into abstract AFs) therefore provide a natural interface, including systems developed with legal and evidential reasoning in mind (Prakken 1997; Modgil and Prakken 2014; Gordon, Prakken, and Walton 2007). Within such pipelines, tenability can be viewed as a *dialectical evaluation layer* that targets a specific phenomenon: whether a position can be sustained against consistent lines of attack without requiring a uniform defense against mutually incompatible challenges.

Adversarial legal procedure is an obvious arena in which the distinction between uniform and non-uniform defense matters, since legal debates intertwine *presumptions*, *burdens of proof*, and varying *standards of proof* (Gordon, Prakken, and Walton 2007; Prakken and Sartor 2006; Walton 2014; Prakken 2009). Tenability does not aim to encode such standards directly; rather, given an AF extracted from an evidential record (possibly via a structured formalism), it can support analyses of whether a party's position is dialectically maintainable against legitimate challenges, allowing conditional rebuttals of the kind highlighted by $\mathcal{U}$. In particular, *static tenability* naturally matches "one-shot" settings where an opponent advances an internally consistent body of evidence and the proponent must be able to extend her commitments with a consistent counter-case. This resonates with the organisation of evidential reasoning around competing hypotheses or "case theories", and with hybrid approaches combining argumentative and narrative structure (Bex et al. 2010).

The game-based notions—*strong tenability* and (especially) *tenability*—are better suited to modeling extended exchanges in which commitments accumulate over time. Here, monotone commitment protocols capture the idea that argumentative moves incur liabilities that cannot simply be ignored later, a theme studied extensively in dialogue theory and commitment-based models of argumentation (Walton and Krabbe 1995; Prakken 2005). In this setting, the distinction between strong tenability and tenability is not merely technical: it separates debates in which one party is held to a stricter standard of sustained defensibility from debates in which both parties are held to this higher standard. This provides a principled handle on the *Master versus Ignoramus* effect isolated in Section 2: strong tenability blocks spurious attacks, giving a suitable model of *Master versus Novice* settings, whereas tenability blocks forms of "attack switching" that would be unreasonable under a defensibility requirement on lines of challenge, while still allowing the proponent's defense to be non-uniform, thus giving a suitable model of *Master versus Master* debates.

Outside legal domains, tenability seems relevant whenever deliberation proceeds against a background of defaults and asymmetric costs of revision. In many practical settings, an agent has a current stance that is not abandoned unless a challenger can articulate a plausible defeating case, an idea closely connected to presumption and burden allocation (Walton 2014; Prakken 2009). This suggests potential uses in interactive explanation and decision support: a system may justify maintaining a position by exhibiting a defense strategy tailored to a user's reasonable objections, without committing to a single all-purpose extension that anticipates mutually incompatible challenges.

[2]Consider whether a_0 is weakly admissible in $(F, \rightarrowtail)$ where $F = \{a_i : i \in \mathbb{N}\}$ and $a_j \rightarrowtail a_i$ for all $i < j$.

Conclusions and Future Work

Our comparison with existing weak semantics (Theorem 5.5 and Figure 6) does not exhaust the landscape. A natural next step is to extend the map to additional weak semantics—for instance, the weakly-preferred and UB-preferred semantics studied by Dondio and Longo (Dondio and Longo 2021). Another direction concerns principle-based evaluation. Since each tenability notion is closed under subset, several standard principles for extension-based semantics will fail for reasons that are orthogonal to the intended reading. Two ways to make such a study informative are to analyze *maximal* variants of the tenability notions, and to extract from tenability new principles that isolate non-uniformity as a design dimension applicable beyond our specific games.

A particularly promising direction concerns the extension of tenability to richer argumentation formalisms, where the strategic character of defense becomes even more salient. Bipolar Argumentation Frameworks (BAFs) (Amgoud et al. 2008) are a natural target: by distinguishing attack from support, BAFs better capture the structure of real dialogues, in which a proponent typically marshals not only counterattacks but also corroborating arguments. Lifting tenability to this setting, however, is not a routine exercise. The design choices that make tenability work in the purely conflict-based setting—legitimacy as conflict-freeness, monotone commitment, unanswered attack as the trigger for proponent moves, counterattack as the source of vindication—must all be rethought once support is present. One must decide, for instance, whether commitments should be closed under supported acceptance, whether the proponent may introduce supporters of arguments already in play as a distinct kind of move, and whether the opponent may legitimately attack a position by attacking its supporters. Analogous questions arise for structured formalisms such as ASPIC+ (Modgil and Prakken 2014), where tenability's strategy-based, monotone character appears especially well-suited to incremental dialectical evaluation.

Looking forward, a promising application of this work lies in the generation of transparent explanations in legal AI. Implementing tenability semantics in automated legal reasoning tools offers a natural solution for this domain, where courtroom disputes inherently require non-uniform, strategic defenses tailored to specific lines of attack. By translating the monotone commitment games introduced here into applied computational frameworks, we aim to develop dialectical tools that assist legal practitioners in evaluating the defensibility of complex claims and formulating targeted, context-aware rebuttals.

To conclude, tenability isolates and formalizes *non-uniform* defense against consistent attacks as a principled alternative to uniform-witness notions of acceptance. We expect this perspective to be most useful when combined with established modeling layers for structure, priorities, and proof standards, rather than as a replacement for them.

Acknowledgements

The first author's work was supported by the National Science Foundation under Grant DMS-2348792. San Mauro is a member of INDAM-GNSAGA.

AI Declaration

AI-based language tools were used for minor editorial assistance. No research content or results were generated using these tools.

References

Amgoud, L.; Cayrol, C.; Lagasquie-Schiex, M.-C.; and Livet, P. 2008. On bipolarity in argumentation frameworks. *International Journal of Intelligent Systems* 23(10):1062–1093.

Baroni, P., and Giacomin, M. 2007. On principle-based evaluation of extension-based argumentation semantics. *Artificial Intelligence* 171(10–15):675–700.

Baroni, P., and Giacomin, M. 2009. Semantics of abstract argument systems. In Rahwan, I., and Simari, G. R., eds., *Argumentation in Artificial Intelligence*. Springer. 25–44.

Baroni, P.; Gabbay, D. M.; Giacomin, M.; and van der Torre, L., eds. 2018. *Handbook of Formal Argumentation*. College Publications.

Baroni, P.; Caminada, M.; and Giacomin, M. 2011. An introduction to argumentation semantics. *The Knowledge Engineering Review* 26(4):365–410.

Baroni, P.; Dunne, P. E.; and Giacomin, M. 2011. On the resolution-based family of abstract argumentation semantics and its grounded instance. *Artificial Intelligence* 175(3–4):791–813.

Baumann, R.; Brewka, G.; and Ulbricht, M. 2020. Revisiting the foundations of abstract argumentation – semantics based on weak admissibility and weak defense. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2742–2749.

Bench-Capon, T. J. M., and Dunne, P. E. 2007. Argumentation in artificial intelligence. *Artificial Intelligence* 171(10–15):619–641.

Bex, F. J.; van Koppen, P. J.; Prakken, H.; and Verheij, B. 2010. A hybrid theory of arguments, stories and criminal evidence. *Artificial Intelligence and Law* 18(2):123–152.

Blümel, L., and Ulbricht, M. 2022. Defining Defense and Defeat in Abstract Argumentation From Scratch – A Generalizing Approach. In *Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning*, 63–72.

Bodanza, G. A., and Tohmé, F. A. 2009. Two approaches to the problems of self-attacking arguments and general odd-length cycles of attack. *Journal of Applied Logic* 7(4):403–420.

Bodanza, G. A.; Tohmé, F. A.; and Simari, G. R. 2016. Beyond admissibility: Accepting cycles in argumentation with game protocols for cogency criteria. *Journal of Logic and Computation* 26(4):1235–1255.

Caminada, M., and Wu, Y. 2009. An argument game for stable semantics. *Logic Journal of the IGPL* 17(1):77–90.

Caminada, M. W. A.; Carnielli, W. A.; and Dunne, P. E. 2012. Semi-stable semantics. *Journal of Logic and Computation* 22(5):1207–1254.

Dauphin, J.; Rienstra, T.; and van der Torre, L. 2020. A principle-based analysis of weakly admissible semantics. In *Computational Models of Argument (COMMA 2020)*, 167–178.

Dondio, P., and Longo, L. 2021. Weakly complete semantics based on undecidedness blocking. https://arxiv.org/abs/2103.10701.

Dung, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence* 77:321–357.

Dunne, P. E., and Bench-Capon, T. J. M. 2002. Coherence in finite argument systems. *Artif. Intell.* 141(1/2):187–203.

Dvořák, W.; Ulbricht, M.; and Woltran, S. 2022. On the complexity of semantics based on weak admissibility. *Journal of Artificial Intelligence Research* 75:203–251.

Gordon, T. F.; Prakken, H.; and Walton, D. 2007. The Carneades model of argument and burden of proof. *Artificial Intelligence* 171(10–15):875–896.

Horty, J. F. 2001. Argument construction and reinstatement in logics for defeasible reasoning. *Artificial Intelligence and Law* 9:1–28.

Jech, T. 2003. *Set Theory: The Third Millennium Edition*. Springer.

Larson, P. B., and Shelah, S. 2008. The stationary set splitting game. *Mathematical Logic Quarterly* 54(2):187–193.

McBurney, P., and Parsons, S. 2009. Dialogue games for agent argumentation. In Rahwan, I., and Simari, G. R., eds., *Argumentation in Artificial Intelligence*. Springer. 261–280.

Modgil, S., and Prakken, H. 2014. The ASPIC+ framework for structured argumentation: A tutorial. *Argument & Computation* 5(1):31–62.

Prakken, H., and Sartor, G. 2006. Presumptions and burdens of proof. In van Engers, T. M., ed., *Legal Knowledge and Information Systems. JURIX 2006: The Nineteenth Annual Conference*, volume 152 of *Frontiers in Artificial Intelligence and Applications*. IOS Press. 21–30.

Prakken, H. 1997. *Logical Tools for Modelling Legal Argument: A Study of Defeasible Reasoning in Law*, volume 32 of *Law and Philosophy Library*. Dordrecht: Kluwer Academic Publishers.

Prakken, H. 2005. Coherence and flexibility in dialogue games for argumentation. *Journal of Logic and Computation* 15(6):1009–1040.

Prakken, H. 2009. A logical analysis of burdens of proof. In Kaptein, H.; Prakken, H.; and Verheij, B., eds., *Legal Evidence and Proof: Statistics, Stories, Logic*. Farnham, UK: Ashgate Publishing. 223–253.

Rahwan, I., and Simari, G. R., eds. 2009. *Argumentation in Artificial Intelligence*. Springer.

van der Torre, L., and Vesic, S. 2017. The principle-based approach to abstract argumentation semantics. *IfCoLog Journal of Logics and Their Applications* 4(8):2735–2778.

Verheij, B. 1996. Two approaches to dialectical argumentation: Admissible sets and argumentation stages. In *Proceedings of NAIC*, 357–368.

Vreeswijk, G. A. W., and Prakken, H. 2000. Credulous and sceptical argument games for preferred semantics. In *Proceedings of JELIA*, 239–253. Springer.

Walton, D. N., and Krabbe, E. C. W. 1995. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. Albany, NY: State University of New York Press.

Walton, D. 2014. *Burden of Proof, Presumption and Argumentation*. Cambridge, UK: Cambridge University Press.