Proceedings of the ESSLLI 2022 Student Session, 128–140
doi.org/10.21942/uva.20370810

Subjective Logic as a Model for Social Networks

José C. Oliveira

University of Minnesota Duluth, Department of Computer Science
josecarlosdeoliveirajr@gmail.com

Abstract. Our goal is to use *subjective logic* (SL), a quantitative logic with uncertainty, to model agents' opinions and uncertainty, and their changes over time, in social networks. In this paper, we discuss the desired properties for an opinion update function in this setting. We show that SL's predefined belief update functions do not satisfy these properties, and we define a new belief update function satisfying the desired properties. We show that a special case of opinions in SL with our new update function corresponds to earlier (non-logical) work on social networks [1], and that the inclusion of uncertainty strictly extends this earlier work.

1 Introduction

Recently, social networks have begun to influence every aspect of our lives, with unprecedented, unanticipated consequences, especially in politics and public opinion. Research on social networks has studied opinions and their change over time, but to accurately model real people and their opinions and beliefs, we must include information about *uncertainty*, or confidence, in formal models of social networks.

To achieve this goal, we use *subjective logic* (SL), which includes information about agents' uncertainty, to develop a more nuanced model of social networks and their changes over time. More precisely, the contributions of this paper are the following:

- We identify four desiderata of properties for an update function with SL. Those desiderata are our interpretation of rationality when updating opinions.
- We propose an opinion update function using SL's trust discount and belief fusion. We show through examples that our update function, with cumulative, averaging, and weighted belief fusions, does not satisfy our desiderata.
- We propose a new update function using an earlier (non-logical) work on social networks. We show that the new update function strictly extends this earlier work and satisfies our desiderata.

2 Related work

Research about the dynamics of opinions in social networks using formal models and logic-based approaches is fairly new, and much of it uses non-quantitative (yes or no) opinions or beliefs, rather than qualitative opinions which can take on a spectrum of values between 0 and 1. Sîrbu et al. [9] investigate the effects of algorithmic bias on polarization by counting the number of opinions clusters. Gargiulo et al. [4] develop simulated social network and observe group formation over time. Liu et al. [8] use ideas from doxastic and dynamic epistemic logics to qualitatively model influence and belief changes in social networks. Christoff [3] develops several non-quantitative logics for social networks, and Young Pedersen [10] develops a non-quantitative logic concerned specifically with polarization. Hunter [5] introduces a logic of belief updates over social networks with varying levels of influence and trust. Using dynamic epistemic logic, Baltag et al. [2] created a threshold model where agents' behavior changes when the proportion of supporters changes.

Alvim et al. [1] develop a formal model for social networks where agents have quantitative opinions and quantitative influence on each other, with a function for agents' belief update over time. The goal of the current paper is to extend this model by adding the possibility of *uncertainty* to the agents' quantitative opinions.

3 Background: Subjective Logic

This section describes all the elements of subjective logic that we use in our model. *Subjective Logic* is a logic developed by Josang [6] that extends probabilistic logic by adding *uncertainty* and *subjectivity*. In probabilistic logic, a uniform distribution does not express *"we don't know"*, because a uniform distribution says that we know that the distribution over the domain is uniform. Subjective logic can express *uncertainty* by the lack of confidence about the distribution. The *subjectivity* comes from the fact that we can assign an opinion about a proposition to an agent.

The main object of subjective logic is the *opinion*. We represent an opinion by ω_X^A, where A is an agent, X a random variable, and ω_X^A is A's opinion about X. An opinion expresses support for none, one or many states of a domain. This section presents the elementary definitions composing an opinion. A *domain* is a state space consisting of a set of values called states, events, outcomes, hypotheses or propositions. The values are assumed to be exclusive and exhaustive.

Belief mass is a distribution over a domain representing an agent's confidence in each value in the domain. The belief mass assigned to a value $x \in \mathbb{X}$ expresses support for x being TRUE. Belief mass is sub-additive, i.e. $\sum\limits_{x\in\mathbb{X}} \mathbf{b}_X(x) \leq 1$. The sub-additivity is complemented by *uncertainty mass* u_X and it represents the lack of support or evidence for the variable X having any specific value.

Definition 1. (Belief Mass Distribution) *Let $\mathbb{X}$ be a domain of size $k \geq 2$, and let X be a variable over that domain. A belief mass distribution denoted $\mathbf{b}_X : \mathbb{X} \to [0, 1]$ assigns belief mass to possible values of the variable X. Belief mass and uncertainty mass sum to one, i.e., $u_X + \sum_{x \in \mathbb{X}} \mathbf{b}_X(x) = 1$.*

Opinions can be semantically different, depending on the situation they apply to. An *aleatory* opinion applies to a variable governed by a frequentist process, and representing the likelihood of values of the variable in any unknown past or future instance of the process. "*The bias of a coin is $p = 0.6$*" is an aleatory opinion. An *epistemic* opinion applies to a variable that is assumed to be non-frequentist, and that represents the likelihood of the variables in a specific unknown instance. "*Beatriz killed Evandro*" is an epistemic opinion. In an epistemic opinion, opposite/different pieces of evidence should cancel each other out. Therefore, it must be uncertainty-maximized.

Base rate distribution represents a *prior* probability distribution over a domain: the probability distribution before considering evidence about the domain.

Definition 2. (Opinion) *Let $\mathbb{X}$ be a domain of size $k \geq 2$, and X a random variable in $\mathbb{X}$. An opinion over the random variable X is the ordered triple $\omega_X = (\mathbf{b}_X, u_X, \mathbf{a}_X)$ where*

- *$\mathbf{b}_X$ is a belief mass distribution over X,*
- *u_X is the uncertainty mass which represents the vacuity of evidence,*
- *$\mathbf{a}_X$ is a base rate distribution (a probability distribution) over $\mathbb{X}$.*

The *projected probability distribution* of an opinion is the *posterior* probability distribution after updating the base rate distribution with the belief mass distribution. The more an opinion depends on the belief mass, the less it depends on the base rate. The projected probability distribution is defined by $\mathbf{P}_X(x) = \mathbf{b}_X(x) + \mathbf{a}_X(x)u_X, \ \forall x \in \mathbb{X}$

The definition of opinion is useful for our model since it is more expressive than the belief state of an agent about a proposition in [1], which is similarly an opinion with domain $\mathbb{X} = \{true, false\}$, with no uncertainty mass. The agent must commit all of their mass to the values of the domain with no uncertainty.

Example 1. Let $\mathbb{X} = \{x, \overline{x}\}$ be a domain where x is "global warming is happening" and $\overline{x}$ is "global warming is *not* happening". Let X be a random variable in $\mathbb{X}$. An opinion about X must be epistemic, because it is about a fact in the present instance that is true or false. Let the base rate be uniform. With no evidence, an agent A will hold the opinion $\omega_X^A = ((0, 0), 1, (0.5, 0.5))$ with $\mathbf{P}_X^A(x) = 0.5$, meaning that A is 50% sure that the global warming is happening, but their opinion is relying only on the base rate, with no evidence supporting either of the values.

After gathering evidence from newspapers, scientific studies and other people, A assigns a belief mass to x. If agent A agrees with x by 80%, A holds the opinion

$\omega^A_X = ((0.6, 0), 0.4, (0.5, 0.5))$ with $\mathbf{P}^A_X(x) = 0.8$, meaning that A is 80% sure that global warming is happening, and has evidence that corresponds to 60% of their mass. The uncertainty mass means that A is relying 40% on the base rate.

To model influence that one agent has over another, subjective logic has *trust opinion*, an opinion that an agent has about another agent as a good source of information.

Definition 3. (Trust opinion) *Let* $\mathbb{T}_B = \{t_B, \bar{t}_B\}$ *be a* trust domain, *where* t_B *means "B is a good source of information" and* $\bar{t}_B$ *means "B is not a good source of information". Then* $\omega^A_{t_B}$*, or* ω^A_B *for short, is the (trust) opinion that A has about the trustworthiness of B as a source of information.*

We use trust opinions to model an agent's updated opinion after communication: $\omega^{[A;B]}_X$ is a new opinion generated by taking belief mass ω^B_X proportional to the belief mass of the trust opinion ω^A_B. $\omega^{[A;B]}_X$ represents A's opinion about X after communicating with B, ω^A_B represents A's opinion about B's trustworthiness, and ω^B_X represents B's opinion about X. The operation is denoted $\omega^{[A;B]}_X = \omega^A_B \otimes \omega^B_X$.

Example 2. Let $\omega^A_B = ((1, 0), 0, \mathbf{a}^A_B)$ with $\mathbf{P}^A_B(t_B) = 1$ and $\omega^B_X = ((0.6, 0), 0.4, (0.5, 0.5))$ with $\mathbf{P}^B_X(x) = 0.8$. Here, A completely trusts B and B is 80% sure that x is true with 60% of their mass assigned to x.

$\mathbf{P}^A_B(t_B) = 1$, i.e. A completely trusts B. Then, A by trusting B (in short $[A; B]$) will hold the same opinion as B about X. Therefore, $\omega^{[A;B]}_X = \omega^B_X$. By the opinion that A has about X by trusting B, A is 60% sure that x is true with 80% of their mass assigned to x.

Example 3. Let $\omega^A_B = ((0.5, 0.5), 0, \mathbf{a}^A_B)$ with $\mathbf{P}^A_B(t_B) = 0.5$ and $\omega^B_X = ((0.8, 0), 0.2, (0.5, 0.5))$ with $\mathbf{P}^A_X(x) = 0.9$. Here, A trusts B by 50% and B is 80% sure that x is true with 60% of their mass assigned to x.

$\mathbf{P}^A_B(t_B) = 0.5$. Then, $[A; B]$ will hold 50% of the belief mass of each value from B. Therefore, $\omega^{[A;B]}_X = ((0.4, 0), 0.6, (0.5, 0.5))$ with $\mathbf{P}^{[A;B]}_X(x) = 0.7$. By the opinion that A has about X by trusting B, A is 70% sure that x is true with 40% of their mass assigned to x.

To model A's concurrent interactions with multiple other agents, we can use belief fusion [6, 7]. Belief fusion combines a set of opinions into a single opinion which then represents the opinion of the collection of sources. There is more than one possible definition for the belief fusion operator. They differ by their properties and applications. For our model, we consider the following operators from [6, 7]:

- *Cumulative belief fusion* ($\omega^{(A \diamond B)}_X = \omega^A_X \oplus \omega^B_X$): It is used when it is assumed that the amount of independent evidence increases by including more and

more sources. The idea is to sum the amount of evidence of the opinions. It is non-idempotent. E.g. a set of agents flips a coin a number of times and produce an opinion about the bias of the coin. An opinion produced by cumulative belief fusion represents all the experiments made by the agents.
- *Averaging belief fusion* ($\omega_X^{(A \underline{\diamond} B)} = \omega_X^A \underline{\oplus} \omega_X^B$): It is used when including more sources does not mean that more evidence is supporting the conclusion. The idea is to take the average of the amount of evidence of the opinions. It is idempotent, but it has no neutral element. E.g. After observing the court proceedings, each member of a jury produces an opinion about the same evidence. The verdict is the fusion between those opinions.
- *Weighted belief fusion* ($\omega_X^{(A \hat{\diamond} B)} = \omega_X^A \hat{\oplus} \omega_X^B$): It is used when we take the average of the amount of evidence of the opinions weighted by their lack of uncertainty. In particular, opinions with no belief mass are rejected. It is idempotent and it has a neutral element $u_X = 1$. E.g. a group of medical doctors needs to decide on a diagnosis. Each of them has an opinion, but some of them are more certain (assigned more belief mass) than others. Those opinions must have more weight than the others upon fusion.

4 Update function for two agents

Our goal is to model agents' opinions in social networks and to develop an update function for opinions after agents interact. This section describes our proposal for an update function for two agents with belief fusion and trust opinions.

Definition 4. (Desiderata for the properties of an update function) *Let* $\mathbb{X} = \{x, \overline{x}\}$ *be a domain,* X *a random variable in* $\mathbb{X}$*, and* A *and* B *agents holding opinions* $\omega_X^{A[t]}$ *and* $\omega_X^{B[t]}$ *at time* t*. Let* ω_B^A *be a trust opinion. We need an update function* $\omega_X^{A[t+1]} = f(\omega_X^{A[t]}, \omega_B^A, \omega_X^{B[t]})$ *satisfying the following properties:*

1. Weak convergence*:* $\mathbf{P}_X^{A[t]} \leq \mathbf{P}_X^{A[t+1]} \leq \mathbf{P}_X^{B[t]}$ *or* $\mathbf{P}_X^{A[t]} \geq \mathbf{P}_X^{A[t+1]} \geq \mathbf{P}_X^{B[t]}$*. A cannot move further from B upon update.*
2. Markovian behaviour*:* $\omega_X^{A[t+1]}$ *must not depend on the opinions from time less than* t*. Each step of the update must be independent.*
3. Idempotence*: If* $\omega_X^{A[t]} = \omega_X^{B[t]}$*, then* $\omega_X^{A[t]} = \omega_X^{A[t+1]}$*. No matter the trust, if A and B have the same opinion, A must not change their opinion.*
4. Non-increasing uncertainty*: we have two versions of this desideratum.*
   (a) Non-increasing group uncertainty: *The sum of the uncertainty of all agents must not increase over time.*
   (b) Non-increasing individual uncertainty: *Each agent's individual uncertainty never increases.*

These desiderata are our interpretation of rationality when updating opinions. The *weak convergence* property says that when an agent receives new information, the belief is updated to a new value closer to the new information they receive, and they do not become more extreme. The *Markovian behavior* property says that our current belief contains all relevant information, so past beliefs are not taken into account to compute a new one. The *idempotence* property says that, if the new information an agent receives is identical to their current belief, then the belief remains unchanged. The *non-increasing uncertainty* property says that at least some agents do not become totally uncertain as they receive new information and communicate. We aim for one of the two versions of this property because we hope to avoid the situation where all agents become totally uncertain in the limit.

The intuition behind our planned update function is to fuse agent A's current opinion with all the opinions that A can gather by trusting other agents. Define a *dogmatic opinion* as an opinion with no uncertainty, i.e. $u_X = 0$. For this update function, we are not considering situations with dogmatic opinions, because it means the agent has an infinite amount of evidence and the belief fusion operators remove non-dogmatic opinions when at least one is present.

Next, we investigate whether subjective logic's predefined update functions satisfy our desired properties.

Definition 5. (Update function for 2 agents with Belief Fusion and Trust) *Let $\omega_X^{A[t]}$ and $\omega_X^{B[t]}$ be non-dogmatic opinions. Let $\oplus$ be a belief fusion operator. Define the update function for $\omega_X^{A[t+1]}$ as $\omega_X^{A[t+1]} = \omega_X^{A[t]} \oplus (\omega_B^A \otimes \omega_X^{B[t]})$.*

We call $(\omega_B^A \otimes \omega_X^{B[t]})$ the opinion that A *will* learn by interacting with B. $\omega_X^{A[t+1]}$ is the opinion that A holds after merging their previous opinion ($\omega_X^{A[t]}$) with the opinion that A learned $(\omega_B^A \otimes \omega_X^{B[t]})$.

By definition, the update function for $\omega_X^{A[t+1]}$ is Markovian. Note that, if $\oplus$ is the averaging or weighted belief fusion operator, and $\omega_X^{A[t]} = \omega_B^A \otimes \omega_X^{B[t]}$, then $\omega_X^{A[t]} = \omega_X^{A[t+1]}$. This is different from our desired idempotence property, which does not depend on trust opinion.

Example 4. For brevity, we write the value of $\mathbf{P}_B^A(t_B)$ where it should be ω_B^A.

$$\begin{aligned}\omega_X^{A[t+1]} &= ((0,0),1,(0.5,0.5)) \oplus (0.5 \otimes ((0.8,0),0.2,(0.5,0.5))) \\ &= ((0,0),1,(0.5,0.5)) \oplus ((0.4,0),0.6,(0.5,0.5))\end{aligned} \tag{1}$$

where

$$\begin{array}{llll}\omega_X^{A[t]} & = ((0,0),1,(0.5,0.5)), & \mathbf{P}_X^{A[t]}(x) & = 0.5 \\ \omega_B^A & = ((0.5,0.5),0,\mathbf{a}_B^A), & \mathbf{P}_B^A(t_B) & = 0.5 \\ \omega_X^{B[t]} & = ((0.8,0),0.2,(0.5,0.5))) & \mathbf{P}_X^{A[t]}(x) & = 0.9\end{array} \tag{2}$$

Here, agent A has an opinion with only uncertainty mass; they are 50% sure that x is true, and trusts B by 50%. Because the trust opinion has no uncertainty mass, the base rate is not relevant. Agent B is 90% sure that x is true with 80% of their mass to x. The final result depends on the choice of belief fusion operator.

Example 5. Now consider these two agents updating their opinion over time.

$$\begin{array}{lll} \omega_X^{A[0]} = ((0.2, 0), 0.8, (0.5, 0.5)) & \mathbf{P}_X^{A[0]}(x) = 0.6 & \mathbf{P}_B^A(t_B) = 0.5 \\ \omega_X^{B[0]} = ((0.8, 0), 0.2, (0.5, 0.5)) & \mathbf{P}_X^{B[0]}(x) = 0.9 & \mathbf{P}_A^B(t_A) = 0.5 \end{array} \tag{3}$$

Here, agent A is 60% sure about x and B is 90% about x. Both trust each other by 50%. The evolution of $\mathbf{P}_X^A(x)$ and $\mathbf{P}_X^B(x)$ is shown in Fig. 1.

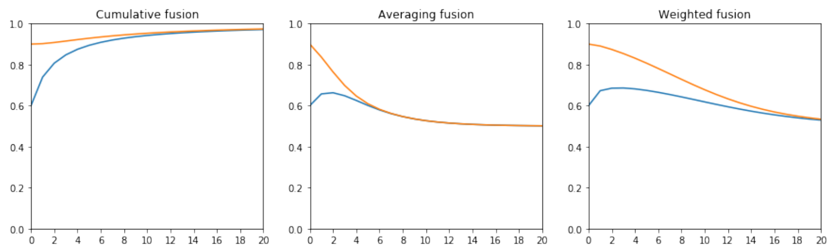

Fig. 1. $\mathbf{P}_X^A(x)$ (blue) and $\mathbf{P}_X^B(x)$ (orange) updated 20 times as in Example 5.

We expected that the update function would satisfy the weak convergence property by $\mathbf{P}_X^A(x)$ and $\mathbf{P}_X^B(x)$ converging to some value between $\mathbf{P}_X^{A[0]}(x) = 0.6$ and $\mathbf{P}_X^{B[0]}(x) = 0.9$. But because evidence keeps accumulating over time, with cumulative belief fusion, $\mathbf{P}_X^A(x)$ and $\mathbf{P}_X^B(x)$ converge to 1.

For averaging and weighted belief fusion, $\mathbf{P}_X^A(x)$ and $\mathbf{P}_X^B(x)$ converge to 0.5, violating weak convergence. With epistemic opinions, increasing uncertainty over time is expected, but the same happens with aleatory opinions.

Example 6. In this case, agent A trusts B by 0.5, but B does not trust A.

$$\begin{array}{lll} \omega_X^{A[0]} = ((0.2, 0), 0.8, (0.5, 0.5)) & \mathbf{P}_X^{A[0]}(x) = 0.6 & \mathbf{P}_B^A(t_B) = 0.5 \\ \omega_X^{B[0]} = ((0.8, 0), 0.2, (0.5, 0.5)) & \mathbf{P}_X^{B[0]}(x) = 0.9 & \mathbf{P}_A^B(t_A) = 0 \end{array} \tag{4}$$

Here, agent A is 60% and B 90% sure about x. A trusts B by 50%, but B does not trust A. The evolution of $\mathbf{P}_X^A(x)$ and $\mathbf{P}_X^B(x)$ is shown in Fig. 2.

For cumulative belief fusion, B does not change their opinion because $\omega_B^A \otimes \omega_X^{B[t]}$ has no belief mass, therefore, B cannot learn anything. However A keeps learning from B over time, and $\mathbf{P}_X^A(x)$ converges to 1.

The weighted belief fusion case shows how the idempotency works for its operator and how it is different from the idempotence property that we are aiming for. When $\omega_X^A = ((0.4, 0), 0.6, (0.5, 0.5))$ with $\mathbf{P}_X^A(x) = 0.7$, $\omega_X^A = \omega_B^A \otimes \omega_X^B$. Therefore, $\mathbf{P}_X^A(x)$ converges to 0.7.

Example 7. In this case, A and B have the same opinion and the same trust.

$$\begin{array}{lll} \omega_X^{A[0]} = ((0.6, 0), 0.4, (0.5, 0.5)) & \mathbf{P}_X^{A[0]}(x) = 0.8 & \mathbf{P}_B^A(t_B) = 0.5 \\ \omega_X^{B[0]} = ((0.6, 0), 0.4, (0.5, 0.5)) & \mathbf{P}_X^{B[0]}(x) = 0.8 & \mathbf{P}_A^B(t_A) = 0.5 \end{array} \tag{5}$$

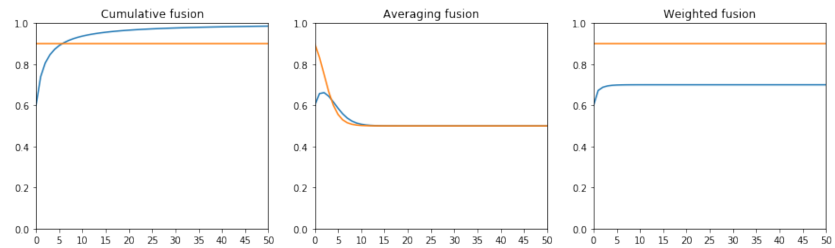

Fig. 2. $\mathbf{P}_X^A(x)$ (blue) and $\mathbf{P}_X^B(x)$ (orange) updated 20 times as in Example 6.

Here, A and B have the same opinion. They are 80% sure about x. Both trust each other by 50%. The evolution of $\mathbf{P}_X^A(x)$ and $\mathbf{P}_X^B(x)$ is shown in Fig. 3.

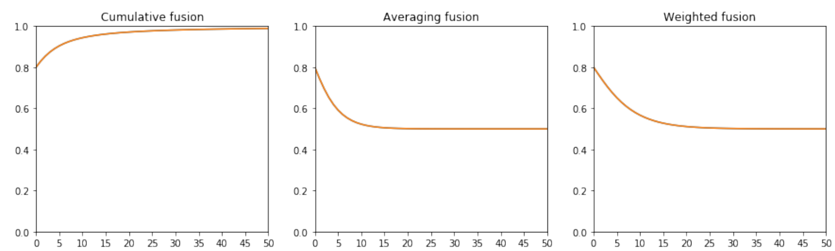

Fig. 3. $\mathbf{P}_X^A(x)$ and $\mathbf{P}_X^B(x)$ (both orange) updated 20 times as in Example 7.

Even starting with the same opinion, agents A and B do not keep the same opinion over time. With cumulative belief fusion, A and B keep accumulating evidence and $\mathbf{P}_X^A(x)$ and $\mathbf{P}_X^B(x)$ converge to 1. With averaging or weighted belief fusion, uncertainty keeps increasing, and $\mathbf{P}_X^A(x)$ and $\mathbf{P}_X^B(x)$ converge to 0.5.

To sum up, these examples show that an update function with trust discount and belief fusion does not satisfies our desiderata. In the next section, we propose a update function using an earlier work on social networks that satisfies our desiderata.

5 Alternative update function

In the previous section, we showed through some examples that our update function defined in Def. 5 does not satisfy the desired properties. This section shows an alternative update function with elements from [1].

One of our goals is to represent belief about multiple issues. We can update each state of the belief mass distribution using a definition similar to the classical belief update function from [1]. Let $\mathscr{A} = \{A_1, \cdots, A_n\}$ be a set of n agents. Let p be a proposition. Let $\mathrm{Bel}_p^t : A \to [0,1]$ be the belief state of an agent at the time t. Let $\mathrm{In} : \mathscr{A} \times \mathscr{A} \to [0,1]$ be the influence graph s.t. $\mathrm{In}(A_i, A_j)$ represents the influence of agent A_i on A_j, and $\mathrm{In}(A_i, A_j) = 1$ if $i = j$.

Definition 6. (Update function from [1]) *The belief of agent A_i given the interaction with all other agents in A at time t is defined as*

$$\mathrm{Bel}_p^{t+1}(A_i) = \frac{1}{n} \sum_{A_j \in \mathscr{A}} \left(\mathrm{Bel^t_p(A_i) + In(A_j, A_i)} \left(\mathrm{Bel^t_p(A_j) - Bel^t_p(A_i)}\right)\right) \tag{6}$$

In particular, the update function for two agents A and B is defined as

$$\mathrm{Bel}_p^{t+1}(A) = \frac{\mathrm{Bel}_p^t(A)}{2} + \frac{\mathrm{Bel}_p^t(A) + \mathrm{In}(B, A)(\mathrm{Bel}_p^t(B) - \mathrm{Bel}_p^t(A)}{2} \tag{7}$$

The idea behind this update function is to move the belief state of A to the belief state of B proportionally by the influence that B has on A. So, the new belief state is the average between the belief state that A has with the belief that A has when influenced by B. We define a new update function with subjective logic's opinion and the Def. 6

Definition 7. (Update Function for 2 agents with Constant Group Uncertainty) *Let $\mathbb{X}$ be a domain of size $k \geq 2$, X a random variable in $\mathbb{X}$, A and B agents holding opinions $\omega_X^{A[t]}$ and $\omega_X^{B[t]}$ at time t, and ω_B^A a trust opinion. Define the update function for $\omega_X^{A[t+1]}$ as*

$$\omega_X^{A[t+1]} = \begin{cases} \mathbf{b}_X^{A[t+1]} = \dfrac{\mathbf{b}_X^{A[t]}}{2} + \dfrac{\mathbf{b}_X^{A[t]} + \mathbf{P}_B^A(t_B)(\mathbf{b}_X^{B[t]} - \mathbf{b}_X^{A[t]})}{2} \\ u_X^{A[t+1]} = \dfrac{u_X^{A[t]}}{2} + \dfrac{u_X^{A[t]} + \mathbf{P}_B^A(t_A)(u_X^{B[t]} - u_X^{A[t]})}{2} \\ \mathbf{a}_X^{A[t+1]} = \mathbf{a}_X^{A[t]} \end{cases} \tag{8}$$

Example 8. Let $\mathbb{X} = \{x_1, x_2, x_3\}$ be the domain. Agents A and B have the following opinions and trust opinions:

$$\begin{array}{ll} \omega_X^{A[0]} = ((0, 0.4, 0.6), 0, \mathbf{a}_X^{A[0]}) & \mathbf{P}_B^A(t_B) = 0.5 \\ \omega_X^{B[0]} = ((0.4, 0.4, 0.2), 0, \mathbf{a}_X^{B[0]}) & \mathbf{P}_A^B(t_A) = 0.5 \end{array} \tag{9}$$

After one step, A and B have the following opinions:

$$\begin{array}{l} \omega_X^{A[0]} = ((0.1, 0.4, 0.5), 0, \mathbf{a}_X^{A[0]}) \\ \omega_X^{B[0]} = ((0.3, 0.4, 0.3), 0, \mathbf{a}_X^{B[0]}) \end{array} \tag{10}$$

Note that, for x_1 and x_3, the belief masses does not move further from each agent upon update. Also, $\mathbf{b}_X^{A[0]}(x_2) = \mathbf{b}_X^{B[0]}(x_2) = \mathbf{b}_X^{A[1]}(x_2) = \mathbf{b}_X^{A[1]}(x_2) = 0.4$. Because A and B have dogmatic opinions, $\mathbf{b}_X = \mathbf{P}_X$.

Theorem 1. *If $\mathbb{X}$ is a domain of size $k = 2$, and $\omega_t^{A[0]}$ and $\omega_t^{A[0]}$ are dogmatic opinions, then the update function for 2 agents with constant group uncertainty from Def. 7 corresponds to the Def. 6 for 2 agents.*

Proof. Since $\omega_X^{A[0]}$ (and $\omega_X^{B[0]}$) is a dogmatic opinion, $\omega_X^{A[1]}$ is defined as

$$\omega_X^{A[1]} = \begin{cases} \mathbf{b}_X^{A[1]} = \dfrac{\mathbf{b}_X^{A[0]}}{2} + \dfrac{\mathbf{b}_X^{A[0]} + \mathbf{P}_B^A(t_B)(\mathbf{b}_X^{B[0]} - \mathbf{b}_X^{A[0]})}{2} \\ u_X^{A[1]} = 0 \\ \mathbf{a}_X^{A[1]} = \mathbf{a}_X^{A[0]} \end{cases} \tag{11}$$

Also, $\mathbf{b}_X^{A[t]} = \mathbf{P}_X^{A[t]}$ for any t.

We want to map the $\mathbf{P}_X^{A[1]}(x)$ to $\mathrm{Bel}_p^{t+1}(A)$, because we consider that A's take decisions using $\mathbf{P}_X^{A[1]}$ in subjective logic and $\mathrm{Bel}_p^{t+1}(A)$ in [1].

If $x = p$, the belief mass $\mathbf{b}_X^{A[0]}$ corresponds to the belief state $\mathrm{Bel}_p^0(A)$, the belief mass $\mathbf{b}_X^{B[0]}$ corresponds to the belief state $\mathrm{Bel}_p^0(B)$, and the posterior probability $\mathbf{P}_B^A(t_B)$ from the trust opinion ω_B^A corresponds to the influence $\mathrm{In}(B, A)$, then $\mathbf{b}_X^{A[t]}(x)$ and $\mathrm{Bel}_p^{t+1}(A)$ have the same definition and $\mathbf{P}_X^{A[1]} = \mathrm{Bel}_p^{t+1}(A)$.

It is easy to extend the update function to arbitrarily many agents, using the same Def. 6. It is possible to extend the update function for domains of size $k > 2$, but with no clear correspondence for belief state.

Theorem 2. *If $\mathbb{X}$ is a domain of size $k = 2$, the update function from Def. 7 satisfies the desiderata from Def. 4: 1. weak convergence, 2. Markovian behaviour, 3. idempotence, and 4.(a) non-increasing group uncertainty.*

Proof. 1. weak convergence: Suppose that $\mathbf{b}_X^{A[t]}(x) \leq \mathbf{b}_X^{B[t]}(x)$, then

$$\mathbf{b}_X^{A[t+1]}(x) = \mathbf{b}_X^{A[t]}(x) + \left(\frac{\mathbf{P}_B^A(t_B)(\mathbf{b}_X^{B[t]}(x) - \mathbf{b}_X^{B[t]}(x))}{2} \right). \tag{12}$$

Since $\mathbf{b}_X^{A[t]}(x) \leq \mathbf{b}_X^{B[t]}(x)$ and $\mathbf{P}_B^A(t_B) \geq 0$, $\mathbf{b}_X^{A[t]}(x) \leq \mathbf{b}_X^{A[t+1]}(x)$. Now,

$$\mathbf{b}_X^{A[t+1]}(x) = \mathbf{b}_X^{B[t]}(x) - \left(\mathbf{b}_X^{B[t]}(x) - \mathbf{b}_X^{A[t]} - \frac{\mathbf{P}_B^A(t_B)(\mathbf{b}_X^{B[t]}(x) - \mathbf{b}_X^{B[t]}(x))}{2} \right). \tag{13}$$

Since $\mathbf{b}_X^{A[t]}(x) \leq \mathbf{b}_X^{B[t]}(x)$ and $\mathbf{P}_B^A(t_B) \leq 1$, $\mathbf{b}_X^{A[t+1]}(x) \leq \mathbf{b}_X^{B[t]}(x)$. Therefore, $\mathbf{b}_X^{A[t]}(x) \leq \mathbf{b}_X^{A[t+1]}(x) \leq \mathbf{b}_X^{B[t]}(x)$. With a similar proof, we can show that $u_X^{A[t]} \leq u_X^{A[t+1]} \leq u_X^{B[t]}$.

Now, we need to show that $\mathbf{P}_X^{A[t]}(x) \leq \mathbf{P}_X^{A[t+1]}(x) \leq \mathbf{P}_X^{B[t]}(x)$. By definition of projected probability,

$$\mathbf{b}_X^{A[t]}(x) + u_X^{A[t]}\mathbf{a}_X^{A[t]}(x) \leq \mathbf{b}_X^{A[t+1]}(x) + u_X^{A[t+1]}\mathbf{a}_X^{A[t]}(x) \leq \mathbf{b}_X^{B[t]}(x) + u_X^{B[t]}\mathbf{a}_X^{B[t]}(x). \tag{14}$$

Since $\mathbf{b}_X^{A[t]}(x) \leq \mathbf{b}_X^{A[t+1]}(x) \leq \mathbf{b}_X^{B[t]}(x)$ and $u_X^{A[t]} \leq u_X^{A[t+1]} \leq u_X^{B[t]}$,

$$\mathbf{P}_X^{A[t]}(x) \leq \mathbf{P}_X^{A[t+1]}(x) \leq \mathbf{P}_X^{B[t]}(x). \tag{15}$$

The proof for $\mathbf{P}_X^{A[t]}(x) \geq \mathbf{P}_X^{A[t+1]}(x) \geq \mathbf{P}_X^{B[t]}(x)$ is similar.

2. Markovian behavior: The definition does not uses opinions from time before t. We also consider the agents does not have a history of acquired evidence over time.
3. idempotence: Suppose that $\omega_X^{A[t]} = \omega_X^{B[t]}$, then

$$\begin{aligned} \mathbf{b}_X^{A[t+1]}(x) &= \frac{\mathbf{b}_X^{A[t]}(x) + \mathbf{b}_X^{A[t]}(x) + \mathbf{P}_B^A(\mathbf{b}_X^{A[t]}(x) + \mathbf{b}_X^{A[t]}(x))}{2} \\ &= \frac{\mathbf{b}_X^{A[t]}(x) + \mathbf{b}_X^{A[t]}(x)}{2} = \mathbf{b}_X^{A[t]}(x) \end{aligned} \tag{16}$$

4.(a) non-increasing group uncertainty i.e. $u_X^{A[t]} + u_X^{B[t]} = u_X^{A[t+1]} + u_X^{B[t+1]}$.

$$\begin{aligned} u_X^{A[t]} + u_X^{B[t]} &= u_X^{A[t+1]} + u_X^{B[t+1]} \\ &= \frac{u_X^{A[t]}}{2} + \frac{u_X^{A[t]} + \mathbf{P}_B^A(t_A)(u_X^{B[t]} - u_X^{A[t]})}{2} \\ &\quad + \frac{u_X^{B[t]}}{2} + \frac{u_X^{B[t]} + \mathbf{P}_A^B(t_B)(u_X^{A[t]} - u_X^{B[t]})}{2} \\ &= u_X^{A[t]} + u_X^{B[t]} \end{aligned} \tag{17}$$

To sum up, in this section we first presented the belief update function from [1]. Then we proposed a subjective logic belief update function, and in Theorem 1 we showed that our SL belief update function corresponds to the belief update function from [1]. Next in Theorem 2 we proved that our new belief update function satisfies the four desiderata from Def. 4. This provides us with several valuable pieces of information: first, it is possible to define at least one belief update function in SL satisfying the four desiderata. Second, we can define a belief update function in SL that meaningfully corresponds to the belief update function from [1], even though SL is in fact more expressive than the formalism used in [1], which does not provide for uncertainty. This shows that SL is promising for modelling social networks, since it can already capture the information used in [1], and also is equipped with an additional method for reasoning about agents' uncertainty.

6 Conclusions and Future Work

In this paper, we have used subjective logic (SL) to consistently extend the dynamic models from [1] with *uncertainty*, allowing a more nuanced and realistic model of complex opinions in social networks.

We identified four desiderata of properties that describe our interpretation of rationality when updating opinions. From SL, we selected the trust discount and belief fusion operators to define our update function. We showed through examples that our function does not satisfy the desiderata. Our next step for this part of our work will be to analyse to what situations in social networks our function applies. We conjecture that an update function using cumulative belief fusion might model a situation in which agents that are exposed to the same opinion many times will develop more confidence in their opinion than the source of the information itself.

One of our goals is to represent belief about multiple issues. We defined a new update function using SL that is similar to a definition of classical belief update function from [1] that already satisfies our desiderata. We showed that the new update function satisfies our desiderata, showing that it is possible to define at least one belief update function in SL. We also showed that our SL belief update function corresponds to the belief update function from [1]. This shows that SL is promising for modelling social networks, since it can already capture the information used in [1], and also is equipped with an additional method for reasoning about agents' uncertainty. Since SL is more expressive than the formalism used in [1], the next step for this function is to extend the model using SL operators and what these operations mean from the classic formalism perspective.

Acknowledgments. I would like to thank Dr. Sophia Knight and Dr. Mário S. Alvim for the exceptional advising and guidance through this work.

References

1. Alvim, M.S., Knight, S., Valencia, F.: Toward a formal model for group polarization in social networks. In: The Art of Modelling Computational Systems. LNCS, vol. 11760, pp. 419–441. Springer (2019)
2. Baltag, A., Christoff, Z., Rendsvig, R.K., Smets, S.: Dynamic epistemic logics of diffusion and prediction in social networks. Studia Logica **107**(3), 489–531 (2019)
3. Christoff, Z.L.: Dynamic logics of networks: information flow and the spread of opinion. Ph.D. thesis, Universiteit van Amsterdam (2016)
4. Gargiulo, F., Gandica, Y.: The role of homophily in the emergence of opinion controversies. CoRR **abs/1612.05483** (2016), http://arxiv.org/abs/1612.05483
5. Hunter, A.: Reasoning about trust and belief change on a social network: A formal approach. In: Liu, J.K., Samarati, P. (eds.) Information Security Practice and Experience 2017. Lecture Notes in Computer Science, vol. 10701, pp. 783–801. Springer (2017)
6. Jøsang, A.: Subjective Logic - A Formalism for Reasoning Under Uncertainty. Artificial Intelligence: Foundations, Theory, and Algorithms, Springer (2016)
7. Jøsang, A.: Categories of belief fusion. Journal of Advances in Information Fusion **13**(2), 235–254 (2018)
8. Liu, F., Seligman, J., Girard, P.: Logical dynamics of belief change in the community. Synthese **191**(11), 2403–2431 (2014)

9. Sîrbu, A., Pedreschi, D., Giannotti, F., Kertész, J.: Algorithmic bias amplifies opinion polarization: A bounded confidence model. CoRR **abs/1803.02111** (2018)
10. Young Pedersen, M.: Polarization and Echo Chambers: A Logical Analysis of Balance and Triadic Closure in Social Networks. Master's thesis, Universiteit van Amsterdam (2019)