

# Forecasting Argumentation Frameworks

Benjamin Irwin, Antonio Rago and Francesca Toni

Imperial College London

{benjamin.irwin19, a.rago, ft}@imperial.ac.uk

## Abstract

We introduce *Forecasting Argumentation Frameworks (FAFs)*, a novel argumentation-based methodology for forecasting informed by recent judgmental forecasting research. FAFs comprise update frameworks which empower (human or artificial) agents to argue over time about the probability of outcomes, e.g. the winner of a political election or a fluctuation in inflation rates, whilst flagging perceived *irrationality* in the agents' behaviour with a view to improving their forecasting accuracy. FAFs include five argument types, amounting to standard *pro/con* arguments, as in bipolar argumentation, as well as novel *proposal* arguments and *increase/decrease amendment* arguments. We adapt an existing gradual semantics for bipolar argumentation to determine the aggregated dialectical strength of proposal arguments and define irrational behaviour. We then give a simple aggregation function which produces a final group forecast from rational agents' individual forecasts. We identify and study properties of FAFs and conduct an empirical evaluation which signals FAFs' potential to increase the forecasting accuracy of participants.

## 1 Introduction

Historically, humans have performed inconsistently in judgemental forecasting (Makridakis, Hogarth, and Gaba 2010; Tetlock 2017), which incorporates subjective opinion and probability estimates to predictions (Lawrence et al. 2006). Yet, human judgement remains essential in cases where pure statistical methods are not applicable, e.g. where historical data alone is insufficient or for one-off, more 'unknown' events (Petropoulos, Fildes, and Goodwin 2016; Arvan et al. 2019; Baets and Harvey 2020). Judgemental forecasting is widely relied upon for decision-making (Nikolopoulos 2021), in myriad fields from epidemiology to national security (Nikolopoulos et al. 2015; Litsiou et al. 2019). Effective tools to help humans improve their predictive capabilities thus have enormous potential for impact. Two recent global events – the COVID-19 pandemic and the US withdrawal from Afghanistan – underscore this by highlighting the human and financial cost of predictive deficiency. A multi-purpose system which could improve our ability to predict the incidence and impact of events by as little as 5%, could save millions of lives and be worth trillions of dollars per year (Tetlock and Gardner 2016).

Research on judgemental forecasting (see (Lawrence et al. 2006; Zellner et al. 2021) for overviews), including the recent, groundbreaking 'Superforecasting Experiment' (Tetlock and Gardner 2016), is instructive in establishing the desired properties for systems for supporting forecasting. In addition to reaffirming the importance of fine-grained probabilistic reasoning (Mellers et al. 2015b), this literature points to the benefits of some group techniques versus solo forecasting (Landeta, Barrutia, and Lertxundi 2011; Tetlock et al. 2014), of synthesising qualitative and quantitative information (Lawrence et al. 2006), of combating agents' irrationality (Chang et al. 2016) and of high agent engagement with the forecasting challenge, e.g. robust debating (Landeta, Barrutia, and Lertxundi 2011) and frequent prediction updates (Mellers et al. 2015b).

Meanwhile, *computational argumentation* (see (Atkinson et al. 2017; Baroni et al. 2018) for recent overviews) is a field of AI which involves reasoning with uncertainty and resolving conflicting information, e.g. in natural language debates. As such, it is an ideal candidate for aggregating the broad, polymorphous set of information involved in judgemental group forecasting. An extensive and growing literature is based on various argumentation frameworks – rule-based systems for aggregating, representing and evaluating sets of arguments, such as those applied in the contexts of *scheduling* (Cyras et al. 2019), *fact checking* (Kotonya and Toni 2020) or in various instances of *explainable AI* (Cyras et al. 2021). Subsets of the requirements for forecasting systems are addressed by individual formalisms, e.g. *probabilistic argumentation* (Dung and Thang 2010; Thimm 2012; Hunter 2013; Fazzinga, Flesca, and Furfaro 2018) may effectively represent and analyse uncertain arguments about the future. However, we posit that a purpose-built argumentation framework for forecasting is essential to effectively utilise argumentation's reasoning capabilities in this context.

In this paper, we attempt to cross-fertilise these two as of yet unconnected academic areas. We draw from forecasting literature to inform the design of a new computational argumentation approach: *Forecasting Argumentation Frameworks (FAFs)*. FAFs empower (human and artificial) agents to structure debates in real time and to deliver argumentation-based forecasting. They offer an approach in the spirit of *deliberative democracy* (Bessette 1980) to respond to a forecasting problem over time. The steps which

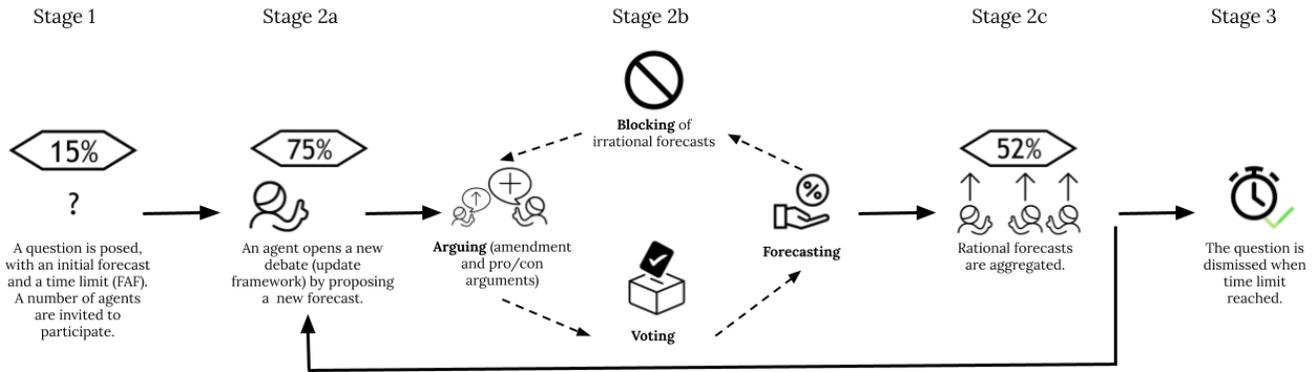


Figure 1: The step-by-step process of a FAF over its lifetime.

underpin FAFs are depicted in Figure 1 (referenced throughout) and can be described in simple terms as follows: a FAF is initialised with a time limit (for the overall forecasting process and for each iteration therein) and a pre-agreed ‘base-rate’ forecast  $\mathcal{F}$  (Stage 1), e.g. based on historical data. Then, the forecast is revised by one or more (non-concurrent) debates, in the form of ‘update frameworks’ (Stage 2), opened and resolved by participating agents (until the specified time limit is reached). Each update framework begins with a proposed revision to the current forecast (Stage 2a), and proceeds with a cycle of argumentation (Stage 2b) about the proposed forecast, voting on said argumentation and forecasting. Forecasts deemed ‘irrational’ with a view to agents’ argumentation and voting are blocked. Finally, the rational forecasts are aggregated and the result replaces the current group forecast (Stage 2c). This process may be repeated over time in an indefinite number of update frameworks (thus continually revising the group forecast) until the (overall) time limit is reached. The composite nature of this process enables the appraisal of new information relevant to the forecasting question as and when it arrives. Rather than confronting an unbounded forecasting question with a diffuse set of possible debates open at once, all agents concentrate their argumentation on a single topic (a proposal) at any given time.

After giving the necessary background on forecasting and argumentation (§2), we formalise our update frameworks for Step 2a (§3). We then give our notion of rationality (Step 2b), along with our new method for aggregating rational forecasts (Step 2c) from a group of agents (§4) and FAFs overall. We explore the underlying properties of FAFs (§5), before describing an experiment with a prototype implementing our approach (§6). Finally, we conclude and suggest potentially fruitful avenues for future work (§7).

## 2 Background

### 2.1 Forecasting

Studies on the efficacy of judgemental forecasting have shown mixed results (Makridakis, Hogarth, and Gaba 2010; Tetlock 2017; Goodwin, Gönül, and Önkal-Atay 2019). Limitations of the judgemental approach are a result of well-documented cognitive biases (Kahneman 2012), irrational-

ities in human probabilistic reasoning which lead to distortion of forecasts. Manifold methodologies have been explored to improve judgemental forecasting accuracy to varying success (Lawrence et al. 2006). These methodologies include, but are not limited to, prediction intervals (Lawrence and Makridakis 1989), decomposition (MacGregor and Armstrong 1994), structured analogies (Green and Armstrong 2007; Nikolopoulos et al. 2015) and unaided judgement (Litsiou et al. 2019). Various group forecasting techniques have also been explored (Linstone and Turoff 1975; Delbecq, Van den Ven, and Gustafson 1975; Landeta, Barrutia, and Lertxundi 2011), although the risks of groupthink (McNees 1987) and the importance of maintaining the independence of each group member’s individual forecast are well established (Armstrong 2001). Recent advances in the field have been led by Tetlock and Mellers’ superforecasting experiment (Tetlock and Gardner 2016), which leveraged geopolitical forecasting tournaments and a base of 5000 volunteer forecasters to identify individuals with consistently exceptional accuracy (top 2%). The experiment’s findings underline the effectiveness of group forecasting orientated around debating (Tetlock et al. 2014), and demonstrate a specific cognitive-intellectual approach conducive for forecasting (Mellers et al. 2015a; Mellers et al. 2015b), but stop short of suggesting a concrete universal methodology for higher accuracy. Instead, Tetlock draws on his own work and previous literature to crystallise a broad set of methodological principles by which superforecasters abide (Tetlock and Gardner 2016, pg.144):

- *Pragmatic*: not wedded to any idea or agenda;
- *Analytical*: capable of stepping back from the tip-of-your-nose perspective and considering other views;
- *Dragonfly-eyed*: value diverse views and synthesise them into their own;
- *Probabilistic*: judge using many grades of maybe;
- *Thoughtful updaters*: when facts change, they change their minds;
- *Good intuitive psychologists*: aware of the value of checking thinking for cognitive and emotional biases.

Subsequent research after the superforecasting experiment has included exploring further optimal forecasting tournament preparation (Horowitz et al. 2021; Katsagounos et al. 2021) and extending Tetlock and Mellers’ approach to answer broader, more time-distant questions (Page, Aiken, and Murdick 2020). It should be noted that there have been no recent advances on computational toolkits for the field similar to that proposed in this paper.

## 2.2 Computational Argumentation

We posit that existing argumentation formalisms are not well suited for the aforementioned future-based arguments, which are necessarily semantically and structurally different from arguments about present or past concerns. Specifically, forecasting arguments are inherently probabilistic and must deal with the passage of time and its implications for the outcomes at hand. Further, several other important characteristics can be drawn from the forecasting literature which render current argumentation formalisms unsuitable, e.g. the paramountcy of dealing with bias (in data and cognitive), forming granular conclusions, fostering group debate and the co-occurrence of qualitative and quantitative arguing.

Nonetheless, several of these characteristics have been previously explored in argumentation and our formalisation draws from several existing approaches. First and foremost, it draws in spirit from abstract argumentation frameworks (AAFs) (Dung 1995), in that the arguments’ inner contents are ignored and the focus is on the relationships between arguments. However, we consider arguments of different types and an additional relation of support (pro), rather than attack (con) alone as in (Dung 1995). Past work has also introduced probabilistic constraints into argumentation frameworks. Probabilistic AAFs (prAAFs) propose two divergent ways for modelling uncertainty in abstract argumentation using probabilities - the constellation approach (Dung and Thang 2010; Li, Oren, and Norman 2011) and the epistemic approach (Hunter 2013; Hunter and Thimm 2014; Hunter, Polberg, and Thimm 2020). These formalisations use probability as a means to assess uncertainty over the validity of arguments (epistemic) or graph topology (constellation), but do not enable reasoning *with* or *about* probability, which is fundamental in forecasting. In exploring temporality, (Cobo, Martínez, and Simari 2010) augment AAFs by providing each argument with a limited lifetime. Temporal constraints have been extended in (Cobo, Martínez, and Simari 2011) and (Baron et al. 2014). Elsewhere, (Rago and Toni 2017) have used argumentation to model irrationality or bias in agents. Finally, a wide range of gradual evaluation methods have gone beyond traditional qualitative semantics by measuring arguments’ acceptability on a scale (normally  $[0,1]$ ) (Leite and Martins 2011; Evripidou and Toni 2012; Amgoud et al. 2017; Amgoud and Ben-Naim 2018; Amgoud and Ben-Naim 2016). Many of these approaches have been unified as Quantitative Bipolar Argumentation Frameworks (QBAFs) in (Baroni, Rago, and Toni 2018).

Amongst existing approaches, of special relevance in this paper are Quantitative Argumentation Debate (QuAD) frameworks (Baroni et al. 2015), i.e. 5-tuples  $\langle \mathcal{X}^a, \mathcal{X}^c, \mathcal{X}^p, \mathcal{R}, \tau \rangle$  where  $\mathcal{X}^a$  is a finite set of *answer* arguments (to im-

PLICIT *issues*);  $\mathcal{X}^c$  is a finite set of *con* arguments;  $\mathcal{X}^p$  is a finite set of *pro* arguments;  $\mathcal{X}^a, \mathcal{X}^c$  and  $\mathcal{X}^p$  are pairwise disjoint;  $\mathcal{R} \subseteq (\mathcal{X}^c \cup \mathcal{X}^p) \times (\mathcal{X}^a \cup \mathcal{X}^c \cup \mathcal{X}^p)$  is an acyclic binary relation;  $\tau : (\mathcal{X}^a \cup \mathcal{X}^c \cup \mathcal{X}^p) \rightarrow [0, 1]$  is a total function:  $\tau(a)$  is the *base score* of  $a$ . Here, attackers and supporters of arguments are determined by the pro and con arguments they are in relation with. Formally, for any  $a \in \mathcal{X}^a \cup \mathcal{X}^c \cup \mathcal{X}^p$ , the set of *con* (*pro*) arguments of  $a$  is  $\mathcal{R}^-(a) = \{b \in \mathcal{X}^c | (b, a) \in \mathcal{R}\}$  ( $\mathcal{R}^+(a) = \{b \in \mathcal{X}^p | (b, a) \in \mathcal{R}\}$ , resp.). Arguments in QuAD frameworks are scored by the *Discontinuity-Free QuAD* (DF-QuAD) algorithm (Rago et al. 2016), using the argument’s intrinsic base score and the *strengths* of its pro/con arguments. Given that DF-QuAD is used to define our method (see Def. 4), for completeness we define it formally here. DF-QuAD’s *strength aggregation function* is defined as  $\Sigma : [0, 1]^* \rightarrow [0, 1]$ , where for  $\mathcal{S} = (v_1, \dots, v_n) \in [0, 1]^*$ : if  $n = 0$ ,  $\Sigma(\mathcal{S}) = 0$ ; if  $n = 1$ ,  $\Sigma(\mathcal{S}) = v_1$ ; if  $n = 2$ ,  $\Sigma(\mathcal{S}) = f(v_1, v_2)$ ; if  $n > 2$ ,  $\Sigma(\mathcal{S}) = f(\Sigma(v_1, \dots, v_{n-1}), v_n)$ ; with the *base function*  $f : [0, 1] \times [0, 1] \rightarrow [0, 1]$  defined, for  $v_1, v_2 \in [0, 1]$ , as:  $f(v_1, v_2) = v_1 + (1 - v_1) \cdot v_2 = v_1 + v_2 - v_1 \cdot v_2$ . After separate aggregation of the argument’s pro/con descendants, the combination function  $c : [0, 1] \times [0, 1] \times [0, 1] \rightarrow [0, 1]$  combines  $v^-$  and  $v^+$  with the argument’s base score ( $v^0$ ):  $c(v^0, v^-, v^+) = v^0 - v^0 \cdot |v^+ - v^-|$  if  $v^- \geq v^+$  and  $c(v^0, v^-, v^+) = v^0 + (1 - v^0) \cdot |v^+ - v^-|$  if  $v^- < v^+$ , resp. The inputs for the combination function are provided by the *score function*,  $\sigma : \mathcal{X}^a \cup \mathcal{X}^c \cup \mathcal{X}^p \rightarrow [0, 1]$ , which gives the argument’s strength, as follows: for any  $x \in \mathcal{X}^a \cup \mathcal{X}^c \cup \mathcal{X}^p$ :  $\sigma(x) = c(\tau(x), \Sigma(\sigma(\mathcal{R}^-(x))), \Sigma(\sigma(\mathcal{R}^+(x))))$  where  $(x_1, \dots, x_n)$  is an arbitrary permutation of the ( $n \geq 0$ ) con arguments in  $\mathcal{R}^-(x)$ ,  $\sigma(\mathcal{R}^-(x)) = (\sigma(x_1), \dots, \sigma(x_n))$  (similarly for pro arguments). Note that the DF-QuAD notion of  $\sigma$  can be applied to any argumentation framework where arguments are equipped with base scores and pro/con arguments. We will do so later, for our novel formalism.

## 3 Update Frameworks

We begin by defining the individual components of our frameworks, starting with the fundamental notion of a *forecast*. This is a probability estimate for the positive outcome of a given (binary) question.

**Definition 1.** A forecast  $\mathcal{F}$  is the probability  $P(Q = \text{true}) \in [0, 1]$  for a given forecasting question  $Q$ .

**Example 1.** Consider the forecasting question  $Q$ : ‘Will the Tokyo 2020 Summer Olympics be cancelled/postponed to another year?’. Here, the *true* outcome amounts to the Olympics being cancelled/postponed (and false to it taking place in 2020 as planned). Then, a forecast  $\mathcal{F}$  may be  $P(Q = \text{true}) = 0.15$  which amounts to a 15% probability of the Olympics being cancelled/postponed. Note that  $\mathcal{F}$  may have been introduced as part of an update framework (herein described), or as an initial base rate at the outset of a FAF (Stage 1 in Figure 1).

In the remainder, we will often drop  $Q$ , implicitly assuming it is given, and use  $P(\text{true})$  to stand for  $P(Q = \text{true})$ .

In order to empower agents to reason about probabilities and thus support forecasting, we need, in addition to

*pro/con* arguments as in QuAD frameworks, two new argument types:

- *proposal* arguments, each about some forecast (and its underlying forecasting question); each proposal argument  $\mathcal{P}$  has a *forecast* and, optionally, some supporting *evidence*; and
- *amendment* arguments, which suggest a modification to some forecast’s probability by increasing or decreasing it, and are accordingly separated into disjoint classes of *increase* and *decrease* (amendment) arguments.<sup>1</sup>

Note that amendment arguments are introduced specifically for arguing about a proposal argument, given that traditional QuAD *pro/con* arguments are of limited use when the goal is to judge the acceptability of a probability, and that in forecasting agents must not only decide *if* they agree/disagree but also *how* they agree/disagree (i.e. whether they believe the forecast is too low or too high considering, if available, the evidence). Amendment arguments, with their increase and decrease classes, provide for this.

**Example 2.** A proposal argument  $\mathcal{P}$  in the Tokyo Olympics setting may comprise forecast: ‘There is a 75% chance that the Olympics will be cancelled/postponed to another year’. It may also include evidence: ‘A new poll today shows that 80% of the Japanese public want the Olympics to be cancelled. The Japanese government is likely to buckle under this pressure.’ This argument may aim to prompt updating the earlier forecast in Example 1. A decrease amendment argument may be  $x_1^\downarrow$ : ‘The International Olympic Committee and the Japanese government will ignore the views of the Japanese public’. An increase amendment argument may be  $x_1^\uparrow$ : ‘Japan’s increasingly popular opposition parties will leverage this to make an even stronger case for cancellation’.

Intuitively, a proposal argument is the focal point of the argumentation. It typically suggests a new forecast to replace prior forecasts, argued on the basis of some new evidence (as in the earlier example). We will see that proposal arguments remain immutable through each debate (update framework), which takes place via amendment arguments and standard *pro/con* arguments. Note that, wrt QuAD argument types, proposal arguments replace issues and amendment arguments replace answers, in that the former are driving the debates and the latter are the options up for debate. Note also that amendment arguments merely state a direction wrt  $\mathcal{F}^{\mathcal{P}}$  and do not contain any more information, such as *how much* to alter  $\mathcal{F}^{\mathcal{P}}$  by. We will see that alteration can be determined by *scoring* amendment arguments.

<sup>1</sup>Note that we decline to include a third type of amendment argument for arguing that  $\mathcal{F}^{\mathcal{P}}$  is just right. This choice rests on the assumption that additional information always necessitates a change to  $\mathcal{F}^{\mathcal{P}}$ , however granular that change may be. This does not restrict individual agents arguing about  $\mathcal{F}^{\mathcal{P}}$  from casting  $\mathcal{F}^{\mathcal{P}}$  as their own final forecast. However, rather than cohering their argumentation around  $\mathcal{F}^{\mathcal{P}}$ , which we hypothesise would lead to high risk of groupthink (McNees 1987), agents are compelled to consider the impact of their amendment arguments on this more granular level.

Proposal and amendment arguments, alongside *pro/con* arguments, form part of our novel update frameworks (Stage 2 of Figure 1), defined as follows:

**Definition 2.** An update framework is a nonad  $\langle \mathcal{P}, \mathcal{X}, \mathcal{X}^-, \mathcal{X}^+, \mathcal{R}^p, \mathcal{R}, \mathcal{A}, \mathcal{V}, \mathcal{F}^{\mathcal{A}} \rangle$  such that:

- $\mathcal{P}$  is a single proposal argument with forecast  $\mathcal{F}^{\mathcal{P}}$  and, optionally, evidence  $\mathcal{E}^{\mathcal{P}}$  for this forecast;
- $\mathcal{X} = \mathcal{X}^\uparrow \cup \mathcal{X}^\downarrow$  is a finite set of amendment arguments composed of subsets  $\mathcal{X}^\uparrow$  of increase arguments and  $\mathcal{X}^\downarrow$  of decrease arguments;
- $\mathcal{X}^-$  is a finite set of con arguments;
- $\mathcal{X}^+$  is a finite set of pro arguments;
- the sets  $\{\mathcal{P}\}, \mathcal{X}^\uparrow, \mathcal{X}^\downarrow, \mathcal{X}^-$  and  $\mathcal{X}^+$  are pairwise disjoint;
- $\mathcal{R}^p \subseteq \mathcal{X} \times \{\mathcal{P}\}$  is a directed acyclic binary relation between amendment arguments and the proposal argument (we may refer to this relation informally as ‘probabilistic’);
- $\mathcal{R} \subseteq (\mathcal{X}^- \cup \mathcal{X}^+) \times (\mathcal{X} \cup \mathcal{X}^- \cup \mathcal{X}^+)$  is a directed acyclic, binary relation from *pro/con* arguments to amendment/*pro/con* arguments (we may refer to this relation informally as ‘argumentative’);
- $\mathcal{A} = \{a_1, \dots, a_n\}$  is a finite set of agents ( $n > 1$ );
- $\mathcal{V} : \mathcal{A} \times (\mathcal{X}^- \cup \mathcal{X}^+) \rightarrow [0, 1]$  is a total function such that  $\mathcal{V}(a, x)$  is the vote of agent  $a \in \mathcal{A}$  on argument  $x \in \mathcal{X}^- \cup \mathcal{X}^+$ ; with an abuse of notation, we let  $\mathcal{V}_a : (\mathcal{X}^- \cup \mathcal{X}^+) \rightarrow [0, 1]$  represent the votes of a single agent  $a \in \mathcal{A}$ , e.g.  $\mathcal{V}_a(x) = \mathcal{V}(a, x)$ ;
- $\mathcal{F}^{\mathcal{A}} = \{\mathcal{F}_{a_1}^{\mathcal{A}}, \dots, \mathcal{F}_{a_n}^{\mathcal{A}}\}$  is such that  $\mathcal{F}_{a_i}^{\mathcal{A}}$ , where  $i \in \{1, \dots, n\}$ , is the forecast of agent  $a_i \in \mathcal{A}$ .

Note that *pro* (*con*) arguments can be seen as supporting (attacking, resp.) other arguments via  $\mathcal{R}$ , as in the case of conventional QuAD frameworks (Baroni et al. 2015).

**Example 3.** A possible update framework in our running setting may include  $\mathcal{P}$  as in Example 2 as well as (see Table 1)  $\mathcal{X} = \{x_1^\downarrow, x_2^\downarrow, x_1^\uparrow\}$ ,  $\mathcal{X}^- = \{x_1^-, x_2^-, x_3^-\}$ ,  $\mathcal{X}^+ = \{x_1^+, x_2^+\}$ ,  $\mathcal{R}^p = \{(x_1^\downarrow, \mathcal{P}), (x_2^\downarrow, \mathcal{P}), (x_3^-, \mathcal{P})\}$ , and  $\mathcal{R} = \{(x_1^-, x_1^\downarrow), (x_2^-, x_1^\downarrow), (x_3^-, x_1^\downarrow), (x_1^+, x_2^\downarrow), (x_2^+, x_1^\downarrow)\}$ . Figure 2 gives a graphical representation of these arguments and relations. Assuming  $\mathcal{A} = \{\text{alice}, \text{bob}, \text{charlie}\}$ ,  $\mathcal{V}$  may be such that  $\mathcal{V}_{\text{alice}}(x_1^-) = 1$ ,  $\mathcal{V}_{\text{bob}}(x_1^\uparrow) = 0$ , and so on.

Several considerations about update frameworks are in order. Firstly, they represent ‘stratified’ debates: graphically, they can be represented as trees with the proposal argument as root, amendment arguments as children of the root, and *pro/con* arguments forming the lower layers, as shown in Figure 2. This tree structure serves to focus argumentation towards the proposal (i.e. the probability and, if available, evidence) it puts forward. Second, we have chosen to impose a ‘structure’ on proposal arguments, whereby their forecast is distinct from their (optional) evidence. Here the forecast has special primacy over the evidence, because forecasts are the vital reference point and the drivers of debates in FAFs. They are, accordingly, both mandatory and required to ‘stand out’ to participating agents. In the spirit of abstract argumentation (Dung 1995), we nonetheless treat all arguments, including proposal arguments, as

	Content
$\mathcal{P}$	‘A new poll today shows that 80% of the Japanese public want the Olympics to be cancelled owing to COVID-19, and the Japanese government is likely to buckle under this pressure ( $\mathcal{E}^{\mathcal{P}}$ ). Thus, there is a 75% chance that the Olympics will be cancelled/postponed to another year’ ( $\mathcal{F}^{\mathcal{P}}$ ).
$x_1^\downarrow$	‘The International Olympic Committee and the Japanese government will ignore the views of the Japanese public’.
$x_2^\downarrow$	‘This poll comes from an unreliable source.’
$x_1^\uparrow$	‘Japan’s increasingly popular opposition parties will leverage this to make an even stronger case for cancellation.’
$x_1^-$	‘The IOC is bluffing - people are dying, Japan is experiencing a strike. They will not go ahead with the games if there is a risk of mass death.’
$x_2^-$	‘The Japanese government may renege on its commitment to the IOC, and use legislative or immigration levers to block the event.’
$x_3^-$	‘Japan’s government has sustained a high-approval rating in the last year and is strong enough to ward off opposition attacks.’
$x_1^+$	‘This pollster has a track record of failure on Japanese domestic issues.’
$x_2^+$	‘Rising anti-government sentiment on Japanese Twitter indicates that voters may be receptive to such arguments.’

Table 1: Arguments in the update framework in Example 3.

‘abstract’, and focus on relations between them rather between their components. In practice, therefore, amendment arguments may relate to a proposal argument’s forecast but also, if present, to its evidence. We opt for this abstract view on the assumption that the flexibility of this approach better suits judgmental forecasting, which has a diversity of use cases (e.g. including politics, economics and sport) where different argumentative approaches may be deployed (i.e. quantitative, qualitative, directly attacking amendment nodes or raising alternative POVs) and wherein forecasters may lack even a basic knowledge of argumentation. We leave the study of structured variants of our framework (e.g. see overview in (Besnard et al. 2014)) to future work: these may consider finer-grained representations of all arguments in terms of different components, and finer-grained notions of relations between components, rather than full arguments. Third, in update frameworks, voting is restricted to pro/con arguments. Preventing agents from voting directly on amendment arguments mitigates against the risk of arbitrary judgements: agents cannot make off-the-cuff estimations but can only express their beliefs via (pro/con) argumentation, thus ensuring a more rigorous process of appraisal for the proposal and amendment arguments. Note that rather than facilitating voting on arguments using a two-valued perspective (i.e. positive/negative) or a three-valued

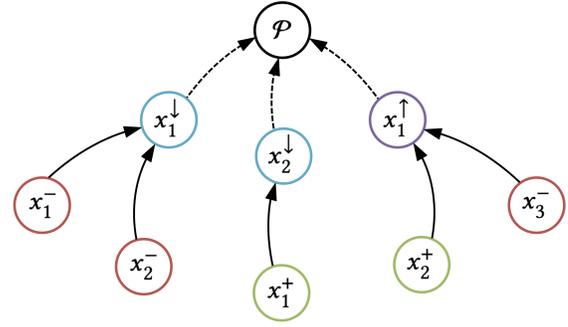


Figure 2: A graphical representation of arguments and relations in the update framework from Example 3. Nodes represent proposal ( $\mathcal{P}$ ), increase ( $\uparrow$ ), decrease ( $\downarrow$ ), pro (+) and con (-) arguments, while dashed/solid edges indicate, resp., the  $\mathcal{R}^{\mathcal{P}}/\mathcal{R}$  relations.

perspective (i.e. positive/negative/neutral),  $\mathcal{V}$  allows agents to cast more granular judgements of (pro/con) argument acceptability, the need for which has been highlighted in the literature (Mellers et al. 2015b). Finally, although we envisage that arguments of all types are put forward by agents during debates, we do not capture this mapping in update frameworks. Thus, we do not capture who put forward which arguments, but instead only use votes to encode and understand agents’ views. This enables more nuanced reasoning and full engagement on the part of agents with alternative viewpoints (i.e. an agent may freely argue both for and against a point before taking an explicit view with their voting). Such conditions are essential in a healthy forecasting debate (Landeta, Barrutia, and Lertxundi 2011; Mellers et al. 2015b).

In the remainder of this paper, with an abuse of notation, we often use  $\mathcal{F}^{\mathcal{P}}$  to denote, specifically, the probability advocated in  $\mathcal{F}^{\mathcal{P}}$  (e.g. 0.75 in Example 2).

## 4 Aggregating Rational Forecasts

In this section we formally introduce (in §4.1) our notion of rationality and discuss how it may be used to identify, and subsequently ‘block’, undesirable behaviour in forecasters. We then define (in §4.2) a method for calculating a revised forecast (Stage 2c of Figure 1), which aggregates the views of all agents in the update framework, whilst optimising their overall forecasting accuracy.

### 4.1 Rationality

Characterising an agent’s view as irrational offers opportunities to refine the accuracy of their forecast (and thus the overall aggregated group forecast). Our definition of rationality is inspired by, but goes beyond, that of QuAD-V (Rago and Toni 2017), which was introduced for the e-polling context. Whilst update frameworks eventually produce a single aggregated forecast on the basis of group deliberation, each agent is first evaluated for their rationality on an individual basis. Thus, as in QuAD-V, in order to define rationality for individual agents, we first reduce frameworks to *delegate frameworks* for each agent, which are the restriction of update frameworks to a single agent.

**Definition 3.** A delegate framework for an agent  $a$  is  $u_a = \langle \mathcal{P}, \mathcal{X}, \mathcal{X}^-, \mathcal{X}^+, \mathcal{R}^p, \mathcal{R}, a, \mathcal{V}_a, \mathcal{F}_a^A \rangle$ .

Note that all arguments in an update framework are included in each agent’s delegate framework, but only the agent’s votes and forecast are carried over.

Recognising the irrationality of an agent requires comparing the agent’s forecast against (an aggregation of) their opinions on the amendment arguments and, by extension, on the proposal argument. To this end, we evaluate the different parts of the update framework as follows. We use the DF-QuAD algorithm (Rago et al. 2016) to score each amendment argument for the agent, in the context of the pro/con arguments ‘linked’ to the amendment argument, using  $\mathcal{R}$ , in the context of the agent’s delegate framework. We refer to the DF-QuAD score function as  $\sigma$ . This requires a choice of base scores for amendment arguments as well as pro/con arguments. We assume the same base score  $\tau(x) = 0.5$  for all  $x \in \mathcal{X}$ ; in contrast, the base score of pro/con arguments is a result of the votes they received from the agent, in the spirit of QuAD-V (Rago and Toni 2017). The intuition behind assigning a neutral (0.5) base score for amendment arguments is that an agent’s estimation of their strength from the outset would be susceptible to bias and inaccuracy. Once each amendment argument has been scored (using  $\sigma$ ) for the agent, we aggregate the scores of all amendment arguments (for the same agent) to calculate the agent’s *confidence score* in the proposal argument (which underpins our rationality constraints), by weighting the mean average strength of this argument’s increase amendment relations against that of the set of decrease amendment relations:

**Definition 4.** Given a delegate framework  $u_a = \langle \mathcal{P}, \mathcal{X}, \mathcal{X}^-, \mathcal{X}^+, \mathcal{R}^p, \mathcal{R}, a, \mathcal{V}_a, \mathcal{F}_a^A \rangle$ , let  $\mathcal{X}^\uparrow = \{x_1^\uparrow, x_2^\uparrow, \dots, x_i^\uparrow\}$  and  $\mathcal{X}^\downarrow = \{x_1^\downarrow, x_2^\downarrow, \dots, x_j^\downarrow\}$ . Then,  $a$ ’s confidence score is as follows:

$$\text{if } i \neq 0, j \neq 0: \quad C_a(\mathcal{P}) = \frac{1}{i} \sum_{k=1}^i \sigma(x_k^\uparrow) - \frac{1}{j} \sum_{l=1}^j \sigma(x_l^\downarrow);$$

$$\text{if } i \neq 0, j = 0: \quad C_a(\mathcal{P}) = \frac{1}{i} \sum_{k=1}^i \sigma(x_k^\uparrow);$$

$$\text{if } i = 0, j \neq 0: \quad C_a(\mathcal{P}) = -\frac{1}{j} \sum_{l=1}^j \sigma(x_l^\downarrow);$$

$$\text{if } i = 0, j = 0: \quad C_a(\mathcal{P}) = 0.$$

Note that  $C_a(\mathcal{P}) \in [-1, 1]$ , which denotes the overall views of the agent on the forecast  $\mathcal{F}^{\mathcal{P}}$  (i.e. as to whether it should be *increased* or *decreased*, and how far). A negative (positive)  $C_a(\mathcal{P})$  indicates that an agent believes that  $\mathcal{F}^{\mathcal{P}}$  should be amended down (up, resp.). The size of  $C_a(\mathcal{P})$  reflects the degree of the agent’s certainty in either direction. In turn, we can constrain an agent’s forecast  $\mathcal{F}_a^A$  if it contradicts this belief as follows.

**Definition 5.** Given a delegate framework  $u_a = \langle \mathcal{P}, \mathcal{X}, \mathcal{X}^-, \mathcal{X}^+, \mathcal{R}^p, \mathcal{R}, a, \mathcal{V}_a, \mathcal{F}_a^A \rangle$ ,  $a$ ’s forecast  $\mathcal{F}_a^A$  is strictly rational

(wrt  $u_a$ ) iff:

$$\begin{aligned} &\text{if } C_a(\mathcal{P}) < 0 \text{ then } \mathcal{F}_a^A < \mathcal{F}^{\mathcal{P}} \\ &\text{if } C_a(\mathcal{P}) > 0 \text{ then } \mathcal{F}_a^A > \mathcal{F}^{\mathcal{P}} \\ &|C_a(\mathcal{P})| \geq \frac{|\mathcal{F}^{\mathcal{P}} - \mathcal{F}_a^A|}{\mathcal{F}^{\mathcal{P}}} \end{aligned}$$

Hereafter, we refer to forecasts which violate the first two constraints as, resp., *irrational increase* and *irrational decrease* forecasts, and to forecasts which violate the final constraint as *irrational scale* forecasts.

This definition of rationality preserves the integrity of group forecast in two ways. First, it prevents agents from forecasting against their beliefs: an agent cannot increase  $\mathcal{F}^{\mathcal{P}}$  if  $C_a(\mathcal{P}) < 0$  and an agent cannot decrease  $\mathcal{F}^{\mathcal{P}}$  if  $C_a(\mathcal{P}) > 0$ ; further, it ensures that agents cannot make forecasts disproportionate to their confidence score – *how far* an agent  $a$  deviates from the proposed change  $\mathcal{F}^{\mathcal{P}}$  is restricted by  $C_a(\mathcal{P})$ ; finally, an agent must have  $C_a(\mathcal{P})$  greater than or equal to the relative change to  $\mathcal{F}^{\mathcal{P}}$  denoted in their forecast  $\mathcal{F}_a^A$ . Note that the *irrational scale* constraint deals with just one direction of proportionality (i.e. providing only a maximum threshold for  $\mathcal{F}_a^A$ ’s deviation from  $\mathcal{F}^{\mathcal{P}}$ , but no minimum threshold). Here, we avoid bidirectional proportionality on the grounds that such a constraint would impose an arbitrary notion of arguments’ ‘impact’ on agents. An agent may have a very high  $C_a(\mathcal{P})$ , indicating their belief that  $\mathcal{F}^{\mathcal{P}}$  is too low, but may, we suggest, rationally choose to increase  $\mathcal{F}^{\mathcal{P}}$  by only a small amount (e.g. if, despite their general agreement with the arguments, they believe the overall issue at stake in  $\mathcal{P}$  to be minor or low impact to the overall forecasting question). Our definition of rationality, which relies on notions of argument strength derived from DF-QuAD, thus informs but does not wholly dictate agents’ forecasting, affording them considerable freedom. We leave alternative, stricter definitions of rationality, which may derive from probabilistic conceptions of argument strength, to future work.

**Example 4.** Consider our running Tokyo Olympics example, with the same arguments and relations from Example 3 and an agent *alice* with a confidence score  $C_{\text{alice}}(\mathcal{P}) = -0.5$ . From this we know that *alice* believes that the suggested  $\mathcal{F}^{\mathcal{P}}$  in the proposal argument  $\mathcal{P}$  should be decreased. Then, under our definition of rationality, *alice*’s forecast  $\mathcal{F}_{\text{alice}}^A$  is ‘rational’ if it decreases  $\mathcal{F}^{\mathcal{P}}$  by up to 50%.

If an agent’s forecast  $\mathcal{F}_a^A$  violates these rationality constraints then it is ‘blocked’ and the agent is prompted to return to the argumentation graph. From here, they may carry out one or more of the following actions to render their forecast rational:

- a. Revise their forecast;
- b. Revise their votes on arguments;
- c. Add new arguments (and vote on them).

Whilst a) and b) occur on an agent-by-agent basis, confined to each delegate framework, c) affects the shared update framework and requires special consideration. Each time new arguments are added to the shared graph, every agent must vote on them, even if they have already made a

rational forecast. In certain cases, after an agent has voted on a new argument, it is possible that their rational forecast is made irrational. In this instance, the agent must resolve their irrationality via the steps above. In this way, the update framework can be refined on an iterative basis until the graph is no longer being modified and all agents' forecasts are rational. At this stage, the update framework has reached a stable state and the agents  $\mathcal{A}$  are collectively rational:

**Definition 6.** Given an update framework  $u = \langle \mathcal{P}, \mathcal{X}, \mathcal{X}^-, \mathcal{X}^+, \mathcal{R}^p, \mathcal{R}, \mathcal{A}, \mathcal{V}, \mathcal{F}^A \rangle$ ,  $\mathcal{A}$  is collectively rational (wrt  $u$ ) iff  $\forall a \in \mathcal{A}$ ,  $a$  is individually rational (wrt the delegate framework  $u_a = \langle \mathcal{P}, \mathcal{X}, \mathcal{X}^-, \mathcal{X}^+, \mathcal{R}^p, \mathcal{R}, a, \mathcal{V}_a, \mathcal{F}_a^A \rangle$ ).

When  $\mathcal{A}$  is collectively rational, the update framework  $u$  becomes immutable and the aggregation (defined next) produces a group forecast  $\mathcal{F}^g$ , which becomes the new  $\mathcal{F}$ .

## 4.2 Aggregating Forecasts

After all the agents have made a rational forecast, an aggregation function is applied to produce one collective forecast. One advantage of forecasting debates vis-a-vis the many other forms of debate, is that a ground truth always exists – an event either happens or does not. This means that, over time and after enough FAF instantiations, data on the forecasting success of different agents can be amassed. In turn, the relative historical performance of forecasting agents can inform the aggregation of group forecasts. In update frameworks, a weighted aggregation function based on Brier Scoring (Brier 1950) is used, such that more accurate forecasting agents have greater influence over the final forecast. Brier Scores are a widely used criterion to measure the accuracy of probabilistic predictions, effectively gauging the distance between a forecaster's predictions and an outcome after it has(n't) happened, as follows.

**Definition 7.** Given an agent  $a$ , a non-empty series of forecasts  $\mathcal{F}_a^A(1), \dots, \mathcal{F}_a^A(\mathcal{N}_a)$  with corresponding actual outcomes  $\mathcal{O}_1, \dots, \mathcal{O}_{\mathcal{N}_a} \in \{\text{true}, \text{false}\}$  (where  $\mathcal{N}_a > 0$  is the number of forecasts  $a$  has made in a non-empty sequence of as many update frameworks),  $a$ 's Brier Score  $b_a \in [0, 1]$  is as follows:

$$b_a = \frac{1}{\mathcal{N}_a} \sum_{t=1}^{\mathcal{N}_a} (\mathcal{F}_a^A(t) - \text{val}(\mathcal{O}_t))^2$$

where  $\text{val}(\mathcal{O}_t) = 1$  if  $\mathcal{O}_t = \text{true}$ , and 0 otherwise.

A Brier Score  $b$  is effectively the mean squared error used to gauge forecasting accuracy, where a low  $b$  indicates high accuracy and high  $b$  indicates low accuracy. This can be used in the update framework's aggregation function via the weighted arithmetic mean as follows. Each Brier Score is inverted to ensure that more (less, resp.) accurate forecasters have higher (lower, resp.) weighted influences on  $\mathcal{F}^g$ :

**Definition 8.** Given a set of agents  $\mathcal{A} = \{a_1, \dots, a_n\}$ , their corresponding set of Brier Scores  $b = \{b_{a_1}, \dots, b_{a_n}\}$  and their forecasts  $\mathcal{F}^A = \{\mathcal{F}_{a_1}^A, \dots, \mathcal{F}_{a_n}^A\}$ , and letting, for  $i \in$

$\{1, \dots, n\}$ ,  $w_i = 1 - b_{a_i}$ , the group forecast  $\mathcal{F}^g$  is as follows:

$$\text{if } \sum_{i=1}^n w_i \neq 0 : \quad \mathcal{F}^g = \frac{\sum_{i=1}^n w_i \mathcal{F}_{a_i}^A}{\sum_{i=1}^n w_i};$$

$$\text{otherwise :} \quad \mathcal{F}^g = 0.$$

This group forecast could be 'activated' after a fixed number of debates (with the mean average used prior), when sufficient data has been collected on the accuracy of participating agents, or after a single debate, in the context of our general Forecasting Argumentation Frameworks:

**Definition 9.** A Forecasting Argumentation Framework (FAF) is a triple  $\langle \mathcal{F}, \mathcal{U}, \mathcal{T} \rangle$  such that:

- $\mathcal{F}$  is a forecast;
- $\mathcal{U}$  is a finite, non-empty sequence of update frameworks with  $\mathcal{F}$  the forecast of the proposal argument in the first update framework in the sequence; the forecast of each subsequent update framework is the group forecast of the previous update framework's agents' forecasts;
- $\mathcal{T}$  is a preset time limit representing the lifetime of the FAF;
- each agent's forecast wrt the agent's delegate framework drawn from each update framework is strictly rational.

**Example 5.** Consider our running Tokyo Olympics example: the overall FAF may be composed of  $\mathcal{F} = 0.15$ , update frameworks  $\mathcal{U} = \{u_1, u_2, u_3\}$  and time limit  $\mathcal{T} = 14$  days, where  $u_3$  is the latest (and therefore the only open) update framework after, for example, four days.

The superforecasting literature explores a range of forecast aggregation algorithms: extremizing algorithms (Baron et al. 2014), variations on logistic and Fourier  $L_2E$  regression (Cross et al. 2018), with considerable success. These approaches aim at ensuring that less certain or less accurate forecasts have a lesser influence over the final aggregated forecast. We believe that FAFs apply a more intuitive algorithm since much of the 'work' needed to bypass inaccurate and erroneous forecasting is expedited via argumentation.

## 5 Properties

We now undertake a theoretical analysis of FAFs by considering mathematical properties they satisfy. Note that the properties of the DF-QuAD algorithm (see (Rago et al. 2016)) hold (for the amendment and pro/con arguments) here. For brevity, we focus on novel properties unique to FAFs which relate to our new argument types. These properties focus on aggregated group forecasts wrt a debate (update framework). They imply the two broad, and we posit, desirable, principles of *balance* and *unequal representation*. We assume for this section a generic update framework  $u = \langle \mathcal{P}, \mathcal{X}, \mathcal{X}^-, \mathcal{X}^\uparrow, \mathcal{R}^p, \mathcal{R}, \mathcal{A}, \mathcal{V}, \mathcal{F}^A \rangle$  with group forecast  $\mathcal{F}^g$ .

**Balance.** The intuition for these properties is that differences between  $\mathcal{F}^g$  and  $\mathcal{F}^P$  correspond to imbalances between the *increase* and *decrease* amendment arguments.

The first result states that  $\mathcal{F}^g$  only differs from  $\mathcal{F}^P$  if  $\mathcal{F}^P$  is the dialectical target of amendment arguments.

**Proposition 1.** If  $\mathcal{X} = \emptyset$  ( $\mathcal{X}^\downarrow = \emptyset$  and  $\mathcal{X}^\uparrow = \emptyset$ ), then  $\mathcal{F}^g = \mathcal{F}^P$ .

*Proof.* If  $\mathcal{X}^\downarrow = \emptyset$  and  $\mathcal{X}^\uparrow = \emptyset$  then  $\forall a \in \mathcal{A}, C_a(\mathcal{P}) = 0$  by Def. 4 and  $\mathcal{F}_a^A = \mathcal{F}^P$  by Def. 5. Then,  $\mathcal{F}^g = \mathcal{F}^P$  by Def. 8.  $\square$

This simple proposition conveys an important property for forecasting: for an agent to put forward a different forecast, amendment arguments must have been introduced.

**Example 6.** *In the Olympics setting, the group of agents could only forecast higher or lower than the proposed forecast  $\mathcal{F}^P$  after the addition of at least one of the amendment arguments  $x_1^\downarrow, x_2^\downarrow$  or  $x_1^\uparrow$ .*

In the absence of increase (decrease) amendment arguments, if there are decrease (increase, resp.) amendment arguments, then  $\mathcal{F}^g$  is not higher (lower, resp.) than  $\mathcal{F}^P$ .

**Proposition 2.** *If  $\mathcal{X}^\downarrow \neq \emptyset$  and  $\mathcal{X}^\uparrow = \emptyset$ , then  $\mathcal{F}^g \leq \mathcal{F}^P$ . If  $\mathcal{X}^\downarrow = \emptyset$  and  $\mathcal{X}^\uparrow \neq \emptyset$ , then  $\mathcal{F}^g \geq \mathcal{F}^P$ .*

*Proof.* If  $\mathcal{X}^\downarrow \neq \emptyset$  and  $\mathcal{X}^\uparrow = \emptyset$  then  $\forall a \in \mathcal{A}, C_a(\mathcal{P}) \leq 0$  by Def. 4 and then  $\mathcal{F}_a^A \leq \mathcal{F}^P$  by Def. 5. Then, by Def. 8,  $\mathcal{F}^g \leq \mathcal{F}^P$ . If  $\mathcal{X}^\downarrow = \emptyset$  and  $\mathcal{X}^\uparrow \neq \emptyset$  then  $\forall a \in \mathcal{A}, C_a(\mathcal{P}) \geq 0$  by Def. 4 and then  $\mathcal{F}_a^A \geq \mathcal{F}^P$  by Def. 5. Then, by Def. 8,  $\mathcal{F}^g \geq \mathcal{F}^P$ .  $\square$

This proposition demonstrates that, if a decrease (increase) amendment argument has an effect on proposal arguments, it can only be as its name implies.

**Example 7.** *In the Olympics setting, the agents could not forecast higher than the proposed forecast  $\mathcal{F}^P$  if either of the decrease amendment arguments  $x_1^\downarrow$  or  $x_2^\downarrow$  had been added, but the increase argument  $x_1^\uparrow$  had not. Likewise, the agents could not forecast lower than  $\mathcal{F}^P$  if  $x_1^\uparrow$  had been added, but neither of  $x_1^\downarrow$  or  $x_2^\downarrow$  had.*

If  $\mathcal{F}^g$  is lower (higher) than  $\mathcal{F}^P$ , there is at least one decrease (increase, resp.) argument.

**Proposition 3.** *If  $\mathcal{F}^g < \mathcal{F}^P$ , then  $\mathcal{X}^\downarrow \neq \emptyset$ . If  $\mathcal{F}^g > \mathcal{F}^P$ , then  $\mathcal{X}^\uparrow \neq \emptyset$ .*

*Proof.* If  $\mathcal{F}^g < \mathcal{F}^P$  then, by Def. 8,  $\exists a \in \mathcal{A}$  where  $\mathcal{F}_a^A < \mathcal{F}^P$ , for which it holds from Def. 5 that  $C_a(\mathcal{P}) < 0$ . Then, irrespective of  $\mathcal{X}^\uparrow, \mathcal{X}^\downarrow \neq \emptyset$ . If  $\mathcal{F}^g > \mathcal{F}^P$  then, by Def. 8,  $\exists a \in \mathcal{A}$  where  $\mathcal{F}_a^A > \mathcal{F}^P$ , for which it holds from Def. 5 that  $C_a(\mathcal{P}) > 0$ . Then, irrespective of  $\mathcal{X}^\downarrow, \mathcal{X}^\uparrow \neq \emptyset$ .  $\square$

We can see here that the only way an agent can decrease (increase) the forecast is by adding decrease (increase, resp.) arguments, ensuring the debate is structured as intended.

**Example 8.** *In the Olympics setting, the group of agents could only produce a group forecast  $\mathcal{F}^g$  lower than  $\mathcal{F}^P$  due to the presence of decrease amendment arguments  $x_1^\downarrow$  or  $x_2^\downarrow$ . Likewise, the group of agents could only produce a  $\mathcal{F}^g$  higher than  $\mathcal{F}^P$  due to the presence of  $x_1^\uparrow$ .*

**Unequal representation.** AFs exhibit instances of unequal representation in the final voting process. In formulating the following properties, we distinguish between two forms of unequal representation. First, *dictatorship*, where a single agent dictates  $\mathcal{F}^g$  with no input from other agents. Second, *pure oligarchy*, where a group of agents dictates

$\mathcal{F}^g$  with no input from other agents outside the group. In the forecasting setting, these properties are desirable as they guarantee higher accuracy from the group forecast  $\mathcal{F}^g$ .

An agent with a forecasting record of *some* accuracy exercises *dictatorship* over the group forecast  $\mathcal{F}^g$ , if the rest of the participating agents have a record of total inaccuracy.

**Proposition 4.** *If  $a_d \in \mathcal{A}$  has a Brier score  $b_{a_d} < 1$  and  $\forall a_z \in \mathcal{A} \setminus \{a_d\}, b_{a_z} = 1$ , then  $\mathcal{F}^g = \mathcal{F}_{a_d}^A$ .*

*Proof.* By Def. 8: if  $b_{a_z} = 1 \forall a_z \in \mathcal{A} \setminus \{a_d\}$ , then  $w_{a_z} = 0$ ; and if  $b_{a_d} < 1$ , then  $w_{a_d} > 0$ . Then, again by Def. 8,  $\mathcal{F}_{a_d}^A$  is weighted at 100% and  $\mathcal{F}_{a_z}^A$  is weighted at 0% so  $\mathcal{F}^g = \mathcal{F}_{a_d}^A$ .  $\square$

This proposition demonstrates how we will disregard agents with total inaccuracy, even in the extreme case where we allow one (more accurate) agent to dictate the forecast.

**Example 9.** *In the running example, if alice, bob and charlie have Brier scores of 0.5, 1 and 1, resp., bob's and charlie's forecasts have no impact on  $\mathcal{F}^g$ , whilst alice's forecast becomes the group forecast  $\mathcal{F}^g$ .*

A group of agents with a forecasting record of *some* accuracy exercises *pure oligarchy* over  $\mathcal{F}^g$  if the rest of the agents all have a record of total inaccuracy.

**Proposition 5.** *Let  $\mathcal{A} = \mathcal{A}_o \cup \mathcal{A}_z$  where  $\mathcal{A}_o \cap \mathcal{A}_z = \emptyset$ ,  $b_{a_o} < 1 \forall a_o \in \mathcal{A}_o$  and  $b_{a_z} = 1 \forall a_z \in \mathcal{A}_z$ . Then,  $\mathcal{F}_{a_o}^A$  is weighted at  $> 0\%$  and  $\mathcal{F}_{a_z}^A$  is weighted at  $0\%$ .*

*Proof.* By Def. 8: if  $b_{a_z} = 1 \forall a_z \in \mathcal{A}_z$ , then  $w_{a_z} = 0$ ; and if  $b_{a_o} < 1 \forall a_o \in \mathcal{A}_o$ , then  $w_{a_o} > 0$ . Then, again by Def. 8,  $\mathcal{F}_{a_o}^A$  is weighted at  $> 0\%$  and  $\mathcal{F}_{a_z}^A$  is weighted at  $0\%$ .  $\square$

This proposition extends the behaviour from Proposition 4 to the (more desirable) case where fewer agents have a record of total inaccuracy.

**Example 10.** *In the running example, if alice, bob and charlie have Brier scores of 1, 0.2 and 0.6, resp., alice's forecast has no impact on  $\mathcal{F}^g$ , whilst bob and charlie's aggregated forecast becomes the group forecast  $\mathcal{F}^g$ .*

## 6 Evaluation

We conducted an experiment using a dataset obtained from the 'Superforecasting' project, Good Judgment Inc (Good Judgment Inc 2014), to simulate four past forecasting debates in FAFs. This dataset contained 1770 datapoints (698 'forecasts' and 1072 'comments') posted by 242 anonymised users with a range of expertise. The original debates had occurred on the publicly available group forecasting platform, the Good Judgment Open (GJO)<sup>2</sup>, providing a suitable baseline against which to compare FAFs' accuracy.

For the experiment, we used a prototype implementation of FAFs in the form of the publicly available web platform called *Arg&Forecast* (see (Irwin, Rago, and Toni 2022) for an introduction to the platform and an additional human experiment with FAFs). Python's Gensim topic modelling library (Řehůřek and Sojka 2010) was used to separate the datapoints for each debate into contextual-temporal groups

<sup>2</sup><https://www.gjopen.com/>

Q	Group $b$	$min(b)$	$max(b)$
Q1	0.1013 (0.1187)	0.0214 (0)	0.4054 (1)
Q2	0.216 (0.1741)	0 (0)	0.3853 (1)
Q3	0.01206 (0.0227)	0.0003 (0)	0.0942 (0.8281)
Q4	0.5263 (0.5518)	0 (0)	0.71 (1)
<b>All</b>	<b>0.2039 (0.217)</b>	<b>0 (0)</b>	<b>1 (1)</b>

Table 2: The accuracy of the platform group versus control, where ‘Group  $b$ ’ is the aggregated (mean) Brier score, ‘ $min(b)$ ’ is the lowest individual Brier score and ‘ $max(b)$ ’ is the highest individual Brier score. Q1-Q4 indicate the four simulated debates.

that could form update frameworks. In each update framework the proposal forecast was set to the mean average of forecasts made in the update framework window and each argument appeared only once. Gensim was further used to simulate voting, matching users to specific arguments they (dis)approved of. Some 4,700 votes were then generated with a three-valued system (where votes were taken from  $\{0,0.5,1\}$ ) to ensure consistency: if a user voiced approval for an argument in the debate time window, their vote for the corresponding argument(s) was set to 1; disapproval for an argument led to a vote of 0, and (in the most common case) if a user did not mention an argument at all, their vote for the corresponding argument(s) defaulted to 0.5.

With the views of all participating users wrt the proposal argument encoded in each update framework’s votes, forecasts could then be simulated. If a forecast was irrational, violating any of the three constraints in Def. 5 (referred to in the following as *increase*, *decrease* and *scale*, resp.), it was blocked and, to mimic real life use, an automatic ‘follow up’ forecast was made. The ‘follow up’ forecast would be the closest possible prediction (to their original choice) a user could make whilst remaining ‘rational’.

Note that evaluation of the aggregation function described in §4.2 was outside this experiment, since the past forecasting accuracy of the dataset’s 242 anonymised users was unavailable. Instead, we used the mean average whilst adopting the GJO’s method for scoring the accuracy of a user and/or group over the lifetime of the question (Roesch 2015). This meant calculating a daily forecast and daily Brier score for each user, for every day of the question. After users made their first rational forecast, that forecast became their ‘daily forecast’ until it was updated with a new forecast. Average and range of daily Brier scores allowed reliable comparison between (individual and aggregated) performance of the GJO versus the FAF implementation.

**Results.** As Table 2 demonstrates, simulating forecasting debates from GJO in *Arg&Forecast* led to predictive accuracy improvements in three of the four debates. This is reflected in these debates by a substantial reduction in Brier scores versus control. The greatest accuracy improvement in absolute terms was in Q4, which saw a Brier score decrease of 0.0255. In relative terms, Brier score decreases ranged from 5% (Q4) to 47% (Q3). The average Brier score decrease was 33%, representing a significant improvement in

Q	$\bar{c}$	Forecasts	Irrational Forecasts		
			<i>Increase</i>	<i>Decrease</i>	<i>Scale</i>
Q1	-0.0418	366	63	101	170
Q2	0.1827	84	11	15	34
Q3	-0.4393	164	53	0	86
Q4	0.3664	84	4	19	15
All	-0.0891	698	131	135	305

Table 3: Auxiliary results from the experiment, where  $\bar{c}$  is the average confidence score, ‘Forecasts’ is number of forecasts made in each question and ‘Irrational Forecasts’ the number in each question which violated each constraint in §4.1.

forecasting accuracy across the board. Table 3 demonstrates how our rationality constraints drove forward this improvement. 82% of forecasts made across the four debates were classified as irrational and subsequently moderated with a rational ‘follow up’ forecast. Notably, there were more *irrational scale* forecasts than *irrational increase* and *irrational decrease* forecasts combined. These results demonstrate how argumentation-based rationality constraints can play an active role in facilitating higher forecasting accuracy, signalling the early promise of FAFs.

## 7 Conclusions

We have introduced the Forecasting Argumentation Framework (FAF), a multi-agent argumentation framework which supports forecasting debates and probability estimates. FAFs are composite argumentation frameworks, comprised of multiple non-concurrent update frameworks which themselves depend on three new argument types and a novel definition of rationality for the forecasting context. Our theoretical and empirical evaluation demonstrates the potential of FAFs, namely in increasing forecasting accuracy, holding intuitive properties, identifying irrational behaviour and driving higher engagement with the forecasting question (more arguments and responses, and more forecasts in the user study). These strengths align with requirements set out by previous research in the field of judgmental forecasting.

There is a multitude of possible directions for future work. First, FAFs are equipped to deal only with two-valued outcomes but, given the prevalence of forecasting issues with multi-valued outcomes (e.g. ‘Who will win the next UK election?’), expanding their capability would add value. Second, further work may focus on the rationality constraints, e.g. by introducing additional parameters to adjust their strictness, or by implementing alternative interpretations of rationality. Third, future work could explore constraining agents’ argumentation. This could involve using past Brier scores to limit the quantity or strength of agents’ arguments and also to give them greater leeway wrt the rationality constraints. Fourth, our method relies upon acyclic graphs: we believe that they are intuitive for users and note that all Good Judgment Open debates were acyclic; nonetheless, the inclusion of cyclic relations (e.g. to allow con arguments that attack each other) could expand the scope of the argumentative reasoning in FAFs. Finally, there is an immediate need for larger scale human experiments.

## Acknowledgements

The authors would like to thank Prof. Anthony Hunter for his helpful contributions to discussions in the build up to this work. Special thanks, in addition, go to Prof. Philip E. Tetlock and the Good Judgment Project team for their warm cooperation and for providing datasets for the experiments. Finally, the authors would like to thank the anonymous reviewers and meta-reviewer for their suggestions, which led to a significantly improved paper.

## References

- Amgoud, L., and Ben-Naim, J. 2016. Evaluation of arguments from support relations: Axioms and semantics. In *IJCAI*, 900–906.
- Amgoud, L., and Ben-Naim, J. 2018. Weighted bipolar argumentation graphs: Axioms and semantics. In *IJCAI*, 5194–5198.
- Amgoud, L.; Ben-Naim, J.; Doder, D.; and Vesic, S. 2017. Acceptability semantics for weighted argumentation frameworks. In *IJCAI*, 56–62.
- Armstrong, J. S. 2001. *Evaluating Forecasting Methods*. Boston, MA: Springer US. 443–472.
- Arvan, M.; Fahimnia, B.; Reisi, M.; and Siemsen, E. 2019. Integrating human judgement into quantitative forecasting methods: A review. *Omega (United Kingdom)* 86.
- Atkinson, K.; Baroni, P.; Giacomini, M.; Hunter, A.; Prakken, H.; Reed, C.; Simari, G. R.; Thimm, M.; and Villata, S. 2017. Towards artificial argumentation. *AI Magazine* 38(3):25–36.
- Baets, S. D., and Harvey, N. 2020. Using judgment to select and adjust forecasts from statistical models. *European Journal of Operational Research* 284(3):882–895.
- Baron, J.; Mellers, B. A.; Tetlock, P. E.; Stone, E.; and Ungar, L. H. 2014. Two reasons to make aggregated probability forecasts more extreme. *Decision Analysis* 11.
- Baroni, P.; Romano, M.; Toni, F.; Aurisicchio, M.; and Bertanza, G. 2015. Automatic evaluation of design alternatives with quantitative argumentation. *Argument and Computation* 6.
- Baroni, P.; Gabbay, D.; Giacomini, M.; and van der Torre, L., eds. 2018. *Handbook of Formal Argumentation*. College Publications.
- Baroni, P.; Rago, A.; and Toni, F. 2018. How many properties do we need for gradual argumentation? In *AAAI*, 1736–1743.
- Besnard, P.; García, A. J.; Hunter, A.; Modgil, S.; Prakken, H.; Simari, G. R.; and Toni, F. 2014. Introduction to structured argumentation. *Argument and Computation* 5(1):1–4.
- Bessette, J. M. 1980. Deliberative democracy: The majority principle in republican government. In *How Democratic is the Constitution?* Washington: American Enterprise Institute for Public Policy Research. 102–116.
- Brier, G. W. 1950. Verification of forecasts expressed in terms of probability. *Monthly Weather Review* 78.
- Chang, W.; Chen, E.; Mellers, B.; and Tetlock, P. E. 2016. Developing expert political judgment: The impact of training and practice on judgmental accuracy in geopolitical forecasting tournaments. *Judgment and Decision Making* 11.
- Cobo, M. L.; Martínez, D. C.; and Simari, G. R. 2010. On admissibility in timed abstract argumentation frameworks. In *ECAI*, 1007–1008.
- Cobo, M. L.; Martínez, D. C.; and Simari, G. R. 2011. Stable extensions in timed argumentation frameworks. In *TAFI*, 181–196.
- Cross, D.; Ramos, J.; Mellers, B.; Tetlock, P. E.; and Scott, D. W. 2018. Robust forecast aggregation: Fourier l2e regression. *Journal of Forecasting* 37.
- Cyras, K.; Letsios, D.; Misener, R.; and Toni, F. 2019. Argumentation for explainable scheduling. In *AAAI*, 2752–2759.
- Cyras, K.; Rago, A.; Albini, E.; Baroni, P.; and Toni, F. 2021. Argumentative XAI: A survey. In *IJCAI*, 4392–4399.
- Delbecq, A.; Van den Ven, A.; and Gustafson, D. 1975. *Group Techniques for Program Planning: A Guide to Nominal Group and Delphi Processes*. Scott Foresman.
- Dung, P. M., and Thang, P. M. 2010. Towards (probabilistic) argumentation for jury-based dispute resolution. In *COMMA*, 171–182.
- 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.
- Evripidou, V., and Toni, F. 2012. Argumentation and voting for an intelligent user empowering business directory on the web. In *RR*, 209–212.
- Fazzinga, B.; Flesca, S.; and Furfaro, F. 2018. Probabilistic bipolar abstract argumentation frameworks: complexity results. In *IJCAI*, 1803–1809.
- Good Judgment Inc. 2014. <https://goodjudgment.com/>. [Online; accessed 09-October-2021].
- Goodwin, P.; Gönül, M. S.; and Önköl-Atay, D. 2019. When providing optimistic and pessimistic scenarios can be detrimental to judgmental demand forecasts and production decisions. *European Journal of Operational Research* 273(3):992–1004.
- Green, K. C., and Armstrong, J. S. 2007. Structured analogies for forecasting. *International Journal of Forecasting* 23.
- Horowitz, M. C.; Ciocca, J.; Kahn, L.; and Ruhl, C. 2021. Keeping score: A new approach to geopolitical forecasting. *Penn Global*.
- Hunter, A., and Thimm, M. 2014. Probabilistic argumentation with incomplete information. In *ECAI*, 1033–1034.
- Hunter, A.; Polberg, S.; and Thimm, M. 2020. Epistemic graphs for representing and reasoning with positive and negative influences of arguments. *Artificial Intelligence* 281.
- Hunter, A. 2013. A probabilistic approach to modelling uncertain logical arguments. *International Journal of Approximate Reasoning* 54.
- Irwin, B.; Rago, A.; and Toni, F. 2022. Argumentative forecasting – extended abstract. In *AAMAS*. (To appear).

- Kahneman, D. 2012. *Thinking, fast and slow*. Farrar, Straus and Giroux.
- Katsagounos, I.; Thomakos, D. D.; Litsiou, K.; and Nikolopoulos, K. 2021. Superforecasting reality check: Evidence from a small pool of experts and expedited identification. *European Journal of Operational Research* 289(1):107–117.
- Kotonya, N., and Toni, F. 2020. Explainable automated fact-checking for public health claims. In *EMNLP*, 7740–7754.
- Landeta, J.; Barrutia, J.; and Lertxundi, A. 2011. Hybrid delphi: A methodology to facilitate contribution from experts in professional contexts. *Technological Forecasting and Social Change* 78.
- Lawrence, M., and Makridakis, S. 1989. Factors affecting judgmental forecasts and confidence intervals. *Organizational Behavior and Human Decision Processes* 43.
- Lawrence, M.; Goodwin, P.; O'Connor, M.; and Önkal, D. 2006. Judgmental forecasting: A review of progress over the last 25 years. *International Journal of Forecasting* 22.
- Leite, J., and Martins, J. G. 2011. Social abstract argumentation. In *IJCAI*, 2287–2292.
- Li, H.; Oren, N.; and Norman, T. J. 2011. Probabilistic argumentation frameworks. In *TAF*, 1–16.
- Linstone, H., and Turoff, M. 1975. *The Delphi Method: Techniques and Applications*. Addison-Wesley Publishing Company.
- Litsiou, K.; Polychronakis, Y.; Karami, A.; and Nikolopoulos, K. 2019. Relative performance of judgmental methods for forecasting the success of megaprojects. *International Journal of Forecasting*.
- MacGregor, D. G., and Armstrong, J. 1994. Judgmental decomposition: when does it work? *International Journal of Forecasting* 10(4):495–506.
- Makridakis, S.; Hogarth, R. M.; and Gaba, A. 2010. Why forecasts fail. what to do instead. *MIT Sloan Management Review* 51.
- McNees, S. K. 1987. Consensus forecasts: tyranny of the majority? *New England Economic Review* 15-21.
- Mellers, B.; Stone, E.; Atanasov, P.; Rohrbaugh, N.; Metz, S. E.; Ungar, L.; Bishop, M. M.; Horowitz, M.; Merkle, E.; and Tetlock, P. E. 2015a. The psychology of intelligence analysis: Drivers of prediction accuracy in world politics. *Journal of Experimental Psychology: Applied* 21.
- Mellers, B.; Stone, E.; Murray, T.; Minster, A.; Rohrbaugh, N.; Bishop, M.; Chen, E.; Baker, J.; Hou, Y.; Horowitz, M.; Ungar, L.; and Tetlock, P. E. 2015b. Identifying and cultivating superforecasters as a method of improving probabilistic predictions. *Perspectives on Psychological Science* 10.
- Nikolopoulos, K.; Litsa, A.; Petropoulos, F.; Bougioukos, V.; and Khammash, M. 2015. Relative performance of methods for forecasting special events. *Journal of Business Research* 68.
- Nikolopoulos, K. 2021. We need to talk about intermittent demand forecasting. *European Journal of Operational Research* 291(2):549–559.
- Page, M.; Aiken, C.; and Murdick, D. 2020. How crowd forecasting can inform the big picture. *Georgetown CSET*.
- Petropoulos, F.; Fildes, R.; and Goodwin, P. 2016. Do 'big losses' in judgmental adjustments to statistical forecasts affect experts' behaviour? *European Journal of Operational Research* 249(3):842–852.
- Rago, A., and Toni, F. 2017. Quantitative argumentation debates with votes for opinion polling. In *PRIMA*, 369–385.
- Rago, A.; Toni, F.; Aurisicchio, M.; and Baroni, P. 2016. Discontinuity-free decision support with quantitative argumentation debates. In *KR*, 63–73.
- Řehůřek, R., and Sojka, P. 2010. Software Framework for Topic Modelling with Large Corpora. In *LREC*, 45–50.
- Roesch, B. 2015. Measuring accuracy in prediction markets and opinion. <https://www.cultivatelabs.com/posts/measuring-accuracy-in-prediction-markets-and-opinion-poll-pools>. [Online; accessed 05-February-2022].
- Tetlock, P. E., and Gardner, D. 2016. *Superforecasting: the art and science of prediction*. Crown Publishers.
- Tetlock, P. E.; Mellers, B. A.; Rohrbaugh, N.; and Chen, E. 2014. Forecasting tournaments: Tools for increasing transparency and improving the quality of debate. *Current Directions in Psychological Science* 23.
- Tetlock, P. E. 2017. *Expert political judgment: How good is it? How can we know?* Princeton University Press.
- Thimm, M. 2012. A probabilistic semantics for abstract argumentation. In *ECAI*, 750–755.
- Zellner, M.; Abbas, A. E.; Budescu, D. V.; and Galstyan, A. 2021. A survey of human judgement and quantitative forecasting methods. *Royal Society open science* 8(2):201187.