Proceedings of the 21st International Conference on Principles of Knowledge Representation and Reasoning
Main Track

Balancing Open-Mindedness and Conservativeness in Quantitative Bipolar Argumentation (and How to Prove Semantical from Functional Properties)

Nico Potyka , Richard Booth
Cardiff University
{potykan, boothr2}@cardiff.ac.uk

Abstract

Quantitative bipolar argumentation frameworks (QBAFs) have various applications in areas like product recommendation, review aggregation and explaining machine learning models. QBAF semantics assign a strength to every argument that is based on an a priori belief and the strength of its attackers and supporters. Intuitively, a QBAF semantics is open-minded when it is unbiased in the sense that a priori beliefs can be given up eventually when sufficient arguments to the contrary are presented. While this behaviour is desirable in many applications, existing open-minded semantics also have the property that even very weak arguments will eventually eliminate the a priori beliefs. In this paper, we will study notions of conservativeness that demand that the deviation from the a priori beliefs is bounded by the strength of pro and contra arguments. We will discuss compatibility and conflicts with existing properties and present two new semantics with interesting semantical guarantees. To do so, we will build up on the framework of modular semantics and prove some general relationships between functional and semantical properties that are useful to simplify the study of new modular semantics.

1 Introduction

Quantitative bipolar argumentation frameworks (QBAFs) are abstract argumentation frameworks (Dung 1995) that represent attack and support relationships between arguments (Amgoud et al. 2008; Amgoud and Ben-Naim 2017; Baroni, Rago, and Toni 2018). Each argument is assigned a base score (an a priori belief) and QBAF semantics assign a final strength to every argument based on its base score and the strength of its attackers and supporters. QBAFs have applications in various domains like product recommendation (Rago, Cocarascu, and Toni 2018), review aggregation (Cocarascu, Rago, and Toni 2019), stance aggregation (Kotonya and Toni 2019), explaining the PageRank of websites (Albini et al. 2020) or the predictions of neural networks (Ayoobi, Potyka, and Toni 2023).

In many applications, we want to balance between *conservativeness* and *open-mindedness* (Potyka 2019b): a semantics should be conservative in the sense that it does not give up a priori beliefs too easily, but it should be open-minded in the sense that a priori beliefs can be given up eventually when sufficient evidence to the contrary is presented. As we will discuss later in detail, existing open-minded semantics have the property that even very weak arguments can make another argument very strong or weak when they occur in large numbers. This can be undesirable. For example, when evaluating a movie based on arguments made in reviews, a large number of mediocre evaluations of an aspect in individual reviews (e.g., the acting was okay) should not result in an outstanding evaluation of the overall aspect (the acting was great). Another example are truth discovery problems, where we want to evaluate the trustworthiness of sources and the believability of claims (Gupta and Han 2011; Berti-Equille and Borge-Holthoefer 2015; Li et al. 2016; Singleton and Booth 2022). As noted in (Singleton 2020), such problems can be naturally represented by bipolar argumentation frameworks where sources and their claims support each other (a claim is more believable if it is supported by trustworthy sources and a source is more trustworthy if it makes believable claims) and contradicting claims attack each other. We recently implemented this idea with QBAFs (Potyka and Booth 2024). In this setting, we want open-mindedness, but we also want to be able to bound the believability of a claim by the trustworthiness of its supporters from above.

We will study two notions of conservativeness. The first notion is called *absolute conservativeness (AC)*. Roughly speaking, AC directly restricts the strength values that an argument can take based on the maximum strength of its attackers and supporters. However, implementing absolute conservativeness turns out to be intricate and may severely restrict the potential final strength values based on the a priori beliefs. We therefore also introduce the weaker notion of *relative conservativeness (RC)* that allows adjusting the base score relative to the strength of attackers and supporters. We discuss limitations of existing semantics and introduce two novel semantics that satisfy conservativeness properties and many other interesting properties from the literature.

In order to study the new semantics efficiently, we build up on the framework of modular semantics (Mossakowski and Neuhaus 2018) that encompasses commonly applied QBAF semantics like Df-QuAD (Rago et al. 2016), Euler-based (Amgoud and Ben-Naim 2017), quadratic energy (Potyka 2018a) and MLP-based semantics (Potyka 2021). All modular semantics define the strength values by means of an update function that is composed of an aggregation function

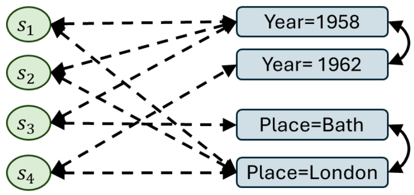

Figure 1: A truth discovery QBAF. Source (green) and claim (blue) arguments have base scores 0.5 and 0, respectively.

(aggregating the strength of an argument's attackers and supporters) and an influence function (adapting the base score based on the aggregate). In recent papers like (Potyka 2021), semantics have been compared based on 14 semantical properties. However, as shown in (Baroni, Rago, and Toni 2018), many of these properties can be seen as instances of more general principles called *Balance* and *Monotonicity*. We will connect Balance and Monotonicity and some other independent semantical principles to properties of aggregation and influence functions to further simplify the study of new modular semantics. Similar connections have been made in (Amgoud and Doder 2018) for gradual semantics of attack frameworks and some of our results can be seen as generalizations of their results to bipolar frameworks. We will discuss the relationship at the end of the paper (Section 6).

The remainder of this paper is structured as follows: We give a quick introduction to QBAFs in Section 2. In Section 3, we will discuss interesting properties of aggregation and influence functions, and explain, in Section 4, how they are related to the semantical properties of the induced QBAF semantics. We will also show that every modular semantics satisfies some basic properties from (Amgoud and Ben-Naim 2017) by definition. These results do not only simplify the investigation of our new semantics, but will hopefully also simplify the design and study of future modular semantics. In Section 5, we will introduce absolute and relative conservativeness formally and discuss limitations of existing semantics. We will introduce a new influence and aggregation function and discuss how they can be combined with existing functions to define novel semantics that can provide a better tradeoff between conservativeness and openmindedness.

2 Background

Formally, QBAFs can be defined as follows (Baroni, Rago, and Toni 2018).

Definition 1 (QBAF)**.** *A QBAF is a quadruple* $Q = (\mathcal{A}, \mathrm{Att}, \mathrm{Sup}, \beta)$ *consisting of a set of arguments* $\mathcal{A}$*, two binary relations* Att *and* Sup *called attack and support and a function* $\beta : \mathcal{A} \to [0, 1]$ *that assigns a* base score $\beta(a)$ *to every argument* $a \in \mathcal{A}$*.*

For all $a \in \mathcal{A}$, we let $\mathrm{Att}(a) = \{b \in \mathcal{A} \mid (b, a) \in \mathrm{Att}\}$ and define $\mathrm{Sup}(a)$ analogously. Figure 1 shows a QBAF that encodes a truth discovery problem as suggested in (Singleton 2020). In this example, we have four sources (green) that make claims (blue) about the place and year of birth of a person. We denote support relationships by dashed and attack relationships by solid edges. We assign base score 0.5 to sources (we are initially indifferent about sources) and 0 to claims (we initially reject all claims).

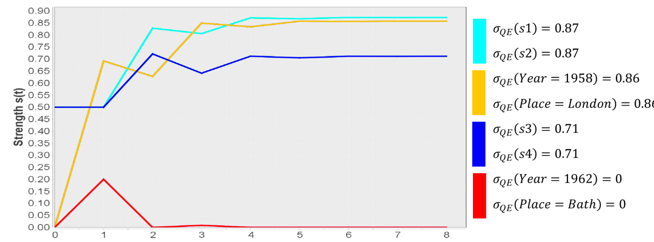

Figure 2: Evolution of strength values under quadratic energy semantics for the QBAF in Figure 1. X-axis shows iteration, Y-axis shows strength values in the iteration.

QBAF semantics associate QBAFs with a strength function that assigns a final strength to every argument.

Definition 2 (QBAF semantics, strength function)**.** *A QBAF semantics is a partial function that maps a QBAF* $Q = (\mathcal{A}, \mathrm{Att}, \mathrm{Sup}, \beta)$ *to a* strength function $\sigma_Q : \mathcal{A} \to D$*.*

QBAF semantics are only partially defined because they are often based on an update procedure that may not terminate for some QBAFs (Mossakowski and Neuhaus 2018). The update procedure of common QBAF semantics (Rago et al. 2016; Amgoud and Ben-Naim 2017; Potyka 2018a; Potyka 2021) initializes the strength values of all arguments with their base scores and then repeatedly applies the following two steps until the strength values converge:

Aggregation: for every argument, use an *aggregation function* to aggregate the strength values of attackers and supporters.

Influence: for every argument, use an *influence function* to set the new strength to a value based on the base score and the aggregate.

We will discuss aggregation and influence functions and the update process in detail in the next section. For now, we just illustrate the process in Figure 2 by plotting the evolution of the strength values of arguments for the QBAF from Figure 1 under the quadratic energy semantics from (Potyka 2018a).

3 Aggregation and Influence Functions

QBAF semantics that can be defined in terms of an aggregation and influence function are called *modular semantics* (Mossakowski and Neuhaus 2018). In this section, we will go into some details of the definitions. Along the way, we will develop some new results that will simplify the study of future modular semantics.

3.1 Aggregation Functions

The *aggregation function* is the most intricate part of modular semantics. Strictly speaking, modular semantics are not based on a single aggregation function, but on a family of aggregation functions of the form $\alpha_{m,n} : \mathbb{R}^{n+m} \to \mathbb{R}$ for all $m, n \in \mathbb{R}^n$. Intuitively, the first n parameters of $\alpha_{m,n}$ are the strength values of an argument's attackers and the last m the strength values of its supporters. We assume that $\alpha_{m,n}$ is *permutation-invariant* with respect to the first m and the last n parameters, that is, the order of the strength values does not matter. This assumption is satisfied for all modular semantics that we are aware of and allows us to simplify the notation $\alpha_{m,n}(a_1, \ldots, a_m, s_1, \ldots, s_n)$ to $\alpha(A, S)$ where $A = \{a_1, \ldots, a_m\}$ and $S = \{s_1, \ldots, s_n\}$ are multisets. We give three examples of aggregation functions from the literature.

$$\alpha_\Sigma(A, S) = \sum_{s \in S} s - \sum_{a \in A} a, \qquad \text{(Sum)}$$

$$\alpha_\Pi(A, S) = \prod_{a \in A}(1 - a) - \prod_{s \in S}(1 - s), \qquad \text{(Product)}$$

$$\alpha_\top(A, S) = \max_{a \in A}(1 - a) - \max_{s \in S}(1 - s). \qquad \text{(Top)}$$

α_Σ is used in the Euler-based (Amgoud and Ben-Naim 2017), quadratic energy (Potyka 2018a) and MLP-based (Potyka 2021) semantics, α_Π in the Df-QuAD (Rago et al. 2016) semantics. The intuition of α_Π can be understood by interpreting the first term as a positive bias that can be decreased by the attackers and the second term as a negative bias that can be decreased by the supporters. $\alpha_\top$ has been introduced in (Mossakowski and Neuhaus 2018) to define a semantics that guarantees well-defined strength values for all QBAFs. While the introduced semantics has some semantical shortcomings caused by the choice of the influence function, $\alpha_\top$ itself has some nice properties that we will come back to later.

Aggregation functions usually ignore arguments with strength 0. In order to simplify referring to the non-zero strength values, we make the following definition.

Definition 3 (Core)**.** *For a multiset S of real numbers, we let $Core(S) = \{x \in S \mid x \neq 0\}$ denote the sub-multiset that contains only the non-zero elements.*

We call an aggregation function that depends only on the core of its inputs *neutral*.

Definition 4 (Neutrality)**.** *An aggregation function α satisfies* Neutrality *iff $\alpha(A, S) = \alpha(Core(A), Core(S))$.*

Intuitively, an aggregation function should return a positive (negative) value if the supporters (attackers) dominate the attackers (supporters) and zero if they are in balance. The following definition will be helpful later to make these ideas precise.

Definition 5 (Domination and Balance)**.** *Let S, T be multisets of real numbers.*

- *S* dominates *T if $Core(S) = Core(T) = \emptyset$ or there is a sub-multiset $S' \subseteq Core(S)$ and a bijective function $f : Core(T) \to S'$ such that $x \leq f(x)$ for all $x \in Core(T)$. If, in addition, $|Core(S)| > |Core(T)|$ or there is an $x \in Core(T)$ such that $x < f(x)$, S* strictly dominates *T.*
- *S and T are* balanced *if $Core(S) = Core(T)$.*

We write $S \succeq T$ ($S \succ T$) if S (strictly) dominates T and $S \cong T$ if they are balanced.

As we explain in the following lemma, $\succeq$ and $\cong$ are related as the symbols suggest.

Lemma 6. *$S \cong T$ if and only if $S \succeq T$ and $T \succeq S$.*

Proof. First assume that $S \cong T$. To see that $S \succeq T$, just let $S' = Core(S)$ in the definition of domination and let f be the identity. $T \succeq S$ follows symmetrically.

Now assume that $S \succeq T$ and $T \succeq S$. By domination, there are bijective mappings from $Core(S)$ to a subset of $Core(T)$ and from $Core(T)$ to a subset of $Core(S)$, which implies that $|Core(S)| = |Core(T)|$. Let $n = |Core(T)|$ and order the elements in S and T increasingly and denote them by $s_1, \ldots, s_n$ and $t_1, \ldots, t_n$, respectively. If $s_1 < t_1$, then S could not dominate T and if $s_1 > t_1$, T could not dominate S. Hence, $s_1 = t_1$. Continuing in this way, we can show that $s_i = t_i$ for all $i = 1, \ldots, n$. Thus, the mapping that maps $f(s_i) = t_i$ shows that S and T are balanced. □

Many aggregation functions satisfy the following monotonicity properties.

Definition 7 ((Strict) Monotonicity)**.** *An aggregation function α satisfies* Monotonicity *iff*

1. *$\alpha(A, S) \leq 0$ if $A \succeq S$,*
2. *$\alpha(A, S) \geq 0$ if $S \succeq A$,*
3. *$\alpha(A, S) \leq \alpha(A', S)$ if $A \succeq A'$,*
4. *$\alpha(A, S) \geq \alpha(A, S')$ if $S \succeq S'$.*

If the inequalities can be replaced by strict inequalities for strict domination, α satisfies Strict Monotonicity.

Monotonicity has an important implication that we call *Balance*.

Lemma 8 (Balance)**.** *If α satisfies monotonicity, then*

1. *$\alpha(A, S) = 0$ whenever $A \cong S$.*
2. *$\alpha(A, S) = \alpha(A', S')$ whenever $A \cong A'$ and $S \cong S'$.*

Proof. 1. Lemma 6 implies that A dominates S and vice versa. Hence, monotonicity implies $0 \leq \alpha(A, S) \leq 0$, that is, $\alpha(A, S) = 0$.

2. Lemma 6 implies that $A \succeq A' \succeq A$ and $S \succeq S' \succeq S$. Hence, $\alpha(A, S) \geq \alpha(A', S) \geq \alpha(A', S') \geq \alpha(A', S) \geq \alpha(A, S)$ by monotonicity of α and thus $\alpha(A, S) = \alpha(A', S')$. □

There are two more technical properties of aggregation functions that are relevant for the duality and open-mindedness properties that we will discuss later. Roughly speaking, duality (Potyka 2018a) requires that the positive effect of supporters equals the negative effect of attackers. Open-mindedness requires that adding supporters/attackers indefinitely will always result in the same upper/lower bound independent of the original set of supporters/attackers.

Definition 9. *An aggregation function α satisfies*

- Duality *if $\alpha(A, S) = -\alpha(S, A)$,*
- Open-Mindedness with range (l, u) *if there are $l, u \in \mathbb{R}_\infty$ such that for all finite multisets A, S of real numbers, $\lim_{n\to\infty} \alpha(A \sqcup E_n, S) = l$ and $\lim_{n\to\infty} \alpha(A, S \sqcup E_n) = u$, where E_n is the multiset that contains n copies of 1.*

All of our aggregation functions satisfy duality (Potyka 2019a)[Proposition 5.4] and it is easy to check that α_Σ satisfies open-mindedness with range $(-\infty, \infty)$. We have some additional guarantees that follow from elementary properties of the arithmetic operations involved in the definitions.

Lemma 10. *α_Σ, α_Π and $\alpha_\top$ satisfy monotonicity, balance, neutrality and duality. α_Σ also satisfies strict monotonicity and open-mindedness with range* $(-\infty, \infty)$.

To see that α_Π and $\alpha_\top$ do not satisfy Strict Monotonicity, consider $A = \{1, 1\}$, $S = \{1\}$. Since A dominates S, strict monotonicity demands that the aggregate is negative, but $\alpha_\Pi(A, S) = 0 - 0 = 0 = 1 - 1 = \alpha_\top(A, S)$. α_Π and $\alpha_\top$ do not satisfy open-mindedness because the upper and lower bounds depend on A and S. For example, if $A = \emptyset$, a single supporter with strength 1 will make the aggregate 1, but if $A = \{0.5\}$, the aggregate can never become larger than 0.5 no matter how many supporters with strength 1 we add.

3.2 Influence Functions

The influence function of modular semantics is relatively straightforward. Formally, it has the form $\iota : \mathbb{R}^2 \to \mathbb{R}$. The first parameter is the base score of an argument and the second is the aggregate obtained from the aggregation function. Intuitively, a positive aggregate should increase and a negative aggregate should decrease the base score. ι has to make sure that the result falls into the desired strength domain again. We give four examples of aggregation functions from the literature.

$$\iota_l(b, a) = b + b \cdot \min\{0, a\} + (1 - b) \cdot \max\{0, a\}, \quad \text{(Linear)}$$

$$\iota_e(b, a) = 1 - (1 - b^2)/(1 + b \cdot \exp(a)), \quad \text{(Euler)}$$

$$\iota_q(b, a) = b - b \cdot h(-a) + b \cdot h(a), \quad \text{(QE)}$$

where $h(a) = (\max\{0, a\})^2/(1 + \max\{0, a\})^2)$,

$$\iota_m(b, a) = \phi_l(\ln(b/(1 - b)) + a), \quad \text{(MLP)}$$

where $\phi_l(x) = 1/(1 + \exp(-z))$ is the logistic function, and we let $\ln 0 := -\infty$, $\ln 1/0 := \infty$, $\phi_l(-\infty) := 0$ and $\phi_l(\infty) := 1$.

ι_l is used with α_Π in Df-QuAD (Rago et al. 2016), but could also naturally be combined with $\alpha_\top$ since both of them create aggregates between -1 and 1. ι_e has been used with α_Σ for the Euler-based (Amgoud and Ben-Naim 2017) semantics and with $\alpha_\top$ in (Mossakowski and Neuhaus 2018) to create a semantics that always guarantees well-defined strength values. However, since ι_e has been defined for an unbounded aggregation function, the strength values under this semantics are very conservative in the sense that they cannot differ much from the base scores (Potyka 2019a). ι_q and ι_m can be combined with α_Σ to define the quadratic energy (Potyka 2018a) and MLP-based semantics (Potyka 2021), respectively.

Balance and monotonicity have the following counterpart for influence functions.

Definition 11 ((Strict) Monotonicity)**.** *An influence function ι satisfies* monotonicity *if*

1. *$\iota(b, a) \leq b$ if $a < 0$,*
2. *$\iota(b, a) \geq b$ if $a > 0$,*
3. *$\iota(b_1, a) \leq \iota(b_2, a)$ if $b_1 < b_2$,*
4. *$\iota(b, a_1) \leq \iota(b, a_2)$ if $a_1 < a_2$.*

If the inequalities can be replaced by strict inequalities, when excluding $b = 0$ for the first and $b = 1$ for the second item, ι satisfies strict monotonicity.

Definition 12 (Balance)**.** *An influence function ι satisfies* Balance *if $\iota(b, 0) = b$.*

We have again two properties of influence functions that are relevant for duality and open-mindedness of modular semantics.

Definition 13. *An influence function ι satisfies*

- Duality *if $1 - \iota(b, a) = \iota(1 - b, -a)$,*
- Open-mindedness *for range $(l, u) \in \mathbb{R}^2_\infty$ if $\lim_{a \to l} \iota(b, a) = 0$ and $\lim_{a \to u} \iota(b, a) = 1$.*

In natural language, duality states that the difference between 1 and the strength for base score b and aggregate a $(1 - \iota(b, a))$ should be equal to the difference between 0 and the strength for base score $1 - b$ and aggregate $-a$ $(\iota(1 - b, -a) - 0)$. Intuitively, this means that a positive aggregate brings the strength equally fast to 1 as a negative aggregate brings it to 0. The open-mindedness property basically guarantees that when we combine aggregation and influence functions with the same range, then adding strong attackers/supporters can bring the strength arbitrarily close to 0/1 independent of our a priori beliefs.

ι_l, ι_q and ι_m satisfy duality as shown in (Potyka 2019a, Proposition 5.1) and (Potyka 2021, Theorem 3), respectively. The remaining claims in the following lemma follow again from basic arithmetic properties and some simple limit considerations for open-mindedness.

Lemma 14. *$\iota_l, \iota_e, \iota_q, \iota_m$ satisfy monotonicity, strict monotonicity and balance. ι_l, ι_q and ι_m satisfy duality. ι_l and $\iota_\top$ are open-minded for range $(-1, 1)$ and ι_q for range* $(-\infty, \infty)$.

We refer to (Potyka 2019a) for a detailled discussion of why ι_e satisfies neither duality nor open-mindedness. ι_m almost satisfies open-mindedness, but fails to do so for $b \in \{0, 1\}$. We refer to (Potyka 2021) for a more thorough discussion of the reasons.

4 Modular Semantics

4.1 Discrete and Continuous Semantics

We will refer to a pair of aggregation and influence function as the *kernel* of a modular semantics in the following. We call it *elementary* if it satisfies some useful properties that we summarize in the following definition.

Definition 15 ((Elementary) Kernel)**.** *A* kernel *is a pair (α, ι) consisting of an aggregation and incluence function. (α, ι) is called* elementary *if both are Lipschitz-continuous, satisfy monotonicity and balance, and α additionally satisfies neutrality.*

Lipschitz-continuity is a technical condition that is helpful for proving convergence guarantees of modular semantics (Potyka 2019a). Intuitively, the growth of a Lipschitz-function is bounded by a constant at any point. The precise definition is not important for this paper, but we mention it here because it is satisfied by all commonly considered modular semantics and important for some technical results.

From each kernel, we can define a *discrete* and a *continuous* modular semantics. Roughly speaking, both can be defined as the limit of a strength evolution process $s : \mathbb{R}_{\geq 0} \rightarrow [0,1]^{|\mathcal{A}|}$, where $s_i(t)$ is the strength of the i-th argument at time t (Potyka 2018a; Potyka 2019a). For discrete semantics, we consider discrete time steps and the strength values at time $t \in \mathbb{N}$ are obtained by updating the base score t times with respect to the aggregation and influence function. This is illustrated in Figure 2, where the x-axis shows t and we plot the components $s(t)_i$ in different colors against the y-axis that shows the strength values. More formally, given a kernel (α, ι) and a QBAF $Q = (\mathcal{A}, \mathrm{Att}, \mathrm{Sup}, \beta)$, we can assume without loss of generality that $\mathcal{A} = \{1, \dots, n\}$. We can then define s by letting

- $s(0)_i = \beta(i)$,
- $s(t+1)_i = \iota(\beta(i), \alpha(\{s(t)_i \mid i \in \mathrm{Att}(a)\}, \{s(t)_i \mid i \in \mathrm{Sup}(a)\})$

for all $i \in \mathcal{A}$ and $t \in N$. The strength function under discrete modular semantics is then defined by the limit

$$\sigma^d(i) = \lim_{t \rightarrow \infty} s(t)_i. \quad (1)$$

To explain the continuous semantics, let u_Q denote the *update function* that maps $s(t)$ to $s(t+1)$ for a QBAF Q as described above and consider the system of differential equations

$$\frac{ds_i}{dt} = u_Q(s)_i - s_i, \quad (i \in \mathcal{A}). \quad (2)$$

with initial condition $s(0)_i = \beta(i)$. As shown in (Potyka 2019a, Proposition 4.1), if the aggregation and influence functions are (Lipschitz-)continuous, there is a unique solution s^* and the continuous semantics is defined by the limit

$$\sigma^c(i) = \lim_{t \rightarrow \infty} s^*(t)_i. \quad (3)$$

Intuitively, s^* continizes s, which can improve stability. That is, in cases where the discrete semantics fails to converge, the continuous semantics may still have a well-defined limit. To illustrate this, Figure 3 shows, on the left, the evolution of strength values for the QBAF in Figure 1 under discrete quadratic energy (top) and Df-QuAD semantics (bottom). Notably, the strength function under Df-QuAD is not defined because the strength values start cycling. On the right, we show the evolution under continuous semantics. Note that the continuous quadratic energy semantics converges to the same limit as its discrete counterpart in a smoother way. The continuous Df-QuAD semantics converges to a meaningful limit even though its discrete counterpart does not.

We are currently unaware of examples where continuous semantics fail to converge and in all examples that we are aware of where both converge, they converge to the same limit. However, we have not been able to prove that this happens in general. The technical details of discrete and continuous semantics are not important for this paper and we refer to (Potyka 2019a) for more details. The most important result for us is that if the limits exist, they correspond to a fixed-point of the update function defined by aggregation and influence function. This allows us to study semantical properties of continuous and discrete semantics in a uniform way even if we cannot guarantee that the fixed-points are equal in general.

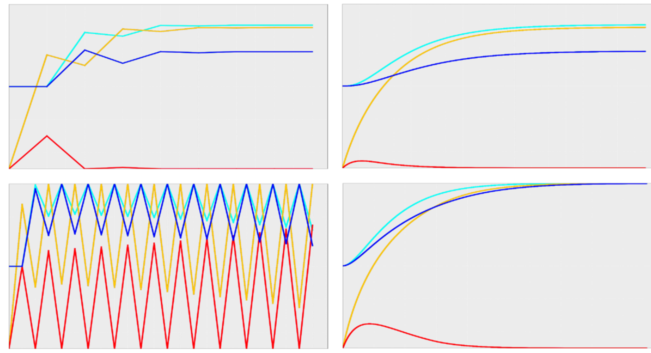

Figure 3: Evolution of strength values under discrete (left) and continuous (right) quadratic energy (top) and Df-QuAD (bottom) semantics for the QBAF in Figure 1.

Definition 16 ((Elementary) Discrete/Continuous Modular Semantics)**.** *The discrete/continuous modular semantics generated by the kernel* (α, ι) *maps a QBAF* Q *to the limit* (1)/(3) *if it exists. The semantics is called* elementary *if the kernel is elementary.*

Proposition 17 ((Potyka 2019a))**.** *If the strength function* σ *under discrete/continuous elementary modular semantics is defined for a QBAF* Q*, then it is a fixed point of the update function* u_Q*.*

It is interesting to note that, in cases where we can prove convergence of discrete semantics, u_Q has a unique fixed-point and so the discrete and continuous semantics are guaranteed to be equal in these cases (Potyka 2019a).

4.2 General Properties of Modular Semantics

Before introducing our new semantics and studying their properties, it is useful to notice some general relationships between properties of aggregation and influence functions and the modular semantics that they induce. As noted in (Baroni, Rago, and Toni 2018), many properties considered in the literature come down to two principles that have been called *balance* and *monotonicity*. Intuitively, if the attackers and supporters of an argument are balanced, then its final strength should just be its base score (*balance*). If one dominates the other, the base score should be increased or decreased accordingly (*monotonicity*). To make this intuition more precise, we first define the notions of domination and balance between sets of arguments.

Definition 18 (Domination and Balance)**.** *Let* $Q = (\mathcal{A}, \mathrm{Att}, \mathrm{Sup}, \beta)$ *be a QBAF such that* σ_Q *is well-defined.*

For all $X \subseteq \mathcal{A}$, let $\sigma_Q(X) = \{\sigma(a) \mid a \in X\}$ denote the multiset of strength values of arguments in X. For all $X_1, X_2 \subseteq \mathcal{A}$, we say that

- X_1 (strictly) dominates X_2 *if* $\sigma_Q(X_1) \succeq \sigma_Q(X_2)$ *($\sigma_Q(X_1) \succ \sigma_Q(X_2)$) and denote this by $X_1 \succeq_Q X_2$ ($X_1 \succ_Q X_2$).*
- *X_1 and X_2 are* balanced *if $\sigma_Q(X_1) \cong \sigma_Q(X_2)$ and denote this by $X_1 \cong_Q X_2$.*

As we note in the following theorem, all modular semantics satisfy some basic properties introduced in (Amgoud and Ben-Naim 2017), and every elementary modular semantics satisfies non-strict balance and monotonicity properties (Baroni, Rago, and Toni 2018). For the anonymity property from (Amgoud and Ben-Naim 2017), we need one additional definition.

Definition 19 (QBAF isomorphism)**.** *A* (QBAF) isomorphism *π between two QBAFs $Q = (\mathcal{A}, \mathrm{Att}, \mathrm{Sup}, \beta)$ and $Q' = (\mathcal{A}', \mathrm{Att}', \mathrm{Sup}', \beta')$ is a bijective function $\pi : \mathcal{A} \rightarrow \mathcal{A}'$ such that $\mathrm{Att}' = \{(\pi(a), \pi(b)) \mid (a,b) \in \mathrm{Att}\}$, $\mathrm{Sup}' = \{(\pi(a), \pi(b)) \mid (a,b) \in \mathrm{Sup}\}$ and $\beta(a) = \beta'(\pi(a))$ for all $a \in \mathcal{A}$.*

Theorem 20. *Let $Q = (\mathcal{A}, \mathrm{Att}, \mathrm{Sup}, \beta)$ and $Q' = (\mathcal{A}', \mathrm{Att}', \mathrm{Sup}', \beta')$ be QBAFs such that σ_Q and $\sigma_{Q'}$ are well-defined. Every modular semantics satisfies*

Anonymity: *If there is an isomorphism π between Q and Q', then $\sigma_Q(a) = \sigma_{Q'}(\pi(a))$ for all $a \in \mathcal{A}$.*

Independence: *If $\mathcal{A} \cap \mathcal{A}' = \emptyset$ and $Q'' = (\mathcal{A} \cup \mathcal{A}', \mathrm{Att} \cup \mathrm{Att}', \mathrm{Sup} \cup \mathrm{Sup}', \beta'')$, where $\beta''(a) = \beta(a)$ for all $a \in \mathcal{A}$ and $\beta''(a') = \beta'(a')$ for all $a' \in \mathcal{A}'$, then $\sigma_{Q''}$ is well-defined and $\sigma_{Q''}(a) = \sigma_Q(a)$ for all $a \in \mathcal{A}$ and $\sigma_{Q''}(a') = \sigma_{Q'}(a')$ for all $a' \in \mathcal{A}'$.*

Directionality: *If $\mathcal{A} = \mathcal{A}'$, $\beta = \beta'$ and either ($\mathrm{Att}' = \mathrm{Att} \cup \{(a,b)\}$ and $\mathrm{Sup}' = \mathrm{Sup}$) or ($\mathrm{Att}' = \mathrm{Att}$ and $\mathrm{Sup}' = \mathrm{Sup} \cup \{(a,b)\}$), then for all $c \in \mathcal{A}$ such that there is no directed path from b to c, $\sigma(c) = \sigma'(c)$.*

Every elementary modular semantics satisfies additionally

Individual A-Monotonicity: *For all $a \in \mathcal{A}$, if $\mathrm{Att}(a) \succeq_Q \mathrm{Sup}(a)$, then $\sigma_Q(a) \leq \beta(a)$.*

Individual S-Monotonicity: *For all $a \in \mathcal{A}$, if $\mathrm{Sup}(a) \succeq_Q \mathrm{Att}(a)$, then $\sigma_Q(a) \geq \beta(a)$.*

Relative Monotonicity: *For all $a, b \in \mathcal{A}$, if $\beta(a) \leq \beta(b)$, $\mathrm{Att}(a) \succeq_Q \mathrm{Att}(b)$ and $\mathrm{Sup}(b) \succeq_Q \mathrm{Sup}(a)$, then $\sigma_Q(a) \leq \sigma_Q(b)$.*

Individual Balance: *For all $a \in \mathcal{A}$, if $\mathrm{Att}(a) \cong_Q \mathrm{Sup}(a)$, then $\sigma_Q(a) = \beta(a)$.*

Relative Balance: *For all $a, b \in \mathcal{A}$, if $\beta(a) = \beta(b)$, $\mathrm{Att}(a) \cong_Q \mathrm{Att}(b)$ and $\mathrm{Sup}(a) \cong_Q \mathrm{Sup}(b)$, then $\sigma_Q(a) = \sigma_Q(b)$.*

Proof. For brevity, for all $x \in \mathcal{A}$, we let $B_x = \beta(x)$, $S_x = \sigma_Q(\mathrm{Sup}(x))$ and $A_x = \sigma_Q(\mathrm{Att}(x))$ in the following.

Anonymity: Follows immediately by observing that the arguments $a \in \mathcal{A}$ and there counterparts $\pi(a) \in \mathcal{A}'$ will be evaluated in the exact same way.

Independence: Follows immediately by observing that the arguments from Q and Q' will be evaluated independently.

Directionality: Follows by observing that the aggregation function depends only on the parents of an argument. Hence, the new edge can affect the evaluation of c only if b is a predecessor of c, which is the case iff there is a directed path from b to c

Individual A-Monotonicity: Since, σ_Q is a fixed-point of u_Q, we have $\sigma_Q(a) = \iota(B_a, \alpha(A_a, S_a))$. $A_a \succeq S_a$ and monotonicity of α imply $\alpha(A_a, S_a) \leq 0$. Hence, $\iota(B_a, \alpha(A_a, S_a)) \leq B_a$ by monotonicity of ι.

Individual S-Monotonicity: follows analogously.

Rel. Monotonicity: $\sigma_Q(a) = \iota(B_a, \alpha(A_a, S_a)$ and $\sigma_Q(b) = \iota(B_b, \alpha(A_b, S_b)$ because σ_Q is a fixed-point of u_Q. Since $A_a \succeq A_b$ and $S_b \succeq S_a$, monotonicty of α implies $\alpha(A_a, S_a) \leq \alpha(A_b, S_a) \leq \alpha(A_b, S_b)$. Then monotonicity of ι implies $\iota(B_a, \alpha(A_a, S_a)) \leq \iota(B_b, \alpha(A_a, S_a)) \leq \iota(B_b, \alpha(A_b, S_b))$. Hence, $\sigma_Q(a) \leq \sigma_Q(b)$.

Individual Balance: Proposition 6 together with Individual A- and S-monotonicity implies $\sigma_Q(a) \leq \beta(a)$ and $\sigma_Q(a) \geq \beta(a)$. Hence, $\sigma_Q(a) = \beta(a)$.

Relative Balance: Proposition 6 together with Relative monotonicity implies $\sigma_Q(a) \leq \sigma_Q(b)$ and $\sigma_Q(a) \geq \sigma_Q(b)$. Hence, $\sigma_Q(a) = \sigma_Q(b)$. □

Let us emphasize that, in the proof of individual and relative balance, we also proved the following relationship.

Corollary 21. *1. Individual A- and S-monotonicity together imply Individual Balance.*

2. Relative Monotonicity implies Relative Balance.

We explain next, what semantical properties the remaining properties of aggregation and influence functions entail.

Theorem 22. *Let Q and σ_Q be defined as in Theorem 20. Every elementary modular semantics such that α and ι satisfy strict monotonicity satisfies*

Strict Individual A-Monotonicity: *For all $a \in \mathcal{A}$, if $\beta(a) > 0$ and $\mathrm{Att}(a) \succ_Q \mathrm{Sup}(a)$, then $\sigma_Q(a) < \beta(a)$.*

Strict Individual S-Monotonicity: *For all $a \in \mathcal{A}$, if $\beta(a) < 1$ and $\mathrm{Sup}(a) \succ_Q \mathrm{Att}(a)$, then $\sigma_Q(a) > \beta(a)$.*

Strict Relative Monotonicity: *For all $a, b \in \mathcal{A}$, if $0 < \beta(a) \leq \beta(b) < 1$, $\mathrm{Att}(a) \succeq_Q \mathrm{Att}(b)$ and $\mathrm{Sup}(b) \succeq_Q \mathrm{Sup}(a)$ and ($\beta(a) < \beta(b)$ or $\mathrm{Att}(a) \succ_Q \mathrm{Att}(b)$ or $\mathrm{Sup}(b) \succ_Q \mathrm{Sup}(a)$), then $\sigma_Q(a) < \sigma_Q(b)$.*

Every elementary modular semantics such that α and ι satisfy duality satisfies

Duality: *For all $a, b \in \mathcal{A}$, if $\beta(a) = 1 - \beta(b)$, $\mathrm{Att}(a) \cong_Q \mathrm{Sup}(b)$ and $\mathrm{Sup}(a) \cong_Q \mathrm{Att}(b)$, then $\sigma_Q(a) = 1 - \sigma_Q(b)$.*

Every elementary modular semantics such that ι and α satisfy open mindedness with the same range (l, u) satisfies

Open-Mindedness: *Let $\mathcal{A}_0 = \mathcal{A}$, $\mathrm{Att}_0 = \mathrm{Att}$, $\mathrm{Sup}_0 = \mathrm{Sup}$, $\beta_0 = \beta$. For $i \geq 1$, let $\mathcal{A}_i = \mathcal{A}_{i-1} \cup \{a_i\}$ for some new argument $a_i \notin \mathcal{A}_{i-1}$ and $\beta_i(a) = \beta_{i-1}(a)$ for all $a \in \mathcal{A}_{i-1}$ and $\beta_i(a_i) = 1$. For all $a \in \mathcal{A}$, let Let $E^a_0 = \emptyset$ and $E^a_i = E^a_{i-1} \cup \{(a_i, a)\}$. If, for all $a \in \mathcal{A}$ such that the strength of all QBAFs $Q^{a,-}_i = (\mathcal{A}_i, \mathrm{Att}_i \cup E_i, \mathrm{Sup}_i, \beta_i)$ and $Q^{a,+}_i = (\mathcal{A}_i, \mathrm{Att}_i, \mathrm{Sup}_i \cup E_i, \beta_i)$ is well-defined, then $\lim_{i \rightarrow \infty} \sigma_{Q^{a,-}_i}(a) = 0$ and $\lim_{i \rightarrow \infty} \sigma_{Q^{a,+}_i}(a) = 1$.*

Proof. Strict Monotonicity: The proofs are analogous to the proofs of the non-strict properties. We demonstrate this for A-Monotonicity: Since, σ_Q is a fixed-point of u_Q, we have $\sigma_Q(a) = \iota(B_a, \alpha(A_a, S_a))$. $A_a \succ S_a$ implies $\alpha(A_a, S_a) < 0$ by strict monotonicity of α. Hence, $\iota(B_a, \alpha(A_a, S_a)) < B_a$ by strict monotonicity of ι.

Duality: Note that $A_x \cong_Q S_y$ implies $Core(A_x) = Core(S_y)$. Since α satisfies neutrality, we can assume without loss of generality that $A_x = Core(A_x) = Core(S_y) = S_y$ in this case. Hence, $\alpha(A_a, S_a) = -\alpha(S_a, A_a) = -\alpha(A_b, S_b)$ by duality of α and $1 - \iota(B_a, \alpha(A_a, S_a)) = \iota(1 - B_a, -\alpha(A_a, S_a)) = \iota(1 - B_b, \alpha(A_b, S_b))$ by duality of ι. From this, the claim follows because σ_Q is a fixed-point of u_Q.

Open-Mindedness: For all new arguments a_j in all QBAFs, we have $\mathrm{Att}(a_j) = \mathrm{Sup}(a_j) = \emptyset$. Hence, stability implies $\sigma_{Q_i^{a,-}}(a_j) = \sigma_{Q_i^{a,+}}(a_j) = \beta_i(a_j) = 1$ for all $i \geq j$. Let us focus on $Q_i^{a,-}$ and just write Q_i for ease of notation. Since, σ_{Q_i} is a fixed-point of the update function, we have $\sigma_{Q_i}(a) = \iota(\beta_i(a), \alpha(\sigma_{Q_i}(\mathrm{Att}_i(a)), \sigma_{Q_i}(\mathrm{Sup}_i(a)))$. By construction $\sigma_{Q_i}(\mathrm{Att}_i(a))$ contains at least i 1s. Hence, open-mindedness of α implies that $\lim_{i\to\infty} \alpha(\sigma_{Q_i}(\mathrm{Att}_i(a)), \sigma_{Q_i}(\mathrm{Sup}_i(a)) = l$. Then open-mindedness of ι implies that $\lim_{i\to\infty} \sigma_{Q_i}(a) = 0$. The proof for $Q_i^{a,+}$ is analogous. □

Intuitively, duality states that attackers and supporters behave symmetrically. For example, if the attackers and supporters increase the base score 0.7 of an argument by 0.2 to 0.9, we should expect that when we switch the roles of attackers and supporters and change the base score to $1-0.7 = 0.3$, it should decrease by 0.2 to 0.1. Open-mindedness states that if we keep adding undefeated strong (base score 1) attackers/supporters of an argument, its strength should eventually go to 0/1.

5 Conservativeness and Open-Mindedness

5.1 Absolute Conservativeness

As discussed in the introduction, in applications like recommender systems or truth discovery, we may want both open-mindedness and some notion of conservativeness. At the moment, we can satisfy open-mindedness only with α_Σ. However, since α_Σ is unbounded, even arguments with marginal strength, say 0.01 instead of 1, can significantly impact other arguments when they occur in large numbers. We want to be able to bound the minimal/maximal achievable strength by the base score and the maximal strength of attackers/supporters. As a first attempt to axiomatize this idea, we may ask that an argument's strength is bounded from below/above by its initial weight and the strength of its attackers/supporters as follows.

Definition 23 (Absolute Conservativeness)**.** *A QBAF semantics satisfies* absolute conservativeness *if for every QBAF Q such that σ_Q is well-defined, for all arguments $a \in \mathcal{A}$, we have* $\min\{\beta(a), 1 - \max_{b\in\mathrm{Att}(a)} \sigma_Q(b)\} \leq \sigma_Q(a) \leq \max\{\beta(a), \max_{b\in\mathrm{Sup}(a)} \sigma_Q(b)\}$.

Intuitively, absolute conservativeness demands that the negative/positive deviation of the final strength from the base score cannot be larger than the maximum strength of an attacker/supporter. We are not aware of any existing semantics that satisfies this property. To see that none of the semantics that we introduced earlier satisfy absolute conservativeness, think of a QBAF with two arguments a, b both with base score 0.5 and a support edge from a to b. Then a's final strength must be 0.5 by individual balance. Hence, all aggregation functions will return a positive aggregate for b and all influence functions will increase b's base score to a strength strictly greater than 0.5. Hence, b will violate the upper bound $\max\{\beta(b), \max_{c\in\mathrm{Sup}(b)} \sigma_Q(c)\} = 0.5$.

There are many trivial ways to satisfy absolute conservativeness. For example, we could use an aggregation function that always returns 0 or an influence function that always returns the initial weight. A more natural solution is to use $\alpha_\top$ (which guarantees that only the maximal values are considered) and to define a new influence function guided by the bound that we want to satisfy. We call this influence function *AC*.

$$\iota_{ac}(b, a) = \begin{cases} \min\{b, 1 + a\} & \textit{if } a < 0, \\ \max\{b, a\} & \textit{else.} \end{cases} \tag{AC}$$

ι_{ac} satisfies the following properties.

Lemma 24. *ι_{ac} is Lipschitz-continuous and satisfies monotonicity, balance, duality and open-mindedness for range $(-1, 1)$.*

Proof. ι_{ac} is Lipschitz-continuous with Lipschitz constant 1 because it is a piece-wise linear function composed of linear functions whose slope is bounded by 1.

Monotonicity and balance follow from elementary arithmetic. For duality, assume that $a > 0$ and note that $\iota_{ac}(1 - b, -a) = \min\{1-b, 1-a\} = 1-\max\{b, a\} = 1-\iota_{ac}(b, a)$. The case $a \leq 0$ is similar. Open-mindedness for range $(-1, 1)$ follows immediately from the definition. □

ι_{ac} does not satisfy strict monotonicity because a change in a will be ignored when it is not sufficiently large compared to $\beta(a)$ and vice versa. We can combine ι_{ac} with $\alpha_\top$ to satisfy absolute conservativeness and call the corresponding semantics *TAC*.

Definition 25 (TAC Semantics)**.** *The* TAC Semantics *is the modular semantics defined by the kernel $(\alpha_\top, \iota_{ac})$.*

Theorem 26. *The TAC semantics is an elementary modular semantics. For every QBAF Q such that σ_Q is well-defined, it satisfies all properties from Theorem 20, duality and absolute conservativeness. It does not satisfy the strict monotonicity and open-mindedness properties.*

Proof. Everything but absolute conservativeness follows from the properties of $(\alpha_\top, \iota_{ac})$ and Theorems 20 and 22. For absolute conservativeness, since σ_Q is a fixed-point of the update function, we have $\sigma_Q(a) = \iota_{ac}(B_a, \alpha_\top(A_a, S_a))$. By definition of $\alpha_\top$, $-\max_{s\in A_a} s \leq \alpha_\top(A_a, S_a) \leq \max_{s\in S_a} s$. Hence, the definition of ι_{ac} implies that $\min\{B_a, 1 - \max_{s\in A_a} s\} \leq \sigma_Q(a) \leq \max\{B_a, \max_{s\in S_a} s\}$. □

The *TAC semantics* does not satisfy the strict monotonicity properties for the same reasons that we had for top and product aggregation before. Once there is an attacker/supporter with strength 1, the maximum impact is obtained and additional attackers/supporters cannot bring any additional negative/positive effect. The lack of open-mindedness is caused by the top aggregation function that cannot take $-1/1$ anymore if there is a single supporter/attacker with non-zero strength.

Overall, the *TAC semantics* is a reasonable and non-trivial choice when absolute conservativeness is desirable. However, it may be more conservative than we want.

Proposition 27. *When combining ι_{ac} and any aggregation function α such that $-maxA \leq \alpha(A,S) \leq maxS$, then for every QBAF Q such that σ_Q is well-defined, for all $a \in \mathcal{A}_Q$, we have $\min\{\beta(a), 1-B\} \leq \sigma_Q(a) \leq \max\{\beta(a), B\}$ where B is the maximum base score of a predecessor of a.*

Proof. All strength values are initialized with the base scores, so the claim is initially true. When performing a discrete update, the new strength of every argument is bounded from above by the argument's base score or the maximum strength of its parents, so it remains bounded by the maximum base score of its predecessors by the assumption about α. We can check similarly that the lower bound remains intact. For continuous updates, consider the solution s^* of equation (2). Again, it respects the bounds initially by the initial condition of (2). If, for some argument $i \in \mathcal{A}$, $s(t)_i = B$ at some time t, (2) implies $\frac{ds_i}{dt}(t) = \iota(\beta(i), \alpha(\{s(t)_i \mid i \in \mathrm{Att}(a)\}, \{s(t)_i \mid i \in \mathrm{Sup}(a)\}) - s_i(t) \leq B - B = 0$, that is, $s(t)_i$ is non-increasing and thus the bounds remains intact. Similarly, we can check that the lower bound remains intact. □

For example, if all base scores are bounded from above by 0.5 as in our truth discovery example, then the final strength of arguments cannot deviate by more than 0.5 from their base score. In our truth discovery example, this means that claims cannot obtain a strength value larger than 0.5.

5.2 Relative Conservativeness

In many applications, a slightly weaker notion of conservativeness seems sufficient that aligns better better with existing influence functions in the literature. Instead of demanding that the initial weight and the strength of parents define absolute bounds, we allow relative deviations from the base score based on the evidence presented by the pro and contra arguments.

Definition 28 (Relative Conservativeness)**.** *A QBAF semantics satisfies* relative conservativeness *if there is a strictly monotonically increasing function $f: [0,1] \to [0,1]$ with $f(0) = 0$ and $f(1) = 1$ such that for every QBAF Q such that σ_Q is well-defined, for all arguments $a \in \mathcal{A}$, we have $\beta(a) - f(\max_{b \in \mathrm{Att}(a)} \sigma_Q(b)) \cdot \beta(a) \leq \sigma_Q(a) \leq \beta(a) + f(\max_{b \in \mathrm{Sup}(a)} \sigma_Q(b)) \cdot (1 - \beta(a))$.*

Intuitively, relative conservativeness demands that the maximum strength of attackers/supporters bounds how close the final strength can come to 0/1 relative to the base score (only an attacker/supporter with strength 1 allows moving the final strength all the way to 0/1 provided that the base score was not already 0/1). The function f determines how fast this can happen. For example, f could just be the linear function $f(x) = x$. In this case, the condition becomes just $\beta(a) - (\max_{b \in \mathrm{Att}(a)} \sigma_Q(b)) \cdot \beta(a) \leq \sigma_Q(a) \leq \beta(a) + (\max_{b \in \mathrm{Sup}(a)} \sigma_Q(b)) \cdot (1 - \beta(a))$. We will connect relative conservativeness to properties of aggregation and influence functions again.

Definition 29. (i). *An aggregation function α satisfies* relative conservativeness *if $-\max A \leq \alpha(A,S) \leq \max S$.* (ii). *An influence function ι satisfies* relative conservativeness *if there is a strictly monotonically increasing function $f: [0,1] \to [0,1]$ with $f(0) = 0$ and $f(1) = 1$ such that $b - f(-\min\{0,a\}) \cdot b \leq \iota(b,a) \leq b + f(\max\{0,a\}) \cdot (1-b)$.*

Proposition 30. *If α and ι satisfy relative conservativeness, then the associated modular semantics satisfies relative conservativeness for every QBAF Q such that σ_Q is well-defined.*

Proof. Since σ_Q is a fixed-point of the update function, we have $\sigma_Q(a) = \iota(\beta(a), \alpha(\sigma_Q(\mathrm{Att}(a)), \sigma_Q(\mathrm{Sup}(a))))$. Since α satisfies conservativeness, we have $-\max \sigma_Q(\mathrm{Att}(a)) \leq \alpha(A,S) \leq \max \sigma_Q(\mathrm{Sup}(a))$. Hence, relative conservativeness of ι implies the claim. □

Lemma 31. *$\alpha_\top, \alpha_\Pi$ and ι_d satisfy relative conservativeness*

Proof. For $\alpha_\top$, the claim follows immediately from the definition. For α_Π, note that $\alpha_\Pi(A,S) = \prod_{a \in A}(1-a) - \prod_{s \in S}(1-s) \leq 1 - (1 - \max S) = \max S$ and similar for the lower bound. ι_d satisfies relative conservativeness for the identity function $f(x) = x$. □

Our previous results imply that Df-Quad and the semantics defined by $(\alpha_\top, \iota_d)$ satisfy relative conservativeness.

Corollary 32. *The modular semantics defined by (α_Π, ι_d) (Df-QuAD) and $(\alpha_\top, \iota_d)$ satisfy relative conservativeness.*

Let us note that our previous results imply that $(\alpha_\top, \iota_d)$ is another elementary modular semantics that satisfies duality and therefore has properties very similar to Df-QuAD. We can now also clarify the relationship between absolute and relative conservativeness.

Proposition 33. *Absolute conservativeness implies relative conservativeness but not vice versa.*

Proof. Suppose that a semantics satisfies absolute conservativeness. We show that it satisfies the lower and upper bound defined by relative conservativness for the identity function $f(x) = x$. For the lower bound, note that $\beta(a) - f(\max_{b \in \mathrm{Att}(a)} \sigma_Q(b)) \cdot \beta(a) = \beta(a) \cdot (1 - f(\max_{b \in \mathrm{Att}(a)} \sigma_Q(b))) \leq \min\{\beta(a), (1 - f(\max_{b \in \mathrm{Att}(a)} \sigma_Q(b)))\}$ since both factors are bounded by 1. For the upper bound, note that $\beta(a) + f(\max_{b \in \mathrm{Sup}(a)} \sigma_Q(b)) \cdot (1 - \beta(a)) \geq \beta(a)$ because the second term is non-negative and $\beta(a) + f(\max_{b \in \mathrm{Sup}(a)} \sigma_Q(b)) \cdot (1 - \beta(a)) = \beta(a) \cdot (1 - f(\max_{b \in \mathrm{Sup}(a)} \sigma_Q(b))) + f(\max_{b \in \mathrm{Sup}(a)} \sigma_Q(b)) \geq f(\max_{b \in \mathrm{Sup}(a)} \sigma_Q(b))$. because the first term is non-negative. Hence, absolute conservativeness implies relative conservativeness.

Since Df-QuAD satisfies relative but not absolute conservativeness, the other direction cannot hold. □

An immediate consequence is that the previously introduced TAC semantics also satisfies relative conservativeness.

Corollary 34. *The TAC semantics satisfies relative conservativeness.*

5.3 Conservativeness and Open-Mindedness

While we saw a number of reasonable semantics that satisfy notions of conservativeness, none of them satisfies open-mindedness. The problem of the existing semantics is the aggregation function that is based on a difference $P - S$ of positive and negative influences. Both top and product aggregation share the desirable property that the aggregate is bounded from below by the negative maximum strength of the attackers and from above by the maximum strength of supporters. However, if both positive and negative influences exist, the aggregate will be bounded away from the minimum/maximum by the minimal opposite value. While this can be desirable when we only want to take the maximal strength of arguments into account, it is undesirable when we also want to take the number of arguments into account. In order to combine the advantages of sum-aggregation (an overwhelming majority of arguments can influence the final outcome to an arbitrary degree) and product- and top-aggregation (the maximum degree is bounded by the maximum strength of the involved arguments), we propose a new aggregation function that we call *convex-max (CM)*:

$$\alpha_c(A, S) = \lambda \cdot \max S - (1 - \lambda) \cdot \max A, \quad \text{(CM)}$$

where $\lambda = \frac{\sum_{s \in S} s}{\sum_{s \in S} s + \sum_{a \in A} a}$ and we let $\max \emptyset = 0$ and $\frac{0}{0} = 1$. Formally, $\alpha_c(A, S)$ is a convex combination of $\max S$ and $-\max A$. Intuitively, this means that it will always return a value between $\max S$ and $-\max A$. How close it is to either of them depends on the coefficient λ. For $\lambda = 1$, it will take $\max S$, for $\lambda = 0$, it will take $-\max A$. Values between 0 and 1 vary between these extremes.

Lemma 35. *α_c is Lipschitz-continuous and satisfies monotonicity, balance, duality, relative conservativeness and is open-minded for range $(-1, 1)$.*

Proof. For Lipschitz-continuity, first note that the maximum is Lipschitz-continuous because it is a piecewise linear function with maximum slope 1. To see that the coeficients are Lipschitz-continuous, we use the fact that a partially differentiable function with bounded partial derivatives is Lipschitz-continuous. By applying the quotient rule, we can see that the partial derivaties are bounded by 1. Hence, each product of coefficent and maximum is Lipschitz-continuous with constant 1 and their sum is Lipschitz-continuous with constant 2.

Monotonicity and balance follow from elementary arithmetic. For duality, note that $1 - \lambda = \frac{\sum_{s \in S} s + \sum_{a \in A} a}{\sum_{s \in S} s + \sum_{a \in A} a} - \frac{\sum_{s \in S} s}{\sum_{s \in S} s + \sum_{a \in A} a} = \frac{\sum_{a \in A} a}{\sum_{s \in S} s + \sum_{a \in A} a}$. Hence, $\alpha_c(A, S) = \lambda \cdot \max S - (1 - \lambda) \cdot \max A = -((1 - \lambda) \cdot \max A - \lambda \cdot \max S) = -\alpha_c(S, A)$. Relative conservativeness follows immediately from the fact that $\alpha_c(A, S)$ is a convex combination of $\max S$ and $-\max A$. For open-mindedness, note that for $\lim_{n \to \infty} \alpha_c(A \sqcup E_n, S)$, we have $\lambda \to 0$ and thus $\alpha_c(A \sqcup E_n, S) \to -\max(A \sqcup E_n) = -1$. The second limit can be checked similarly. □

	DfQ	EB	QE	MLP	*TAC*	*CRC*
(S)IAM	✓(✗)	✓(✗)	✓(✓)	✓(✗)	✓(✗)	✓(✗)
(S)ISM	✓(✗)	✓(✗)	✓(✓)	✓(✗)	✓(✗)	✓(✗)
(S)RM	✓(✗)	✓(✓)	✓(✓)	✓(✓)	✓(✗)	✓(✗)
IB	✓	✓	✓	✓	✓	✓
RB	✓	✓	✓	✓	✓	✓
Du	✓	✗	✓	✓	✓	✓
OM	✗	✗	✓	✗	✗	✓
AC	✗	✗	✗	✗	✓	✗
RC	✓	✗	✗	✗	✓	✓

Figure 4: Overview of which semantics satisfy (strict) individual A-/S- and relative- monotonicity, individual and relative balance, duality, open-mindedness, absolute and relative conservativeness. DfQ, EB, QE, MLP refer to Df-QuAD (Rago et al. 2016), Euler-based (Amgoud and Ben-Naim 2017), quadratic energy (Potyka 2018a) and MLP-based (Potyka 2021) semantics, respectively.

Hence, our previous results imply that the modular semantics defined by (α_c, ι_l) that we call the *CRC semantics* satisfies both open-mindedness and relative conservativeness.

Definition 36 (CRC Semantics)**.** *The* CRC Semantics *is the modular semantics defined by the kernel* (α_c, ι_l).

Theorem 37. *The CRC semantics is an elementary modular semantics. For every QBAF Q such that σ_Q is well-defined, it satisfies all properties from Theorem 20, duality, open-mindedness and relative conservativeness. It does not satisfy absolute conservativeness and strict monotonicity.*

To see that the CRC semantics does not satisfy absolute conservativeness, just consider a QBAF with two arguments with base score 0.5 and a support from one to the other. Then the supported argument has strength 1, which violates absolute conservativeness. For strict monotonicity, note that if an argument a has no attackers and two supporters with strength 0.5 and 0.4, increasing the strength of the weaker argument from 0.4 to 0.41 does not increase a's strength (which is already 0.5). Similarly adding another supporter with strength smaller than 0.5 will not change a's strength.

6 Discussion and Future Work

We studied two notions of conservativeness that allow bounding the strength of arguments based on the strength of their parents. Absolute conservativeness (AC) is a natural notion but may be too restrictive. The weaker notion of relative conservativeness (RC) may be sufficient for many applications. AC cannot be satisfied by existing modular semantics, but can be satisfied by our new TAC semantics. As we saw, Df-QuAD already satisfies relative conservativeness, but does not satisfy open-mindedness. Our new CRC semantics satisfies both and most existing properties from the literature (except the strict ones), but does not satify AC. Figure 4 summarizes the semantical properties of the most important semantics discussed in this paper. We do not include the anonymity, independence and directionality properties because they are satisfied by every modular semantics independent of the choice of α and ι (Theorem 20). Let us note that we could satisfy open-mindedness and absolute conservativeness simultaneously by combining the AC-influence function ι_{ac} and the CM-aggregation function α_c. However, since the induced modular semantics will suffer from the limitations stated in Proposition 27 we did not discuss it here in the interest of space. Implementations of all modular semantics that we discussed here can be found in the Java library *Attractor*[1] (Potyka 2018b; Potyka 2022).

We spent some time on identifying relevant properties of aggregation and influence functions. The results between the relationship of properties of aggregation and influence functions and the induced modular semantics simplified the investigation of our new semantics significantly and we hope that they can further simplify the study of future modular semantics. As we discussed in the introduction, some of these results can be seen as generalizations of results from (Amgoud and Doder 2018) from the attack-only setting to the bipolar setting. For example, our elementary kernels can be seen as a stricter form of *well-behaved evaluation methods* defined in (Amgoud and Doder 2018, Def. 6) (note that the two conditions in (Amgoud and Doder 2018, Def. 6) can be seen as the attack-only counterpart of the functional counterpart of balance and monotonicity). Our results about anonymity, independence and directionality in Theorem 20 can be seen as a generalization of (Amgoud and Doder 2018, Prop. 2) to the bipolar setting. There is also a slight generalization in the attack-only setting because (Amgoud and Doder 2018, Prop. 2) assumes a unique fixed-point, while we assume only convergence to any fixed-point (well-definedness). The second part of Theorem 20 can probably be seen as a generalization of (Amgoud and Doder 2018, Prop. 3-5) to the bipolar setting. Stating the exact relationship is difficult because the properties in (Amgoud and Doder 2018, Prop. 3-5) are special cases of balance and monotonicity (Baroni, Rago, and Toni 2018) and no proofs are provided in (Amgoud and Doder 2018).

In future work, we will study to which extent different QBAF encodings under different semantics can satisfy the postulates for truth discovery proposed in (Singleton and Booth 2022) and evaluate their performance empirically building up on benchmarks developed in (Elsaesser, Everaere, and Konieczny 2023) and (Potyka and Booth 2024). It would also be interesting to explore how QBAFs compare to probabilistic (Hunter et al. 2021) or combined approaches (Spaans and Doder 2023) in this setting.

[1] https://sourceforge.net/projects/attractorproject/

References

Albini, E.; Baroni, P.; Rago, A.; and Toni, F. 2020. Pagerank as an argumentation semantics. In *COMMA 2020*, 55–66. IOS Press.

Amgoud, L., and Ben-Naim, J. 2017. Evaluation of arguments in weighted bipolar graphs. In *ECSQARU 2017*, 25–35. Springer.

Amgoud, L., and Doder, D. 2018. Gradual semantics for weighted graphs: An unifying approach. In *KR 2018*, 613–614. AAAI Press.

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

Ayoobi, H.; Potyka, N.; and Toni, F. 2023. Sparx: Sparse argumentative explanations for neural networks. In *ECAI 2023*, 149–156. IOS Press.

Baroni, P.; Rago, A.; and Toni, F. 2018. How many properties do we need for gradual argumentation? In *AAAI 2018*, 1736–1743. AAAI Press.

Berti-Equille, L., and Borge-Holthoefer, J. 2015. Veracity of data: From truth discovery computation algorithms to models of misinformation dynamics. *Synthesis Lectures on Data Management* 7(3):1–155.

Cocarascu, O.; Rago, A.; and Toni, F. 2019. Extracting dialogical explanations for review aggregations with argumentative dialogical agents. In *AAMAS 2019*, 1261–1269.

Dung, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial intelligence* 77(2):321–357. Publisher: Elsevier.

Elsaesser, Q.; Everaere, P.; and Konieczny, S. 2023. S&f: Sources and facts reliability evaluation method. In *AAMAS 2023*, 2778–2780. ACM.

Gupta, M., and Han, J. 2011. Heterogeneous network-based trust analysis: A survey. *SIGKDD Explorations Newsletter* 13(1):54–71.

Hunter, A.; Polberg, S.; Potyka, N.; Rienstra, T.; and Thimm, M. 2021. Probabilistic argumentation: A survey. *Handbook of Formal Argumentation* 2.

Kotonya, N., and Toni, F. 2019. Gradual argumentation evaluation for stance aggregation in automated fake news detection. In *Workshop on Argument Mining*, 156–166.

Li, Y.; Gao, J.; Meng, C.; Li, Q.; Su, L.; Zhao, B.; Fan, W.; and Han, J. 2016. A Survey on Truth Discovery. *SIGKDD Explorations Newsletter* 17(2):1–16.

Mossakowski, T., and Neuhaus, F. 2018. Modular semantics and characteristics for bipolar weighted argumentation graphs. *arXiv preprint arXiv:1807.06685*.

Potyka, N., and Booth, R. 2024. An empirical study of the behaviour of quantitative bipolar argumentation frameworksfor truth discovery. In *COMMA 2024*, To appear. IOS Press.

Potyka, N. 2018a. Continuous dynamical systems for weighted bipolar argumentation. In *KR 2018*, 148–157.

Potyka, N. 2018b. A tutorial for weighted bipolar argumentation with continuous dynamical systems and the java library attractor. NMR.

Potyka, N. 2019a. Extending Modular Semantics for Bipolar Weighted Argumentation. In *AAMAS 2019*, 1722–1730. IFAAMAS.

Potyka, N. 2019b. Open-mindedness of gradual argumentation semantics. In *SUM 2019*, volume 11940, 236–249. Springer.

Potyka, N. 2021. Interpreting neural networks as quantitative argumentation frameworks. In *AAAI 2021*, 6463–6470.

Potyka, N. 2022. Attractor - A java library for gradual bipolar argumentation. In *COMMA 2022*, volume 353 of *Frontiers in Artificial Intelligence and Applications*, 369–370. IOS Press.

Rago, A.; Toni, F.; Aurisicchio, M.; and Baroni, P. 2016. Discontinuity-free decision support with quantitative argumentation debates. In *KR 2016*, 63–73.

Rago, A.; Cocarascu, O.; and Toni, F. 2018. Argumentation-based recommendations: Fantastic explanations and how to find them. In *IJCAI 2018*, 1949–1955.

Singleton, J., and Booth, R. 2022. Towards an axiomatic approach to truth discovery. *Autonomous Agents and Multi-Agent Systems* 36(2):1–49.

Singleton, J. 2020. On the link between truth discovery and bipolar abstract argumentation. *Online Handbook of Argumentation for AI* 43–46.

Spaans, J. P., and Doder, D. 2023. Graduality in probabilistic argumentation frameworks. In *ECAI 2023*, volume 372, 2186–2193. IOS Press.