Algebraic Foundations for Non-Monotonic Practical Reasoning

Nourhan Ehab[1] , **Haythem O. Ismail**[2,1]

[1] Department of Computer Science and Engineering, German University in Cairo

[2] Department of Engineering Mathematics, Cairo University

{nourhan.ehab, haythem.ismail}@guc.edu.eg

Abstract

Practical reasoning is a hallmark of human intelligence. We are confronted everyday with situations that require us to meticulously choose among our possibly conflicting desires, and we usually do so with ease guided by beliefs which may be uncertain or even contradictory. The desires we end up choosing to pursue make up intentions which we seamlessly revise whenever our beliefs change. Modelling the intricate process of practical reasoning has attracted a lot of attention in the KR community giving rise to a wide array of logics coming in diverse flavours. However, a robust logic of practical reasoning with adequate semantics representing the preferences among the agent's different mental attitudes, while capturing the intertwined revision of beliefs and intentions, remains missing. In this paper, we aspire to fill this gap by introducing general algebraic foundations for practical reasoning. We present an algebraic logic we refer to as $Log_A\mathbf{PR}$ capable of modelling preferences among the agent's different mental attitudes and capturing their joint revision.

1 Introduction

"What should I do?" is a question that repeatedly poses itself in our everyday lives. The process of reflection and deliberation we undergo to answer this question is what we refer to as practical reasoning (Broome 2002). To demonstrate the intricate process of practical reasoning, consider the following example.

Example 1. ***The Weekend Dilemma***
Ted needs to decide what to do over the weekend. He has to work on a long overdue presentation as his boss will be really mad if Ted does not give the presentation by latest the beginning of next week. Ted also had previous plans with his best friend Marshall to go on a hunting trip during the weekend. Ted thinks that he can go to the trip and dedicate some time to work on the presentation there. Barney, Ted's other best friend who is currently in a fall out with Marshall, told Ted that he heard from his fiancé Robin that the trip location has no internet connectivity so Ted will not be able to work on the presentation there. Ted trusts that Robin usually tells the truth, but suspects that Barney might be lying to make him not go on the trip with Marshall. Ted desires to go to the trip, but he still wishes he desired to not make his boss mad. As Ted wants to start being more rational and responsible, he prefers to give up his desires or obligations whenever they conflict with his beliefs and prefers to give up his desires whenever they conflict with his obligations. What should Ted do?

Since the days of Aristotle in the twelfth century BC, modelling practical reasoning has posed a difficult challenge for philosophers, logicians, and computer scientists alike. Several attempts have been made over the years to come up with logical theories of practical reasoning; however, a comprehensive and adequate theory remains missing (Thomason 2018). Some endeavours at modelling practical reasoning have been successful when the problem was viewed as a mean to model rational agency. In this view, rational agents are thought of as practical reasoners that act based on their *beliefs* about the world and driven by their *desires* (Rao and Wooldridge 1999; Searle 2003). The action-attitudes of the agent representing its commitment to some motivations, as permissible by its beliefs, are classically referred to as *intentions* (Bratman 1987; Cohen and Levesque 1990). In this way, the agent's intentions are evidently dependent on both its beliefs and desires. Taking the trinity of beliefs, desires, and intentions to be the key elements of the agent's mental state is the approach taken by the much renowned BDI model of rational agents (Rao and Georgeff 1995) and its extensions to include other mental attitudes such as obligations (Broersen et al. 2001; Broersen et al. 2002).

In practical settings, the agent's beliefs and desires are often governed by a system of preferences and are continuously revised. Consequently, the revision of beliefs and desires must be reflected on the agent's intentions. The existing logical approaches to modelling preferences within the BDI architecture are the graded-BDI (g-BDI) model (Casali, Godo, and Sierra 2008; Casali, Godo, and Sierra 2011) and TEAMLOG (Dunin-Keplicz, Nguyen, and Szalas 2010). While both approaches propose frameworks for joint reasoning with graded beliefs, desires, and intentions; neither has an account for the joint *revision* of the three mental attitudes. Moreover, the g-BDI model lacks precise semantics and TEAMLOG is based on a normal modal logic providing only a third-person account of reasoning *about* the mental attitudes. On the other hand, the joint revision of beliefs and intentions has been attempted in (Shoham 2009; Icard, Pacuit, and Shoham 2010). These theories, however, do not account for desire or preferences over beliefs and intentions. In this paper, we aspire to address this gap in the literature. Our contribution is twofold. First, we introduce general algebraic foundations for first-person practical reasoning with several mental attitudes where preference and joint revision can be captured. Second, we provide precise semantics for an algebraic logic we refer to as $Log_A\mathbf{PR}$ for joint reasoning with graded beliefs and motivations. "Log" stands for logic, "A" for algebraic, and "$\mathbf{PR}$" for practical reasoning. The grades associated with the beliefs and motivations in $Log_A\mathbf{PR}$ are reified and are taken to represent measures of trust or preference. In $Log_A\mathbf{PR}$, we deviate from the BDI model and its extensions in (at least) two ways: (i) we replace the notion of desire with a more general notion of *motivation* to encompass all the different types of motivational attitudes a rational agent can have including (but not limited to) desires, obligations, and social norms; and (ii) we follow (Castelfranchi and Paglieri 2007; Cohen and Levesque 1990) and treat intention as a mental attitude derived from belief and motivation rather than treating it as a basic attitude.

The rest of the paper is structured as follows. In Section 2, we present the motivations behind employing a non-classical logic like $Log_A\mathbf{PR}$ by highlighting its different capabilities. Since we are taking the algebraic route, we review in Section 3 foundational concepts of Boolean algebra on which $Log_A\mathbf{PR}$ will be based. We also generalize the classical notion of filters in Boolean algebra into what we will refer to as *multifilters* providing a generalized algebraic treatment of reasoning with multiple mental attitudes. Next, in Section 4, we present the syntax and semantics of $Log_A\mathbf{PR}$. In Section 5, we extend multifilters to accommodate reasoning with *graded* beliefs and motivations. Additionally, we present our extended graded consequence relation representing the joint reasoning from graded beliefs and motivations to intentions. Finally, in Section 6 we outline some concluding remarks.

2 Why $Log_A\mathbf{PR}$?

$Log_A\mathbf{PR}$ is the most recent addition to a growing family of algebraic logics (Ismail 2012; Ismail 2013; Ismail 2020; Ehab and Ismail 2020). As such, it is essential for a treatment of practical reasoning within the algebraic framework. Hence, independent motivations for the algebraic approach are also motivations for $Log_A\mathbf{PR}$. Such motivations do exist, and are detailed in (Ismail 2012; Ismail 2013; Ismail 2020; Ehab and Ismail 2020). Furthermore, $Log_A\mathbf{PR}$ is a generalization of $Log_A\mathbf{G}$ which is an algebraic logic we presented earlier for non-monotonic reasoning about graded beliefs. As proven in (Ehab and Ismail 2018; Ehab and Ismail 2019), $Log_A\mathbf{G}$ can capture a wide array of non-monotonic reasoning formalisms such as possibilistic logic, circumscription, default logic, autoepistemic logic, and the principle of negation as failure. Thus, $Log_A\mathbf{G}$ can be considered a unifying framework for non-monotonicity. $Log_A\mathbf{PR}$ naturally inherits all the features of $Log_A\mathbf{G}$ yielding a very powerful system of practical reasoning.

In what follows, we briefly present the different features of $Log_A\mathbf{PR}$ and motivate why they are needed by referring to the introductory example. To the best of our knowledge, there does not exist a formalism for practical reasoning that possesses all the following capabilities of $Log_A\mathbf{PR}$.

1. $Log_A\mathbf{PR}$ is a graded logic. The use of graded propositions in $Log_A\mathbf{PR}$ allows the representation of preferences among the agent's beliefs and motivations. This is useful to represent Ted's different trust degrees in his own supposition that he can go to the trip and work on the presentation and the contradicting assertion attributed to Robin by Barney. Further, preferences among Ted's different motivations can likewise be represented.
2. The nesting of graded propositions in $Log_A\mathbf{PR}$ admits the representation of nested graded beliefs and motivations. This naturally facilitates the representation of information acquired by Ted through a chain of sources (Barney and Robin) with different trust degrees. Permitting the nesting of graded motivations facilitates the representation of higher-order desires, first introduced in (Frankfurt 1988), which is useful for representing Ted's wish to desire not to make his boss mad, for instance.
3. Different scales of graded motivations can be represented in $Log_A\mathbf{PR}$. Having separate scales is useful for modelling agents with contradicting motivations and allows us to circumvent several paradoxes of deontic logic as suggested by (Ismail 2020). Two scales, personal desire and obligation, are needed to account for Ted's desire to go to the trip and his obligation towards working on the presentation. We also provide an account for modelling the *characters* of artificial agents as an ordering over their beliefs and motivation scales. For example, a hedonistic agent will always prefer to pursue its desires over its obligations while a selfless agent will always prefer to pursue its obligations over its desires. Whenever contradictions among the agent's motivations arise, they are resolved by alluding to the grades of the conflicting motivations in addition to the agent's character. In our example, Ted's character is represented as his preference to pursue his obligations whenever they conflict with his desires.
4. The precise semantics of $Log_A\mathbf{PR}$ account for joint reasoning with, and revision of, graded beliefs and motivations. We follow (Rao and Georgeff 1995) and refer to the subset of consistent motivations the agent chooses to pursue as its intentions.

At this point, questions about the grades associated to beliefs and motivations may occur to the reader. The set of grades in $Log_A\mathbf{PR}$ can be any totally ordered set of (reified) particulars. The grades may be numeric or may merely be locations on a *qualitative* scale. As such, only the order among the grades is significant and not their actual nature. The assignment of grades to particular beliefs and motivations is out of the scope of this paper; we briefly remark on this, however. The grades come from the same knowledge source that provides the beliefs and motivations themselves. If the latter are provided by a knowledge engineer, for example, then so must the former be. This should not complicate the task of the knowledge engineer as several studies showed that domain experts are often quite good at setting and subjectively assessing numbers to be used as grades for the beliefs and motivations (Charniak 1991). Alternatively, if the beliefs and motivations are learned by some machine learning procedure, then the, typically numeric, grades can be learned as well. Several attempts for accomplishing this are suggested in (Fern 2010; Paccanaro and Hinton 2001; Richardson and Domingos 2006; Vovk, Gammerman, and Shafer 2005). It is also worth pointing out that any difficulty resulting from the task of assigning grades is a price one is bound to pay to account for non-monotonic reasoning. It can be argued that similar equivalently challenging tasks arise in other non-graded non-monotonic formalisms. For instance, how do we set the priorities among the default rules when using prioritized default logic? It might even be that using quantitative grades simplifies the problem as there are several well-defined computational approaches to setting the grades as we previously pointed out.

3 Boolean Algebras and Multifilters

In this section we lay the algebraic foundations on which $Log_A\mathbf{PR}$ is based. We start by reviewing the algebraic concepts of Boolean algebras and filters underlying classical logic, then we extend the notion of filters to accommodate a practical logic of multiple mental attitudes.

A Boolean algebra is a sextuple $\mathfrak{A} =< \mathcal{P}, +, \cdot, -, \bot, \top >$ where $\mathcal{P}$ is a non-empty set with $\{\bot, \top\} \subseteq \mathcal{P}$. $\mathfrak{A}$ is closed under the two binary operators $+$ and $\cdot$ and the unary operator $-$ with commutativity, associativity, absorption, and complementation properties as detailed in (Sankappanavar and Burris 1981). For the purposes of this paper, we will take the elements of $\mathcal{P}$ to be propositions and the operators $+$, $\cdot$, and $-$ to to be disjunction, conjunction, and negation, respectively.

The following definition of *filters* is an essential notion of Boolean algebras to represent an algebraic counterpart to logical consequence. Filters are defined in pure algebraic terms, without alluding to the notion of truth, by utilizing the natural lattice order $\leq$ on the the algebra: for $p_1, p_2 \in \mathcal{P}$, $p_1 \leq p_2 =_{def} p1 \cdot p2 = p1$. Henceforth, $\mathfrak{A}$ is a Boolean algebra $\langle \mathcal{P}, +, \cdot, -, \bot, \top \rangle$.

Definition 3.1. *A filter of $\mathfrak{A}$ is a subset F of $\mathcal{P}$ where*

1. $\top \in F$;
2. *If $a, b \in F$, then $a \cdot b \in F$; and*
3. *If $a \in F$ and $a \leq b$, then $b \in F$.*

The filter generated by $\mathcal{Q} \subseteq \mathcal{P}$ is the smallest filter $F(\mathcal{Q})$ of which $\mathcal{Q}$ is a subset.

Since practical reasoning typically involves joint reasoning with *multiple* mental attitudes (beliefs, motivations, intentions, wishes, etc.), we extend the notion of filters giving rise to what we will refer to as *multifilters*. In contrast to classical filters that rely on the natural order $\leq$ on the Boolean algebra, multifilters will rely on an order on tuples. (Recall that $\leq$ is the classical lattice order.)

Definition 3.2. *Let k be a positive integer. A k partial-order on $\mathfrak{A}$ is a partial order $\preceq_k$ on $\mathcal{P}^k$ such that, $(a_1, \ldots, a_k) \preceq_k (b_1, \ldots, b_k)$ and $b_i = \bot$, for some $1 \leq i \leq k$, only if $a_j = \bot$, for some $1 \leq j \leq k$. Further, we say that $\preceq_k$ is* classical *in i if and only if: (i) if $(a_1, ..., a_k) \preceq_k (b_1, ..., b_k)$ then $a_i \leq b_i$ and (ii) if $a \leq b$ then $(\{\top\}^{i-1} \times \{a\} \times \{\top\}^{k-i}) \times (\mathcal{P}^{i-1} \times \{b\} \times \mathcal{P}^{k-i}) \subseteq \preceq_k$.*

We will henceforth drop the subscript k in $\preceq_k$ whenever there is no resulting ambiguity.

Definition 3.3. *Let $\preceq$ be a k partial order on $\mathfrak{A}$ and $C \subseteq \{1, ..., k\}$. A $\preceq$-multifilter of $\mathfrak{A}$ with respect to C is a tuple $\mathfrak{F}_{\preceq}(C) = \langle F_1, F_2, ..., F_k \rangle$ of subsets of $\mathcal{P}$ such that*

1. $\top \in F_i$, *for* $1 \leq i \leq k$;
2. *if $i \in C$, $a \in F_i$, and $b \in F_i$, then $a.b \in F_i$; and*
3. *if $(a_1, ..., a_k) \preceq (b_1, ..., b_k)$ and $(a_1, ..., a_k) \in \times_{i=1}^{k} F_i$ then $(b_1, ..., b_k) \in \times_{i=1}^{k} F_i$.*

We can observe at this point that the three conditions on multifilters are just generalizations of the three conditions on filters. The second condition though need not apply to all the sets $F_1, ..., F_k$. The set C specifies the sets which behave *classically* in observing the second condition.

We next define how multifilters can be generated by a tuple of sets of propositions. The intuition is that each set of propositions represents a mental attitude and the tuple of sets represents the collective mental state.

Definition 3.4. *Let $\mathcal{Q}_1, ..., \mathcal{Q}_k \subseteq \mathcal{P}$, $\preceq$ be a k partial order on $\mathfrak{A}$, and $C \subseteq \{1, ..., k\}$. The $\preceq$-multifilter generated by $\langle \mathcal{Q}_1, ..., \mathcal{Q}_k \rangle$ with respect to C, denoted $\mathfrak{F}_{\preceq}(\langle \mathcal{Q}_1, ..., \mathcal{Q}_k \rangle, C)$, is a $\preceq$-multifilter $\langle \mathcal{Q}'_1, ..., \mathcal{Q}'_k \rangle$ with respect to C where $\mathcal{Q}'_i$ is the smallest set containing $\mathcal{Q}_i$, for $1 \leq i \leq k$, $\mathcal{Q}'_i$.*

The following theorem states that, under certain conditions, multifilters can be reduced to classical filters applied to the different sets of propositions representing the different mental attitudes.

Theorem 1. *Let $\mathcal{Q}_1, ..., \mathcal{Q}_k \subseteq \mathcal{P}$, $C \subseteq \{1, ..., k\}$, and $\preceq$ be a k partial order on $\mathfrak{A}$ which is classical in i for some $i \in C$. If $\mathfrak{F}_{\preceq}(\langle \mathcal{Q}_1, ..., \mathcal{Q}_k \rangle, C) = \langle \mathcal{Q}'_1, ..., \mathcal{Q}'_k \rangle$, then $\mathcal{Q}'_i = F(\mathcal{Q}_i)$.*

In the remainder of the paper, we will be assuming that practical reasoning is based on a tuple of sets of propositions, the first set representing the agent's *beliefs* and the rest representing different types of *motivations* that the agent acts upon. When using multifilters, we will assume that only the set of beliefs behaves classically ($C = \{1\}$). We will henceforth use $\mathfrak{F}_{\preceq}(\langle \mathcal{Q}_1, ..., \mathcal{Q}_k \rangle)$ as a shorthand for $\mathfrak{F}_{\preceq}(\langle \mathcal{Q}_1, ..., \mathcal{Q}_k \rangle, \{1\})$.

4 $Log_A\mathbf{PR}$ Languages

In this section, we present the syntax and semantics of $Log_A\mathbf{PR}$ in addition to defining two logical consequence relations one for beliefs and the other for motivations. Utilizing the multifilters presented in Section 3, we show that our logical consequence relations have the distinctive properties of classical Tarskian logical consequence.

4.1 $Log_A\mathbf{PR}$ Syntax

$Log_A\mathbf{PR}$ consists of terms constructed algebraically from function symbols. There are no sentences; instead, we use terms of a distinguished syntactic type to denote propositions. Propositions are included as first-class individuals in the $Log_A\mathbf{PR}$ ontology and are structured in a Boolean algebra. Though non-standard, the inclusion of propositions in the ontology has been suggested by several authors (Church 1950; Bealer 1979; Parsons 1993; Shapiro 1993). *Grades* are also taken to be first-class individuals. As a result, propositions *about* graded beliefs and motivations can be constructed, which are themselves recursively gradable.

A $Log_A\mathbf{PR}$ language is a many-sorted language composed of a set of terms partitioned into three base sorts: σ_P is a set of terms denoting propositions, σ_G is a set of terms denoting grades, and σ_I is a set of terms denoting anything else. A $Log_A\mathbf{PR}$ *alphabet* Ω includes a non-empty, countable set of constant and function symbols each having a syntactic sort from the set $\sigma = \{\sigma_P, \sigma_G, \sigma_I\} \cup \{\tau_1 \longrightarrow \tau_2 \mid \tau_1 \in \{\sigma_P, \sigma_G, \sigma_I\}$ and $\tau_2 \in \sigma\}$ of syntactic sorts. Intuitively, $\tau_1 \longrightarrow \tau_2$ is the syntactic sort of function symbols that take a single argument of sort σ_P, σ_G, or σ_I and produce a functional term of sort τ_2. Given the restriction of the first argument of function symbols to base sorts, $Log_A\mathbf{PR}$ is, in a sense, a first-order language. In addition, an alphabet Ω includes a countably infinite set of variables of the three base sorts; a set of syncategorematic symbols including the comma, various matching pairs of brackets and parentheses, and the symbol $\forall$; and a set of logical symbols defined as the union of the following sets: (i) $\{\neg\} \subseteq \sigma_P \longrightarrow \sigma_P$, (ii) $\{\wedge, \vee\} \subseteq \sigma_P \longrightarrow \sigma_P \longrightarrow \sigma_P$, (iii) $\{\lessdot, \doteq\} \subseteq \sigma_G \longrightarrow \sigma_G \longrightarrow \sigma_P$, and (iv) $\{\mathbf{G}\} \cup \{\mathbf{M}_i\}_{i=1}^k \subseteq \sigma_P \longrightarrow \sigma_G \longrightarrow \sigma_P$. $\mathbf{G}(p, g)$ denotes a belief that the grade of p is g and $\mathbf{M}_i(p, g)$ denotes that p is a motivation of type i with a grade of g. Terms involving $\Rightarrow$ (material implication), $\Leftrightarrow$ (logical equivalence), and $\exists$ are abbreviations defined in the standard way.

A $Log_A\mathbf{PR}$ language $\mathcal{L}$ is the smallest set of terms formed according to the following rules, where t and t_i ($i \in \mathbb{N}$) are terms in $\mathcal{L}$.

- All variables and constants in the alphabet Ω are in $\mathcal{L}$.
- $f(t_1, \ldots, t_m) \in \mathcal{L}$, where $f \in \Omega$ is of type $\tau_1 \longrightarrow \ldots \longrightarrow \tau_m \longrightarrow \tau$ ($m > 0$) and t_i is of type τ_i.
- $\neg t \in \mathcal{L}$, where $t \in \sigma_P$.
- $(t_1 \otimes t_2) \in \mathcal{L}$, where $\otimes \in \{\wedge, \vee\}$ and $t_1, t_2 \in \sigma_P$.
- $\forall x(t) \in \mathcal{L}$, where x is a variable in Ω and $t \in \sigma_P$.
- $t_1 \lessdot t_2 \in \mathcal{L}$, where $t_1, t_2 \in \sigma_G$.
- $t_1 \doteq t_2 \in \mathcal{L}$, where $t_1, t_2 \in \sigma_G$.
- $\mathbf{G}(t_1, t_2) \in \mathcal{L}$, where $t_1 \in \sigma_P$ and $t_2 \in \sigma_G$.
- $\mathbf{M_i}(t_1, t_2) \in \mathcal{L}$, where $t_1 \in \sigma_P$ and $t_2 \in \sigma_G$.

In what follows, we consider two distinguished subsets Φ_G and Φ_M of σ_P. Φ_G is the set of terms of the form $\mathbf{G}(\phi, g)$ and Φ_M is a set of terms of the form $\mathbf{M}_i(\psi, g)$ with ψ not containing any occurrence of $\mathbf{G}$.

Definition 4.1. *A* $Log_A\mathbf{PR}$ *theory* $\mathbb{T}$ *is a triple* $\langle \mathbb{B}, \mathbb{M}, \mathbb{R} \rangle$ *where:*

- $\mathbb{B} \subseteq \sigma_P$ *represents the agent's beliefs;*
- $\mathbb{M} = (\mathbb{M}_1, ..., \mathbb{M}_k)$ *is a k-tuple of subsets of* Φ_M *representing the agent's k motivation types; and*
- $\mathbb{R}$ *is a set of bridge rules each of the form* $B, M_1, ..., M_k \longmapsto B', M_1', ..., M_k'$ *where* $B \subseteq \sigma_P$, $B' \subseteq \Phi_G$, *and* $M_1, ..., M_k, M_1', ..., M_k' \subseteq \Phi_M$.

The bridge rules serve to "bridge" propositions across the different mental attitudes. A bridge rule $B, M_1, ..., M_k \longmapsto B', M_1', ..., M_k'$ means that if B is a subset of the current beliefs and M_i is a subset of the current i motivations, then B' should be added to the current beliefs and each M_i' should be added to the current i motivations.

We now go back to Example 1 showing a corresponding encoding of it as a $Log_A\mathbf{PR}$ theory.

Example 2. *Let "p" denote working on the presentation, "t" denote going to the trip, and "m" denote the boss's getting mad. A possible* $Log_A\mathbf{PR}$ *theory representing Example 1 is* $\mathbb{T} = \langle \mathbb{B}, (\mathbb{M}_1, \mathbb{M}_2), \mathbb{R} \rangle$ *where:*

- $\mathbb{B}$ *is made up of the following terms.*
  b1. $\mathbf{G}(p \wedge t, 5)$
  b2. $\mathbf{G}(\mathbf{G}(p \Leftrightarrow \neg t, 10), 2)$
- $\mathbb{M}_1 = \{\mathbf{M_1}(t, 1)\}$.
- $\mathbb{M}_2 = \{\mathbf{M_2}(p, 1), \mathbf{M_2}(\mathbf{M_1}(\neg m, 2), 3)\}$.
- $\mathbb{R}$ *is the set of instances of the following rule schema where* ϕ *and* g *are variables.*
  r1. $\{\}, \{\}, \{\mathbf{M_1}(\phi, g)\} \longmapsto \{\}, \{\mathbf{M_1}(\phi, g)\}, \{\}$.
  r2. $\{\}, \{\mathbf{M_1}(\neg m, g)\}, \{\} \longmapsto \{\}, \{\mathbf{M_1}(p, g), \mathbf{M_1}(\neg t, g)\}, \{\}$.

b1 *represents Ted's belief that he can work on the presentation. He trusts his belief* ***b1*** *with a degree of 5.* ***b2*** *represents the information Ted acquired through a chain of sources (Barney and Robin) that he cannot work on the presentation while being on the trip. Since Ted trusts Robin more than Barney,* $p \Leftrightarrow \neg t$ *which is acquired through Robin is given the grade 10 and the whole graded belief* $\mathbf{G}(p \Leftrightarrow \neg t, 10)$ *acquired through Barney is given the grade of 2 as Ted trusts Barney the least. There are two types of motivations in this example: Ted's personal desires and his obligations.* $\mathbf{M_1}(\phi, g)$ *represents that Ted desires to* ϕ *with a degree of* g*. Likewise,* $\mathbf{M_2}(\phi, g)$ *represents that Ted is obliged to* ϕ *with a degree of* g*.* $\mathbb{M}_1$ *is made up of Ted's desire to go to the trip with a degree of 1.* $\mathbb{M}_2$ *is made up of Ted's obligations. The first one represents Ted's obligation to work on the presentation with a degree of 1 as well. The second represents his obligation to desire to not make his boss mad.* ***r1*** *is a bridge rule motivated by Ted's character that prefers to pursue his obligations. So whenever Ted is obliged to have a desire, then he has it as a desire.* ***r2*** *represents that if Ted has a desire to make his boss not mad with a degree* g*, then he should desire to work on the presentation with the same degree* g *and not go to the trip with the same grade.*

4.2 From Syntax to Semantics

A key element in the semantics of $Log_A\mathbf{PR}$ is the notion of a $Log_A\mathbf{PR}$ *structure*.

Definition 4.2. *A* $Log_A\mathbf{PR}$ structure *is a quintuple* $\mathfrak{S}_k = \langle \mathcal{D}, \mathfrak{A}, \mathfrak{g}, \mathfrak{M}_k, \ll, \mathfrak{e} \rangle$*, where*

- $\mathcal{D}$*, the domain of discourse, is a set with two disjoint, non-empty, countable subsets: a set of propositions* $\mathcal{P}$*, and a set of grades* $\mathcal{G}$*.*
- $\mathfrak{A} = \langle \mathcal{P}, +, \cdot, -, \perp, \top \rangle$ *is a complete, non-degenerate Boolean algebra (Sankappanavar and Burris 1981).*

- $\mathfrak{g} : \mathcal{P} \times \mathcal{G} \longrightarrow \mathcal{P}$ *is a belief-grading function.*
- $\mathfrak{M}_k = \{\mathfrak{m}_i \mid 1 \leq i \leq k\}$ *is a set of* k *motivation-grading functions such that each* $\mathfrak{m}_i : \mathcal{P} \times \mathcal{G} \longrightarrow \mathcal{P}$.
- $\ll: \mathcal{G} \times \mathcal{G} \longrightarrow \mathcal{P}$ *is a ordering function imposing a total order.*
- $\mathfrak{e} : \mathcal{G} \times \mathcal{G} \longrightarrow \{\perp, \top\}$ *is an equality function, where for every* $g_1, g_2 \in \mathcal{G}$*:* $\mathfrak{e}(g_1, g_2) = \top$ *if* $g_1 = g_2$*, and* $\mathfrak{e}(g_1, g_2) = \perp$ *otherwise.*

A *valuation* $\mathcal{V}$ of a $Log_A\mathbf{PR}$ language is a triple $\langle \mathfrak{S}, \mathcal{V}_f, \mathcal{V}_x \rangle$, where $\mathfrak{S}$ is a $Log_A\mathbf{PR}$ structure, $\mathcal{V}_f$ is a function that assigns to each function symbol an appropriate function on $\mathcal{D}$, and $\mathcal{V}_x$ is a function mapping each variable to a corresponding element of the appropriate block of $\mathcal{D}$. An *interpretation* of $Log_A\mathbf{PR}$ terms is given by a function $[\![\cdot]\!]^{\mathcal{V}}$.

Definition 4.3. Let $\mathcal{L}$ be a $Log_A\mathbf{PR}$ language and let $\mathcal{V}$ be a valuation of $\mathcal{L}$. An *interpretation* of the terms of $\mathcal{L}$ is given by a function $[\![\cdot]\!]^{\mathcal{V}}$:

- $[\![x]\!]^{\mathcal{V}} = \mathcal{V}_x(x)$, for a variable x
- $[\![c]\!]^{\mathcal{V}} = \mathcal{V}_f(c)$, for a constant c
- $[\![f(t_1, \ldots, t_n)]\!]^{\mathcal{V}} = \mathcal{V}_f(f)([\![t_1]\!]^{\mathcal{V}}, \ldots, [\![t_m]\!]^{\mathcal{V}})$, for an m-adic ($m \geq 1$) function symbol f
- $[\![(t_1 \wedge t_2)]\!]^{\mathcal{V}} = [\![t_1]\!]^{\mathcal{V}} \cdot [\![t_2]\!]^{\mathcal{V}}$
- $[\![(t_1 \vee t_2)]\!]^{\mathcal{V}} = [\![t_1]\!]^{\mathcal{V}} + [\![t_2]\!]^{\mathcal{V}}$
- $[\![\neg t]\!]^{\mathcal{V}} = -[\![t]\!]^{\mathcal{V}}$
- $[\![\forall x(t)]\!]^{\mathcal{V}} = \prod_{a \in \mathcal{D}} [\![t]\!]^{\mathcal{V}[a/x]}$
- $[\![t_1 \lessdot t_2]\!]^{\mathcal{V}} = [\![t_1]\!]^{\mathcal{V}} \ll [\![t_2]\!]^{\mathcal{V}}$
- $[\![t_1 \doteq t_2]\!]^{\mathcal{V}} = \mathfrak{e}([\![t_1]\!]^{\mathcal{V}}, [\![t_2]\!]^{\mathcal{V}})$
- $[\![\mathbf{G}(t_1, t_2)]\!]^{\mathcal{V}} = \mathfrak{g}([\![t_1]\!]^{\mathcal{V}}, [\![t_2]\!]^{\mathcal{V}})$
- $[\![\mathbf{M_i}(t_1, t_2)]\!]^{\mathcal{V}} = \mathfrak{m}_i([\![t_1]\!]^{\mathcal{V}}, [\![t_2]\!]^{\mathcal{V}})$

In the rest of the paper, for any $\Gamma \subseteq \sigma_p$, we will use $[\![\Gamma]\!]^{\mathcal{V}}$ to denote $\prod_{p \in \Gamma} [\![p]\!]^{\mathcal{V}}$ for notational convenience.

4.3 Logical Consequence

In this section, we employ our notion of multifilters from Section 3 to define logical consequence for $Log_A\mathbf{PR}$ in algebraic terms. In Section 3, we defined multifilters based on an arbitrary partial order $\preceq$. We start by defining how to construct such an order for the tuples of propositions in $\mathcal{P}$. The intuition is that the order is induced by the bridge rules in a $Log_A\mathbf{PR}$ theory in addition to the natural order $\leq$ among the belief propositions.

Definition 4.4. *Let* $\mathbb{T} = \langle \mathbb{B}, (\mathbb{M}_1, \ldots, \mathbb{M}_k), \mathbb{R} \rangle$ *be a* $Log_A\mathbf{PR}$ *theory and* $\mathcal{V}$ *a valuation. A* $\mathbb{T}^{\mathcal{V}}$*-induced order, denoted* $\preceq_{\mathbb{T}^{\mathcal{V}}}$*, is a partial order over* $\mathcal{P}^{k+1}$ *with the following properties.*

1. *If* $b \leq b'$*, then* $(b, \top, \ldots, \top) \preceq_{\mathbb{T}^{\mathcal{V}}} (b', \top, \ldots, \top)$.
2. *If* $(B, \ldots, M_k \longmapsto B', \ldots, M'_k) \in \mathbb{R}$*, then* $([\![B]\!]^{\mathcal{V}}, \ldots, [\![M_k]\!]^{\mathcal{V}}) \preceq_{\mathbb{T}^{\mathcal{V}}} ([\![B']\!]^{\mathcal{V}}, \ldots, [\![M'_k]\!]^{\mathcal{V}})$

Observation 4.1. *If* $\mathbb{T} = \langle \mathbb{B}, \mathbb{M}, \mathbb{R} \rangle$ *is a* $Log_A\mathbf{PR}$ *theory and* $\mathcal{V}$ *a valuation, then* $\preceq_{\mathbb{T}^{\mathcal{V}}}$ *is a* $k+1$ *partial-order on* $\mathfrak{A}$*. Further, if, for every* $(B, M_1, \ldots, M_k \longmapsto B', M'_1, \ldots, M'_k) \in \mathbb{R}$*,* $B' \neq \{\}$ *only if* $M_i = M'_i = \{\}$ *and* $[\![B]\!]^{\mathcal{V}} \leq [\![B']\!]^{\mathcal{V}}$*, then* $\preceq_{\mathbb{T}^{\mathcal{V}}}$ *is classical in 1.*

We next utilise a multifilter based on a $\mathbb{T}^{\mathcal{V}}$-induced order to define an extended logical consequence relation for beliefs and motivations.

Definition 4.5. *Let* $\mathbb{T} = \langle \mathbb{B}, (\mathbb{M}_1, \ldots, \mathbb{M}_k), \mathbb{R} \rangle$ *a* $Log_A\mathbf{PR}$ *theory. For every* $\phi \in \sigma_P$*,* ϕ *is a belief (or motivation) consequence of* $\mathbb{T}$*, denoted* $\mathbb{T} \models_B \phi$ *(or* $\mathbb{T} \models_M \phi$*), if, for every valuation* $\mathcal{V}$*,* $[\![\phi]\!]^{\mathcal{V}} \in \mathcal{B}$ *(or* $[\![\phi]\!]^{\mathcal{V}} \in \mathcal{M}_i$ *for some* i*,* $1 \leq i \leq k$*), where* $\langle \mathcal{B}, \mathcal{M}_1, \ldots, \mathcal{M}_k \rangle = \mathfrak{F}_{\preceq_{\mathbb{T}^{\mathcal{V}}}}(\langle [\![\mathbb{B}]\!]^{\mathcal{V}}, [\![\mathbb{M}]\!]_1^{\mathcal{V}}, \ldots, [\![\mathbb{M}]\!]_k^{\mathcal{V}} \rangle)$.

Both $\models_B$ and $\models_M$ are *monotonic* and have the distinctive properties of classical Tarskian logical consequence, with $\models_B$ observing a variant of the deduction theorem.

Theorem 2. *Let* $\mathbb{T} = \langle \mathbb{B}, \mathbb{M}, \mathbb{R} \rangle$ *and* $\mathbb{T}' = \langle \mathbb{B}', \mathbb{M}', \mathbb{R}' \rangle$ *be* $Log_A\mathbf{PR}$ *theories.*

1. *If* $\phi \in \mathbb{B}$*, then* $\mathbb{T} \models_B \phi$.
2. *If* $\phi \in \mathbb{M}_i$ *for some* $\mathbb{M}_i \in \mathbb{M}$*, then* $\mathbb{T} \models_M \phi$.
3. *If* $\mathbb{T} \models_B \phi$*,* $\mathbb{B} \subseteq \mathbb{B}'$*,* $\mathbb{M}_i \subseteq \mathbb{M}'_i$ *for all* $1 \leq i \leq k$*, and* $\mathbb{R}' \subseteq \mathbb{R}$*, then* $\mathbb{T}' \models_B \phi$.
4. *If* $\mathbb{T} \models_M \phi$*,* $\mathbb{B} \subseteq \mathbb{B}'$*,* $\mathbb{M}_i \subseteq \mathbb{M}'_i$ *for all* $1 \leq i \leq k$*, and* $\mathbb{R}' \subseteq \mathbb{R}$*, then* $\mathbb{T}' \models_M \phi$.
5. *If* $\mathbb{T} \models_B \psi$ *and* $\langle \mathbb{B} \cup \{\psi\}, \mathbb{M}, \mathbb{R} \rangle \models_B \phi$*, then* $\mathbb{T} \models_B \phi$.
6. *Let* $\mathbb{M}'_i = \mathbb{M}_i \cup \{\psi\}$*, for some* $1 \leq i \leq k$*, and* $\mathbb{M}'_j = \mathbb{M}_j$*, for* $j \neq i$*. If* $\mathbb{T} \models_M \psi$ *and* $\langle \mathbb{B}, \mathbb{M}', \mathbb{R} \rangle \models_M \phi$*, then* $\mathbb{T} \models_M \phi$.
7. *If* $\langle \mathbb{B} \cup \{\phi\}, \mathbb{M}, \mathbb{R} \rangle \models_M \psi$*, then* $\mathbb{T} \models_B \phi \Rightarrow \psi$.

5 Graded Multifilters

Consider the $Log_A\mathbf{PR}$ theory of the weekend dilemma from Example 2. Given that Ted believes $\mathbf{G}(p \wedge t, 5)$, and does not believe $\neg(p \wedge t)$, it would make sense for him to accept $p \wedge t$ despite his uncertainty about it. (Who is ever absolutely certain of their beliefs?) Similarly, it would make sense for Ted to add t to his desires and p to his obligations if they do not conflict with other motivations or beliefs. However, if we only use multifilters, we will never be able to reason with those nested graded beliefs and motivations as they are not themselves in the agent's theory but only grading propositions thereof. For this reason, we extend our notion of multifilters into a more liberal notion of *graded multifilters* to enable the agent to conclude, in addition to the consequences of the initial theory, beliefs and motivations graded by the initial beliefs and motivations (like $p \wedge t$). Should this lead to contradictions, the agent's character and the grades of the contradictory propositions are used to resolve them. Due to nested grading, graded multifilters come in degrees depending on the depth of nesting of the admitted graded propositions.

The rest of this section is dedicated to formalizing graded filters and presenting our definition of graded consequence for beliefs and motivations. We start by introducing some convenient abbreviations and notational conventions.

- Since we are modelling joint reasoning with beliefs and different types of motivations, in the sequel we assume a tuple $\mathcal{Q} = \langle \mathcal{Q}_0, \mathcal{Q}_1, \ldots, \mathcal{Q}_k \rangle$ where $\mathcal{Q}_0, \ldots, \mathcal{Q}_k \subseteq \mathcal{P}$. $\mathcal{Q}_0$ represents a set of believed propositions, and $\mathcal{Q}_1, \ldots, \mathcal{Q}_k$ represent sets of motivation propositions where each $\mathcal{Q}_i$

represents a different type of motivation. We will refer to $\mathcal{Q}$ as the mental state of the agent.

- For every $p \in \mathcal{P}$ and $g \in \mathcal{G}$, $\mathfrak{g}(p, g)$ is referred to as a *belief-grading proposition* that *grades* p and p is a *graded belief*. Similarly, $\mathfrak{m}_\mathfrak{i}(p, g)$ is a *motivation-grading proposition* and p is a *graded motivation*.
- If $\mathfrak{g}(p, g) \in \mathcal{Q}_0$, then p is graded in $\mathcal{Q}_0$. Similarly, if $\mathfrak{m}_\mathfrak{i}(p, g) \in \mathcal{Q}_i$, then p is *graded* in $\mathcal{Q}_i$.
- If $\mathcal{R} \subseteq \mathcal{P}$ and $p \in \mathcal{P}$, then $G_B(p, \mathcal{R}) = \{\mathfrak{g}(p, g) \mid g \in \mathcal{G}$ and $\mathfrak{g}(p, g) \in \mathcal{R}\}$ and $G_{Mi}(p, \mathcal{R}) = \{\mathfrak{m}_\mathfrak{i}(p, g) \mid g \in \mathcal{G}$ and $\mathfrak{m}_\mathfrak{i}(p, g) \in \mathcal{R}\}$, for $1 \leq i \leq k$.

5.1 Embedding and Grading Chains

As a building step towards formalizing graded multifilters, the structure of graded propositions should be carefully scrutinized.

Definition 5.1. *Let $\mathcal{R} \subseteq \mathcal{P}$ and X be any of B or M_i, for $1 \leq i \leq k$. The set $E^n_X(\mathcal{R})$ of X-embedded propositions at depth $n \in \mathbb{N}$ in $\mathcal{R}$ is inductively defined as follows.*

- $E^0_X(\mathcal{R}) = \mathcal{R}$ *and*
- $E^{i+1}_X(\mathcal{R}) = E^i_X(\mathcal{R}) \cup \{p \mid G_X(p, E^i_X(\mathcal{R})) \neq \{\}\}$.

In the sequel, recalling that $\mathcal{Q} = \langle \mathcal{Q}_0, \mathcal{Q}_1, \ldots, \mathcal{Q}_k \rangle$ is the agent's mental state, we let

$$E^n(\mathcal{Q}) = \langle E^n_B(\mathcal{Q}_0), E^n_{M_1}(\mathcal{Q}_1), ..., E^n_{M_k}(\mathcal{Q}_k) \rangle$$

Having carefully defined the notions of embedding and the degree of embedding of a graded proposition, we say that a *grading chain* of a belief (or motivation) p is a non-empty, finite sequence $\langle q_0, q_1, \ldots, q_n \rangle$ where $q_0, q_1, \ldots, q_n$ are belief (or motivation) grading propositions such that q_i grades q_{i+1} for $1 \leq i < n$ and q_n grades p. We next define some properties of sets of propositions based on the grading chains they include.

Definition 5.2. *Let $\mathcal{R} \subseteq \mathcal{P}$.*

1. *$\mathcal{R}$ is depth-bounded if there is some $d \in \mathbb{N}$ such that every belief (or motivation) grading chain in $\mathcal{R}$ has at most d distinct grading propositions.*
2. *$\mathcal{R}$ is fan-out-bounded if there is some $f_{\text{out}} \in \mathbb{N}$ such that every grading belief (or motivation) chain in $\mathcal{R}$ grades at most f_{out} propositions.*
3. *$\mathcal{R}$ is fan-in bounded if there is some $f_{\text{in}} \in \mathbb{N}$ where $|G_B(p, \mathcal{R})| \leq f_{\text{in}}$ ($|G_{M_i}(p, \mathcal{R})| \leq f_{\text{in}}$), for every $p \in \mathcal{R}$.*

Since nested grading is allowed, it is necessary to define the *fused grade* of a graded proposition p in a chain C. Moreover, a proposition p might be graded by more than one grading chain. Accordingly, we also need to fuse the grades of p across all the chains grading it in some $\mathcal{R} \subseteq \mathcal{P}$. The intuition is to compute the fused grade of p for each chain that grades it by some operator $\otimes$, then combine these fused grades together using another operator $\oplus$.

Definition 5.3. *Let $\mathcal{R} \subseteq \mathcal{P}$ be fan-in-bounded , then the fused grade of p in $\mathcal{Q}$ is defined as*

$$\mathfrak{f}_\oplus(p, \mathcal{R}) = \bigoplus \langle \mathfrak{f}_\otimes(p, C_k) \rangle^n_{k=1}$$

where $\oplus : \bigcup^\infty_{i=1} \mathcal{G}^i \longrightarrow \mathcal{G}$ is commutative and $\langle C_k \rangle^n_{k=1}$ is a permutation of the set of longest grading chains of p in $\mathcal{R}$.

5.2 Telescoping and Graded Multifilters

The key to defining graded multifilters is the intuition that the set of consequences of $\mathcal{Q} = \langle \mathcal{Q}_0, \mathcal{Q}_1, ..., \mathcal{Q}_k \rangle$ may be further enriched by *telescoping* $\mathcal{Q}$ and accepting some of the beliefs and motivations embedded therein. We refer to this process as "telescoping" as the set of graded multifilters at increasing depths can be thought of as an inverted telescope. To this end, we need to define (i) the process of telescoping, which is a step-wise process that considers both beliefs and motivations at increasing degrees of embedding, and (ii) a criterion for accepting embedded beliefs and motivations without introducting inconsistencies. In this section, we will be formalizing the process of telescoping and the construction of graded multifilters. A first step towards defining graded multifilters is the notion of telescoping structures.

Definition 5.4. *Let $\mathfrak{S}_k$ be a $Log_A\mathbf{PR}$ structure with a depth- and fan-out-bounded $\mathcal{P}$. A telescoping structure for $\mathfrak{S}_k$ is a sextuple $\mathfrak{T} = \langle \mathcal{T}, \mathfrak{O}, \otimes_B, \oplus_B, \otimes_M, \oplus_M, \mathfrak{C} \rangle$, where*

- *$\mathcal{T} = \langle \mathcal{T}_0, \mathcal{T}_1, ..., \mathcal{T}_k \rangle$, where $\mathcal{T}_0, \mathcal{T}_1, ..., \mathcal{T}_k \subseteq \mathcal{P}$. $\mathcal{T}_0$ is referred to as the set of top beliefs, and each $\mathcal{T}_i$, for $1 \leq i \leq k$, is referred to as a set of top motivations.*
- *$\mathfrak{O}$ is an ultrafilter of the subalgebra induced by $Range(\ll)$ (an ultrafilter is a maximal filter with respect to not including $\bot$ (Sankappanavar and Burris 1981));*
- *$\otimes_B, \oplus_B, \otimes_M$, and $\oplus_M$ are fusion functions from tuples of grades to grades; $\oplus_B$ and $\oplus_M$ are commutative.*
- *$\mathfrak{C}$ is a partial preorder over the set $\{0,...,k\}$ representing the agent's character.*

The telescoping structure provides the top beliefs and motivations that will never be given up together with their consequences. The ultrafilter $\mathfrak{O}$ provides a total ordering over grades to enable comparing them. The operators $\otimes_B, \oplus_B$ and $\otimes_M, \oplus_M$ are used to get fused grades for beliefs and motivations respectively as per Definition 5.3. It is worth noting that, for simplicity, we opted for fusing the grades of all types of motivations using the same pair of operators $\otimes_M$ and $\oplus_M$. The agent's character $\mathfrak{C}$ is defined as an ordering over the set $\{0, ..., k\}$. 0 will be taken to represent the agent's beliefs and $1, .., k$ represent the different types of motivation. The character of the agent in addition to the grades of the motivations will be utilised when picking a consistent set of motivations making up the agent's intentions. For simplicity, in the sequel, we will be assuming that the agent's character is a total order, giving rise to what we will refer to as a *linear character.*

We are now ready to define the $\mathfrak{T}$-induced telescoping of $\mathcal{Q}$. The process of telescoping $\mathcal{Q}$ is made up of first getting the multifilter of $\mathcal{Q}$ then extracting the graded propositions embedded at depth 1. This might introduce inconsistencies. We resolve the inconsistencies by getting the tuple of *kernel survivors* $\kappa(E^1(\mathfrak{F}_\preceq(\mathcal{Q})), \mathfrak{T})$ given the telescoping structure $\mathfrak{T}$. The telescoping structure $\mathfrak{T}$ is useful in getting the survivors as it contains the top beliefs and motivations, fusion operators, and the agent character that will all be used to decide which propositions to keep and which to give up. Since this process can cause some propositions to be given up, other propositions may lose their *support*. For this reason, we only retain the tuple of *supported propositions* $\varsigma(\kappa(E^1(\mathfrak{F}_{\preceq}(\mathcal{Q})), \mathfrak{T}), \mathcal{T})$ amongst the kernel survivors.

Definition 5.5. *Let $\mathfrak{T}$ be a telescoping structure for $\mathfrak{S}_k$. If $\mathcal{Q} = \langle \mathcal{Q}_0, \mathcal{Q}_1, ..., \mathcal{Q}_k \rangle$ where every item in $E^1(\mathfrak{F}_{\preceq}(\mathcal{Q}))$ is fan-in-bounded, then the $\mathfrak{T}$-induced telescoping of $\mathcal{Q}$ is given by*

$$\tau_{\mathfrak{T}}(\mathcal{Q}) = \varsigma(\kappa(E^1(\mathfrak{F}_{\preceq}(\mathcal{Q})), \mathfrak{T}), \mathcal{T})$$

The first two steps of telescoping were already presented in Definitions 3.4 and 5.1 respectively, we only need to define the tuples of kernel survivors and supported propositions. For kernel survival, we generalize the notion of a $\perp$-kernel of a belief base (Hansson 1994) to suit reasoning with multiple sets of propositions. The intuition is that a $\perp$-kernel is a tuple of sets, one for each mental attitude, where the union of the sets is a subset-minimal inconsistent set. This means that if we remove all occurrences of a single proposition from the sets in the $\perp$-kernel, the union becomes consistent. In what follows, we say that a set $\mathcal{R} \subseteq \mathcal{P}$ is inconsistent whenever the classical filter of $\mathcal{R}$ is improper ($F(\mathcal{R}) = \mathcal{P}$).

Definition 5.6. *Let $\mathcal{Q} = \langle \mathcal{Q}_0, ..., \mathcal{Q}_k \rangle$, $\mathcal{X} = \langle \mathcal{X}_0, ..., \mathcal{X}_k \rangle$, and $\langle \mathcal{X}'_0, ..., \mathcal{X}'_k \rangle = \mathfrak{F}_{\preceq}(\mathcal{X})$. $\mathcal{X}$ is a $\perp$-kernel of $\mathcal{Q}$ iff $\mathcal{X}_i \subseteq \mathcal{Q}_i$, $1 \leq i \leq k$, and $\mathcal{X}'_0 \cup \mathcal{X}'_1 \cup ... \cup \mathcal{X}'_k$ is a subset-minimal inconsistent set of propositions.*

Example 3. *We refer back to Example 2. The following are examples of $\perp$-kernels.*[1] *The first set in each $\perp$-kernel represents beliefs of Ted's, the second represents desires thereof, and the third obligations.*

1. $\langle \{p \wedge t, p \Leftrightarrow \neg t\}, \{\}, \{\} \rangle$.
2. $\langle \{p \Leftrightarrow \neg t\}, \{t\}, \{p\} \rangle$.
3. $\langle \{p \Leftrightarrow \neg t\}, \{p, t\}, \{\} \rangle$.
4. $\langle \{\}, \{t, \neg t\}, \{\} \rangle$.

The first $\perp$-kernel shows a contradiction within Ted's beliefs; the second shows a contradiction between a belief, a desire, and an obligation; the third shows a contradiction between a belief and two desires; and the forth shows a contradiction between two desires. Note that the unions resulting from the three $\perp$-kernels are subset minimal.

How do we choose propositions to give up and resolve inconsistency? The intuition is this: The proposition to be given up must be from the least preferred set in the mental state according to the agent's character. If the least preferred set contains more than one proposition, then the proposition to be given up must be the proposition with the least grade in the set. To make finding the least preferred set easier, we reorder $\mathcal{Q}$ such that its items are ordered from the least preferred to the most preferred according to the character, and construct the $\perp$-kernels out of the reordered $\mathcal{Q}$. In this way, the least preferred set in a $\perp$-kernel will be the left-most non-empty set.

Henceforth, we assume $\mathcal{Q} = \langle \mathcal{Q}_0, ..., \mathcal{Q}_k \rangle$ where each set in $\mathcal{Q}$ is fan-in-bounded and a telescoping structure $\mathfrak{T} = \langle \mathcal{T}, \mathfrak{O}, \otimes_B, \oplus_B, \otimes_M, \oplus_M, \mathfrak{C} \rangle$ with $\mathcal{T} = \langle \mathcal{T}_0, \mathcal{T}_1, ..., \mathcal{T}_k \rangle$. We say that $\mathcal{Q}^{\mathfrak{C}}$ is a $\mathfrak{C}$-ordered $\mathcal{Q}$ if $\mathcal{Q}^{\mathfrak{C}}$ is a permutation of $\mathcal{Q}$ where $\forall i, j$ such that $i < j \in \mathfrak{C}$, $\mathcal{Q}_i$ appears before $\mathcal{Q}_j$ in $\mathcal{Q}^{\mathfrak{C}}$. In the sequel, let $\mathcal{Q}^{\mathfrak{C}} = \langle \mathcal{Q}'_0, ..., \mathcal{Q}'_k \rangle$ be a $\mathfrak{C}$-ordered $\mathcal{Q}$ and $\mathcal{Q}^{\mathfrak{C}}_{\perp}$ be the set of $\perp$-kernels in $\mathcal{Q}^{\mathfrak{C}}$ throughout.

Definition 5.7. *Let $\mathcal{X} = \langle \mathcal{X}_0, ..., \mathcal{X}_k \rangle$ be a $\perp$-kernel of $\mathcal{Q}^{\mathfrak{C}}$. p does not survive $\mathcal{X}$ given $\mathfrak{T}$ iff $\exists i$ $1 \leq i \leq k$ such that p is a graded proposition in $\mathcal{X}_i$ where $\mathcal{X}_i$ is the left-most non-empty set in $\mathcal{X}$ and $\forall q \in \mathcal{X}_i$ such that $q \notin \mathcal{T}'_i$ with $\mathfrak{F}_{\preceq}(\mathcal{T}) = \langle \mathcal{T}'_0, ..., \mathcal{T}'_k \rangle$, $(\mathfrak{f}_{\mathfrak{T}}(p, \mathcal{Q}'_i) \ll \mathfrak{f}_{\mathfrak{T}}(q, \mathcal{Q}'_i)) \in \mathfrak{O}$.*

We next define what we refer to as a *next-best* $\perp$-kernel in $\mathcal{Q}^{\mathfrak{C}}_{\perp}$. The next-best $\perp$-kernel is the $\perp$-kernel that must be first examined to pick a proposition to give up from one of its sets to resolve the inconsistency. Our motivation in defining the properties of a next-best $\perp$-kernel is to retain as much beliefs and motivations as possible. As a first condition, a next-best $\perp$-kernel has the longest sequence of empty sets from the left. In addition, a next-best $\perp$-kernel must contain a proposition with the maximum grade amongst the propositions with the minimum grades in the left-most non-empty set of all the $\perp$-kernels satisfying the first condition.

Definition 5.8. *Let $MinGrades = \{\mathfrak{f}_{\oplus}(p, \mathcal{Q}^{\mathfrak{C}}_{\perp}) \mid p$ is a proposition in the left-most non-empty set of a $\perp$-kernel $\mathcal{X}$ of $\mathcal{Q}^{\mathfrak{C}}_{\perp}$ with the longest sequence of empty sets from the left, and $\mathfrak{f}_{\oplus}(p, \mathcal{Q}^{\mathfrak{C}}_{\perp})$ is the minimum fused grade of a proposition in $\mathcal{X}\}$. A next-best $\perp$-kernel $\mathcal{X}^* = \langle \mathcal{X}_0, ..., \mathcal{X}_k \rangle \in \mathcal{Q}^{\mathfrak{C}}_{\perp}$ satisfies the following properties.*

1. *There does not exist another $\perp$-kernel in $\mathcal{Q}^{\mathfrak{C}}_{\perp}$ with a longer sequence of empty sets from the left.*
2. *$\mathcal{X}^*$ contains a proposition with the maximum grade in $MinGrades$.*

We are now ready to present the construction of the tuple of kernel survivors. What we do is, we pick a next-best $\perp$-kernel from $\mathcal{Q}^{\mathfrak{C}}_{\perp}$ and get the propositions that do not survive from it according to Definition 5.7. There might be more that one proposition that does not survive if they all have the same lowest grade. Such propositions are removed from all the beliefs and motivations in $\mathcal{Q}$ to resolve the inconsistency. The intuition behind doing this is that if some propositions in some set in the next-best $\perp$-kernel do not survive, then they can not survive in any other set to guarantee the consistency of the union of the sets in the mental state. We next proceed to getting the kernel survivors from the updated $\mathcal{Q}$ until the union of the sets in the mental state becomes consistent.

Definition 5.9. *The tuple of kernel survivors of $\mathcal{Q}$ given $\mathfrak{T}$ is $\kappa(\mathcal{Q}, \mathfrak{T})$ where $\kappa(\mathcal{Q}, \mathfrak{T})$ is defined as follows:*

1. *if $\mathcal{Q}^{\mathfrak{C}}_{\perp} = \varnothing$, then $\kappa(\mathcal{Q}, \mathfrak{T}) = \mathcal{Q}$; and*
2. *if $\mathcal{X}^*$ is a next-best $\perp$-kernel in $\mathcal{Q}^{\mathfrak{C}}_{\perp}$ and S is the set of propositions that do not survive $\mathcal{X}^*$ given $\mathfrak{T}$, then*

$$\kappa(\mathcal{Q}, \mathfrak{T}) = \kappa(\mathcal{Q}', \mathfrak{T})$$

where $\mathcal{Q}' = \langle \mathcal{Q}_0 - S, \mathcal{Q}_1 - S, ..., \mathcal{Q}_k - S \rangle$.

Example 4. *Suppose we have the same $\perp$-kernels in Example 3. The agent character according to Example 1 $\mathfrak{C} = \{0 < 1, 0 < 2, 2 < 1\}$ with 0 representing the agent's beliefs, 1 representing the desires, and 2 representing the obligations. After getting $\mathcal{Q}^{\mathfrak{C}}$, $\mathcal{Q}^{\mathfrak{C}}_{\perp}$ contains the following*

[1] We use the syntactic $\wedge$ and $\Leftrightarrow$ operators rather than their semantic counterparts for readability.

three kernels. The sets in the kernels are now ordered according to the agent character with the first set containing desires, the second set containing obligations, and the third set containing beliefs.

1. $\langle \{\}, \{\}, \{p \wedge t, p \Leftrightarrow \neg t\} \rangle$
2. $\langle \{t\}, \{p\}, \{p \Leftrightarrow \neg t\} \rangle$
3. $\langle \{p, t\}, \{\}, \{p \Leftrightarrow \neg t\} \rangle$
4. $\langle \{t, \neg t\}, \{\}, \{\} \rangle$

The next-best $\perp$-kernel is the first kernel as it has the longest sequence of empty sets from the left. In the first kernel we are forced to give up beliefs to resolve the contradiction as the desires and obligations sets are empty. The beliefs that will be given up will then be removed from the less preferred obligations and desires. In this way, we treat $\perp$-kernels where we have to give up a proposition from a more preferred set first so the removal of propositions from them affects less preferred set. To decide which propositions to give up, we look at the grades of $p \wedge t$ and $p \Leftrightarrow \neg t$. We consider three cases.

1. *If the grade of $p \Leftrightarrow \neg t$ is less, it will be removed from the first kernel and Ted's beliefs and motivations resolving the inconsistency in the second and third kernels as well. From the updated $\mathcal{Q}^{\mathfrak{C}}$, the forth kernel will be reconstructed and t will be given up from Ted's desires to resolve the contradiction as Ted's character prefers to retain his obligations.*
2. *If the grade of $p \wedge t$ is less, it will be removed from the first kernel and Ted's beliefs and motivations. However, the inconsistency in the second, third, and forth kernels is not resolved by giving up $p \wedge t$. From the updated $\mathcal{Q}^{\mathfrak{C}}$, the last three kernels are reconstructed. The three kernels have the same number of empty sets from the left, so to identify a next-best kernel we look at the kernel where the desires set contains a proposition with the maximum grade amongst the propositions with the minimum grades in the desires set. Suppose that p and $\neg t$ have a grade of 2 and t has a grade of 1. The minimum grade in the three kernels then will be 1. Therefore, all the three kernels are next-best kernels. Suppose we pick the second kernel, t is removed from the desires set in the kernel and Ted's beliefs and the motivations resulting in a consistent mental state. Giving up t from the second kernel resolves the inconsistency in the third and forth kernels.*
   On the other hand, if t has the grade of 2 and p and $\neg t$ have the grade of 1, the second kernel will be the next-best kernel as it contains the proposition with the maximum grade of the minimas in the desires set t. From the second kernel t will be removed from the desires set resolving the inconsistency in the other two kernels as well. Notice here that if we had defined the next-best kernels as the kernel that contains the minimum of the minimas rather than the maximum, the third or forth kernels would have been the next-best kernels. Suppose we choose the third kernel, p would have been chosen to be removed as it has the least grade resolving the contradiction in the second kernel too. However, the contradiction in the forth kernel would not have been resolved and $\neg t$ would have been removed as well from Ted's desires to resolve the contradiction as it has a lower grade than t. We end up here removing p and $\neg t$ from Ted's desires to resolve the contradiction rather than removing only t. This is why we define the next-best kernel as we did in Definition 5.8 to make sure that we retain as many motivations as possible.
3. *If the grades of $p \wedge t$ and $p \Leftrightarrow \neg t$ are equal, then both beliefs are given up. Both propositions are accordingly removed from the motivations resulting in a consistent union of the sets in the mental state, and resolving the inconsistency in the second and third kernels.*

After defining the kernel survivors, what remains for us to fully define the process of telescoping is to present the notion of the tuple of supported propositions in $\mathcal{Q}$ given a tuple of top sets of propositions $\mathcal{T} = \langle \mathcal{T}_0, \mathcal{T}_1, ..., \mathcal{T}_k \rangle$. The motivation for defining this is the following. Suppose in Example 4 we give up $p \wedge t$ from the first $\perp$-kernel. This means that any proposition that was supported by $p \wedge t$ must go away as well as it loses its support. Therefore, the following definition states that the supported propositions are the propositions in the sets of the multifilter of $\mathcal{T}$, or the propositions that are graded by supported propositions in $\mathcal{Q}$.

Definition 5.10. *The set of supported propositions in $\mathcal{Q} = \langle \mathcal{Q}_0, \mathcal{Q}_1, ..., \mathcal{Q}_k \rangle$ given $\mathcal{T} = \langle \mathcal{T}_0, \mathcal{T}_1, ..., \mathcal{T}_k \rangle$ denoted $\varsigma(\mathcal{Q}, \mathcal{T})$ is the tuple $\langle S_0, S_1, ..., S_k \rangle$ where $S_0, S_1, ..., S_k$ are the smallest subsets of, respectively, $\mathcal{Q}_0, \mathcal{Q}_1, ..., \mathcal{Q}_k$ such that, for $0 \leq i \leq k$,*

1. *$\forall p \in S_i$ if $\mathfrak{F}_{\preceq}(\mathcal{T}) = \langle \mathcal{T}'_0, \mathcal{T}'_1, ..., \mathcal{T}'_k \rangle$ and $p \in \mathcal{T}'_i$; and*
2. *$\forall p \in S_i$ if there is a grading chain $\langle q_0, ..., q_n \rangle$ of p in S_i and there is a tuple $(R_0, \ldots, R_k)$ where $R_j \subseteq S_j$, for $0 \leq j \leq k$, such that $q_0 \in R'_i$ where $\mathfrak{F}_{\preceq}(R) = \langle R'_0, \ldots, R'_k \rangle$.*

We can now present an important result. The following theorem states that if the union of the sets in the multifilter of $\mathcal{T}$ is consistent, then the union of the sets in the multifilter after getting the tuple of supported propositions in the kernel survivors of any $\mathcal{Q} \subseteq \mathcal{P}$ given $\mathcal{T}$ is consistent. This basically means that the process of telescoping is consistency preserving. Accordingly, we can do the revision of the agent's beliefs and motivations while maintaining consistency amongst all the beliefs and motivations.

Theorem 3. *Let $\mathfrak{F}_{\preceq}(\mathcal{T}) = \langle \mathcal{B}, \mathcal{M}_1, ..., \mathcal{M}_k \rangle$ and $\mathfrak{F}_{\preceq}(\varsigma(\kappa(\mathcal{Q}, \mathfrak{T}), \mathcal{T})) = \langle \mathcal{B}', \mathcal{M}'_1, ..., \mathcal{M}'_k \rangle$. If $F(\mathcal{B} \cup \mathcal{M}_1 \cup ... \cup \mathcal{M}_k)$ is proper, then $F(\mathcal{B}' \cup \mathcal{M}'_1 \cup ... \cup \mathcal{M}'_k)$ is proper.*

Since graded multifilters come in degrees depending on the nesting level of the telescoped propositions, we need to extend the $\mathfrak{T}$-induced telescoping of $\mathcal{Q}$ to a generalized notion of $\mathfrak{T}$-induced telescoping of the $\mathcal{Q}$ at degree n as follows.

Definition 5.11. *If each set in $\mathcal{Q}$ has finitely-many grading propositions, then $\tau_{\mathfrak{T}}(\mathcal{Q})$ is defined, for every telescoping structure $\mathfrak{T}$. In what follows, provided that the right-hand side is defined, let*

$$\tau^n_{\mathfrak{T}}(\mathcal{Q}) = \begin{cases} \mathcal{Q} & \text{if } n = 0 \\ \tau_{\mathfrak{T}}(\tau^{n-1}_{\mathfrak{T}}(\mathcal{Q})) & \text{otherwise} \end{cases}$$

We now are finally ready define graded multifilters as the multifilter of the $\mathfrak{T}$-induced telescoping of the tuple of sets of top propositions $\mathcal{T}$ at degree n.

Definition 5.12. *Let $\mathfrak{T}$ be a telescoping structure. We refer to $\mathfrak{F}_{\preceq}(\tau^n_{\mathfrak{T}}(\mathcal{T}))$ as a degree n ($\in \mathbb{N}$) graded filter of $\mathcal{T} = \langle \mathcal{T}_0, \mathcal{T}_1, ..., \mathcal{T}_k \rangle$, denoted $\mathcal{F}^n_{\preceq}(\mathcal{T})$.*

It is worth noting here that there might be several graded multifilters of degree n. This is due to the possible existence of several next-best $\perp$-kernels at each step of getting the kernel survivors. The order of considering the possible next-best $\perp$-kernels will affect the graded multifilter we end up with. Nevertheless, according to Theorem 3 the union of the sets in all the possible graded multifilters is consistent.

5.3 Graded Consequence

In what follows, given a $Log_A\mathbf{PR}$ theory $\mathbb{T} = \langle \mathbb{B}, (\mathbb{M}_1, ..., \mathbb{M}_k), \mathbb{R} \rangle$ and a valuation $\mathcal{V} = \langle \mathfrak{S}_k, \mathcal{V}_f, \mathcal{V}_x \rangle$, let the valuation of $\mathbb{T}$ be denoted as $\mathcal{V}(\mathbb{T}) = \langle \mathcal{V}(\mathbb{B}), \mathcal{V}(\mathbb{M}_1), ..., \mathcal{V}(\mathbb{M}_k) \rangle$. Just like we used multifilters to define logical consequence in Section 4, we use graded multifilters to define graded consequence as follows.

Definition 5.13. *Let $\mathbb{T} = \langle \mathbb{B}, (\mathbb{M}_1, ..., \mathbb{M}_k), \mathbb{R} \rangle$ be a $Log_A\mathbf{PR}$ theory and $\preceq_{\mathbb{T}}$ be a $\mathbb{T}$-induced order. For every $\phi \in \sigma_P$, valuation $\mathcal{V} = \langle \mathfrak{S}_k, \mathcal{V}_f, \mathcal{V}_x \rangle$ where $\mathfrak{S}_k$ has a set $\mathcal{P}$ which is depth- and fan-out-bounded, and grading canon $\mathcal{C} = \langle \otimes_B, \oplus_B, \otimes_M, \oplus_M, \mathfrak{C}, n \rangle$, ϕ is a graded belief (or motivation) consequence of $\mathbb{T}$ with respect to $\mathcal{C}$, denoted $\mathbb{T} \approx^{\mathcal{C}}_B \phi$ (or $\mathbb{T} \approx^{\mathcal{C}}_M \phi$), if $\mathcal{F}^n_{\preceq_{\mathbb{T}}}(\mathfrak{T}) = \langle \mathcal{B}, \mathcal{M}_1, ..., \mathcal{M}_k \rangle$ is defined and $[\![\phi]\!]^{\mathcal{V}} \in \mathcal{B}$ (or $\mathcal{M}_i$) for every telescoping structure $\mathfrak{T} = \langle \mathcal{V}(\mathbb{T}), \mathfrak{O}, \otimes_B, \oplus_B, \otimes_M, \oplus_M, \mathfrak{C} \rangle$ where $\mathfrak{O}$ extends $F(\mathcal{V}(\mathbb{T}) \cap Range(\ll))$ (an ultrafilter $\mathfrak{O}$ extends a filter F, if $F \subseteq \mathfrak{O}$).*

Note that $\approx^{\mathcal{C}}_B$ and $\approx^{\mathcal{C}}_M$ are non-monotonic and reduce to $\models_B$ and $\models_M$ respectively if $n = 0$. The set of belief consequences makes up a consistent set of beliefs, and the set of motivation consequences makes up a consistent set of motivations representing the agent's intentions.

5.4 The Weekend Dilemma in $Log_A\mathbf{PR}$

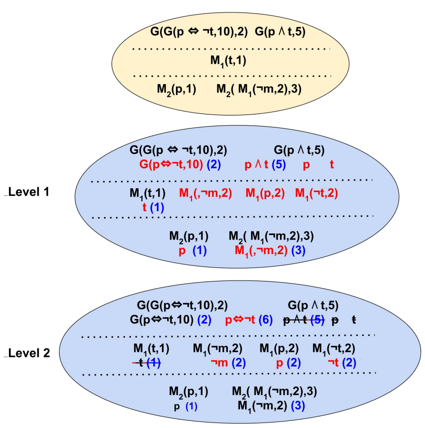

Figure 1: The graded consequences of the $Log_A\mathbf{PR}$ theory in Example 2. The top, middle and bottom portions of each level contains Ted's beliefs, desires, and obligations respectively. The newly added terms in each level is shown in red.

In this section we revisit the weekend dilemma showing how it can be accounted for in $Log_A\mathbf{PR}$ illustrating the joint belief and intention revision. Recall the $Log_A\mathbf{PR}$ theory $\mathbb{T} = \langle \mathbb{B}, (\mathbb{M}_1, \mathbb{M}_2), \mathbb{R} \rangle$ representing the weekend dilemma presented in Example 2. Figure 1 shows the graded belief and motivation consequences of $\mathbb{T}$ with respect to a series of canons with $\otimes_B$=mean, $\oplus_B$=max, and $\oplus_M, \otimes_M$=max, and $0 \leq n \leq 2$ with the agent character $\mathfrak{C} = \{0 < 1, 0 < 2, 2 < 1\}$.

Level 1: Upon telescoping to level 1, the embedded beliefs, desires, and obligations at level 1 are extracted. The classical consequences of the beliefs are added as well including p and t. Once $\mathbf{M}_1(\neg m, 3)$ is extracted in the obligations, the bridge rule **r1** fires to bridge $\mathbf{M}_1(\neg m, 3)$ to Ted's desires. This fires **r2** to add $\mathbf{M}_1(p, 3)$ and $\mathbf{M}_1(\neg t, 3)$ to Ted's desires as well. There are no contradictions between the extracted beliefs, desires and obligations so all the extracted beliefs and motivations survive telescoping and are supported. Hence, at level 1, Ted believes he can work on the presentation and go to the trip, desires to go to the trip, and is obliged to work on the presentation.

Level 2: At level 2, the embedded graded propositions at level 1 in the previous level are extracted adding $p \Leftrightarrow \neg t$ to Ted's beliefs, $\neg m, p$ and $\neg t$ to Ted's desires. Note that $\neg m$ is not extracted in the obligations as it we only telescope obligations in the set of obligations according to Definition 5.1 and $\neg m$ was in a desire term. No new bridge rules are fired in Level 2. However, once we do this, we get several contradictions between Ted's beliefs, desires, and obligations. We get the three $\perp$-kernels in Example 4. Since $p \wedge t$ has a lower grade (5) than $p \Leftrightarrow \neg t$ (with fused grade $\otimes(\langle 10, 2 \rangle) = 6$ as it is graded in a grading chain), the first case in the second scenario explained in Example 4 ensues resulting in removing $p \wedge t$ from the agent's beliefs and t from the desires. This causes both p and t in the agent's beliefs to go away as they lose their support. Hence, at level 2, Ted gives up his belief that he can work on the presentation while being on the trip. He accordingly gives up his desire to go to the trip and ends up desiring to not make his boss mad and consequently desiring working on the presentation and not going to the trip. Ted's obligations to desire to make his boss not mad and to work on the presentation are retained at level 2 as Ted's character prefers to give up desires rather than obligations. Note that a proposition was removed from the beliefs even though it is the highest preferred attitude, but this was necessary in order to resolve the contradiction within the beliefs. This revision only happened at level 2 as we look deeper into the nested graded propositions that contradicted Ted's beliefs and motivations at level 1.

6 Conclusion

Despite the abundance of logical theories in the literature for modelling practical reasoning, a robust theory with adequate semantics remains missing. In this paper, we introduced general algebraic foundations for practical reasoning with several mental attitudes. We also provided semantics for an algebraic logic , $Log_A\mathbf{PR}$, for joint reasoning with graded beliefs and motivations to decide on sets of consistent beliefs and intentions. The $Log_A\mathbf{PR}$ semantics also captures the joint revision of the agent's beliefs and intentions all in one framework. We are currently working on a proof theory for $Log_A\mathbf{PR}$. Reasons for intentions are to computed in the same way reason-maintenance systems computes supports for beliefs. The end result would be a proof theory for practical reasoning augmented with the ability to explain the reasons for choosing to adopt particular intentions to achieve an initial set of motivations giving rise to an explainable AI system.

References

Bealer, G. 1979. Theories of properties, relations, and propositions. *The Journal of Philosophy* 76(11):634–648.

Bratman, M. 1987. *Intention, plans, and practical reason*. Harvard University Press.

Broersen, J.; Dastani, M.; Hulstijn, J.; Huang, Z.; and van der Torre, L. 2001. The BOID architecture: conflicts between beliefs, obligations, intentions and desires. In *Proceedings of the fifth international conference on Autonomous agents*, 9–16.

Broersen, J.; Dastani, M.; Hulstijn, J.; and van der Torre, L. 2002. Goal generation in the BOID architecture. *Cognitive Science Quarterly* 2(3-4):428–447.

Broome, J. 2002. Practical reasoning. In Bermúdez, J. L., and Millar, A., eds., *Reason and nature: Essays in the theory of rationality*. Oxford: Clarendon Press. 85–111.

Casali, A.; Godo, L.; and Sierra, C. 2008. A logical framework to represent and reason about graded preferences and intentions. In Brewka, G., and Lang, J., eds., *Principles of Knowledge Representation and Reasoning: Proceedings of the Eleventh International Conference, KR 2008, Sydney, Australia, September 16-19, 2008*, 27–37. AAAI Press.

Casali, A.; Godo, L.; and Sierra, C. 2011. A graded BDI agent model to represent and reason about preferences. *Artificial Intelligence* 175(7-8):1468–1478.

Castelfranchi, C., and Paglieri, F. 2007. The role of beliefs in goal dynamics: prolegomena to a constructive theory of intentions. *Synthese* 155(2):237–263.

Charniak, E. 1991. Bayesian networks without tears. *AI magazine* 12(4):50–50.

Church, A. 1950. On carnap's analysis of statements of assertion and belief. *Analysis* 10(5):97–99.

Cohen, P. R., and Levesque, H. J. 1990. Intention is choice with commitment. *Artificial intelligence* 42(2-3):213–261.

Dunin-Keplicz, B.; Nguyen, L. A.; and Szalas, A. 2010. A framework for graded beliefs, goals and intentions. *Fundam. Inform.* 100(1-4):53–76.

Ehab, N., and Ismail, H. O. 2018. Towards a unified algebraic framework for non-monotonicity. *Proceedings of the KI 2018 Workshop on Formal and Cognitive Reasoning* 26–40.

Ehab, N., and Ismail, H. O. 2019. A unified algebraic framework for non-monotonicity. In Moss, L. S., ed., *Proceedings Seventeenth Conference on Theoretical Aspects of Rationality and Knowledge, TARK 2019, Toulouse, France, 17-19 July 2019*, volume 297 of *EPTCS*, 155–174.

Ehab, N., and Ismail, H. O. 2020. $Log_A\mathbf{G}$: An algebraic non-monotonic logic for reasoning with graded propositions. In *Annals of Mathematics and Artificial Intelligence*. Springer Nature.

Fern, A. 2010. Weighted logic. Technical report.

Frankfurt, H. G. 1988. Freedom of the will and the concept of a person. In *What is a person?* Springer. 127–144.

Hansson, S. O. 1994. Kernel contraction. *The Journal of Symbolic Logic* 59(03):845–859.

Icard, T.; Pacuit, E.; and Shoham, Y. 2010. Joint revision of belief and intention. In *Proc. of the 12th International Conference on Knowledge Representation*, 572–574.

Ismail, H. O. 2012. $Log_A\mathbf{B}$: A first-order, non-paradoxical, algebraic logic of belief. *Logic Journal of the IGPL* 20(5):774–795.

Ismail, H. O. 2013. Stability in a commonsense ontology of states. *Proceedings of the Eleventh International Symposium on Logical Formalization of Commonsense sense Reasoning (COMMONSENSE 2013)*.

Ismail, H. O. 2020. The good, the bad, and the rational: Aspects of character in logical agents. In ElBolock, A.; Abdelrahman, Y.; and Abdennadher, S., eds., *Character Computing*. Springer. 139–164.

Paccanaro, A., and Hinton, G. E. 2001. Learning distributed representations of concepts using linear relational embedding. *IEEE Transactions on Knowledge and Data Engineering* 13(2):232–244.

Parsons, T. 1993. On denoting propositions and facts. *Philosophical Perspectives* 7:441–460.

Rao, A. S., and Georgeff, M. P. 1995. BDI agents: From theory to practice. In *ICMAS*, volume 95, 312–319.

Rao, A. S., and Wooldridge, M. 1999. Foundations of rational agency. In Rao, A. S., and Wooldridge, M., eds., *Foundations of rational agency*. Springer. 1–10.

Richardson, M., and Domingos, P. 2006. Markov logic networks. *Machine learning* 62(1-2):107–136.

Sankappanavar, H., and Burris, S. 1981. A course in universal algebra. *Graduate Texts Math* 78.

Searle, J. R. 2003. *Rationality in action*. MIT press.

Shapiro, S. C. 1993. Belief spaces as sets of propositions. *Journal of Experimental & Theoretical Artificial Intelligence* 5(2-3):225–235.

Shoham, Y. 2009. Logical theories of intention and the database perspective. *Journal of Philosophical Logic* 38(6):633.

Thomason, R. H. 2018. The formalization of pratical reasoning: Problems and prospects. In Gabbay, D. M., and Guenthner, F., eds., *Handbook of Philosophical Logic: Volume 18*. Cham: Springer International Publishing. 105–132.

Vovk, V.; Gammerman, A.; and Shafer, G. 2005. *Algorithmic learning in a random world*. Springer Science & Business Media.