Preference Dynamics and Reasons for Preference

Fenrong Liu

Tsinghua University & University of Amsterdam

8-10 July 2015, IRIT, Université Paul Sabatier, Toulouse

Outline

1. Introduction
2. Betterness models: statics and dynamics
3. Reasons, priorities and dynamics
4. Relating betterness and priority dynamics
   - Entanglement of preference and belief
   - From individual to group behavior
5. Conclusions

Preference according to Von Wright

- *deontological* or *normative*: notions of right and duty, command, permission and prohibition,
- *axiological*: notions of good and evil, the comparative notion of betterness,
- *anthropological*: notions of need and want, decision and choice, motive, end, and action.

The concept of preference is related to the axiological notion of betterness on one side, but just as well to the anthropological notion of choice.

Issue 1: Intrinsic vs. extrinsic

"... a person says, for example, that he prefers claret to hock, because his doctor has told him or he has found from experience that the first wine is better for his stomach or health in general. In this case a ***judgement of betterness serves as a ground or reason*** *for a preference. I shall call preferences, which hold this relationship to betterness, extrinsic.*

It could, however, also be the case that a person prefers claret to hock, not because he thinks (opines) that the first wine is better for him, but simply because he likes the first better (more). Then his liking the one wine better is not a reason for his preference. ..."

Issue 2: Changes in preference

Consider the following twist to von Wright's scenario:

> *Suppose that before seeing his doctor, he preferred hock to claret. Now the doctor tells him "The second wine is better for your health". He then changes his preference, and will now prefer claret to hock.*

Merely giving information can change preferences!

Other kinds of preference change will be considered below.

“The preferences which we shall study are a subject’s intrinsic preferences on one occasion only. Thus we exclude both ***reasons*** *for preferences and the possibility of* ***changes*** *in preferences."*

([von Wright, 1963], p.23)

Preference logic

- Old tradition: Hallden 1957; Von Wright 1963, 1972.
- Classical modal logics for agents' having a preference: Hansson 2001, *Handbook of Philosophical Logic*.
- Decision theory and game theory: focus mostly on intrinsic preference.
- Computer science and AI: the well-known "*BDI* model" [Rao 1991]; studies on *ceteris paribus* preference"all else being *equal*" [Doyle and Wellman 1994]; social choice on preference aggregation.

Preference logic

- Preference change: [van Benthem et al. 1993] is a first attempt at using dynamic logic for this purpose. Influenced by *AGM*-style belief revision theory, [Hansson1995] proposed postulates governing four basic operations in preference change.
- Reasons for preference an upcoming topic: [Liu 2008], [Liu 2011], [Osherson and Weinstein 2012] in logic, [Dietrich and List 2013] in rational choice theory.

Plan

- Introduce a two-level model for reason-based preference.
- Present a logical account of preference change.
- Relate the dynamics at the two levels of preference.
- Further issues.
- Case study: priority structures in deontic logic.

Logical models: basic ideas

- The intrinsic view is reflected in possible worlds models with a ‘betterness’ ordering.
- ‘Priority graphs’ captures reasons for preference.
- Dynamic epistemic logic (DEL, as well as GDDL) as our main methodology for preference change.

Public Announcement Logic

Let P be a set of proposition letters and I a set of agents, with p ranging over P, a over I. The **dynamic epistemic language** is given by:

$$\varphi ::= \bot \mid p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_a\varphi \mid [\pi]\varphi$$
$$\pi ::= !\varphi.$$

- We read $[!\varphi]K_a\psi$ as: after the truthful public announcement φ, agent a knows ψ.

- A public announcement $!\varphi$ of a true proposition φ turns the current model $(\mathcal{M}, s)$ with actual world s into the model $(\mathcal{M}_{!\varphi}, s)$ whose worlds are just the set $\{w : \mathcal{M}, w \models \varphi\}$.
- The semantic clause for the dynamic modality is the following

  $\mathcal{M}, s \models [!\varphi]\psi$ iff (*if* $\mathcal{M}, s \models \varphi$, then $\mathcal{M}_{!\varphi}, s \models \psi$)

System of PAL

- All instantiations of propositional tautologies
- $K_a(\varphi \to \psi) \to (K_a\varphi \to K_a\psi)$
- $K_a\varphi \to \varphi$
- $K_a\varphi \to K_aK_a\varphi$
- $\neg K_a\varphi \to K_a\neg K_a\varphi$

- $[!\varphi]p \leftrightarrow \varphi \to p$
- $[!\varphi]\neg\psi \leftrightarrow \varphi \to \neg[!\varphi]\psi$
- $[!\varphi](\psi \wedge \chi) \leftrightarrow \varphi \to [!\varphi]\psi \wedge [!\varphi]\chi$
- $[!\varphi]K_a\psi \leftrightarrow (\varphi \to K_a[!\varphi]\psi)$
- From $\vdash \varphi$ and $\vdash \varphi \to \psi$ infer $\vdash \psi$
- From $\vdash \varphi$ infer $\vdash K_a\varphi$

Product Update

An **event model** is a tuple $\mathcal{E} = (E, \sim_a, PRE)$ such that

- E is a non-empty set of events,
- $\sim_a$ is a binary epistemic relation on E,
- PRE is a function from E to the collection of all epistemic formulas.

Let an epistemic model $\mathcal{M} = (S, \sim_a, V)$ and an event model $\mathcal{E} = (E, \sim_a, PRE)$ be given, the **product update model** is defined to be the model $\mathcal{M} \otimes \mathcal{E} = (S \otimes E, \sim'_a, V')$ such that

- $S \otimes E = \{(s, e) \in S \times E : s \models PRE(e)\}$
- (s, e) $\sim'_a$ (t, f) iff both $s \sim_a t$ and $e \sim_a f$
- $V'(p) = \{(s, e) \in S \otimes E : s \in V(p)\}$.

Some references on DEL

- Baltag, Moss and Solecki, *The Logic of Common Knowledge, Public Announcements, and Private Suspicious*, in Proceedings of TARK 1998.
- Van Ditmarsch, van der Hoek and Kooi, *Dynamic Epistemic Logic*, Springer, 2007.
- Van Benthem, Logical *Dynamics of Information and Interaction*, Oxford University Press, 2011.
- Girard, Seligman and Liu, General Dynamic Dynamic Logic, in *Advances in Modal Logic*, Volume 9, pp.239–260. College Publications, London, 2012.

Modal betterness logic

Definition (modal betterness model)

A *modal betterness model* $\mathcal{M} = (W, \preceq, V)$ has a set of worlds W, $\preceq$ is a reflexive and transitive relation (the 'betterness' pre-order), and V is a valuation for proposition letters.

If $s \preceq t$ but not $t \preceq s$, then t is *strictly better* than s ($s \prec t$).

Modal betterness logic

Definition (modal betterness language)

Let $p \in \mathtt{Prop}$, the *modal betterness language* is given by the following inductive syntax rule:

$$\varphi := p \mid \neg\varphi \mid \varphi \wedge \psi \mid \langle\leq\rangle\varphi \mid \langle<\rangle\varphi \mid E\varphi$$

$$[\leq]\varphi := \neg\langle\leq\rangle\neg\varphi$$
$$[<]\varphi := \neg\langle<\rangle\neg\varphi$$
$$A\varphi := \neg E\neg\varphi$$

Truth definition

Definition

Truth conditions for the atomic propositions and Boolean combinations are standard. Modalities work like this:

- $\mathcal{M}, s \models \langle\leq\rangle\varphi$ iff for some t wih $s \preceq t$, $\mathcal{M}, t \models \varphi$.
- $\mathcal{M}, s \models \langle<\rangle\varphi$ iff for some t with $s \prec t$, $\mathcal{M}, t \models \varphi$.
- $\mathcal{M}, s \models E\varphi$ iff for some world t in W, $\mathcal{M}, t \models \varphi$.

Complete logic of weak and strict preference modality axiomatized in van Benthem, Girard and Roy, 2010.

Intrinsic preference change

Example (van Benthem and Liu 2007)

Someone who is indifferent between taking a trip (p) and staying at home ($\neg p$). Now his friend comes along and says

"Let's take a trip!"

'Taking' this suggestion means that any preference we might have had for staying at home is removed from the model.

Information-driven preference change

Example (Lang and van der Torre 2008)

Initially, I desire to eat sushi from this plate.

Then **I learn** that this sushi has been made with old fish.

Now I desire not to eat this sushi.

Dynamics in betterness relations

Definition

Given a modal betterness model $(\mathcal{M}, s)$ and formula φ, the *suggestion upgrade* $(\mathcal{M}_{\sharp\varphi}, s)$ is a model with the same domain, valuation, and actual world as $(\mathcal{M}, s)$, the new betterness relation is

$$\preceq^* = \preceq - \{(s,t) \mid \mathcal{M}, s \models \varphi \text{ and } \mathcal{M}, t \models \neg\varphi\}.$$

In PDL format

$$\sharp\varphi(R) := (?\varphi; R; ?\varphi) \cup (?\neg\varphi; R; ?\neg\varphi) \cup (?\neg\varphi; R; ?\varphi).$$

where R is the given input relation of non-strict betterness, while the operation $?\varphi$ tests whether the proposition φ holds.

Dynamic betterness logics

Definition (dynamic betterness language)

The *dynamic betterness language* is given by a mutually recursive syntax rule:

$\varphi := p \mid \neg\varphi \mid \varphi \wedge \psi \mid \langle\leq\rangle\varphi \mid \langle<\rangle\varphi \mid E\varphi \mid [\pi]\varphi$

$\pi := \sharp\varphi.$

Given a betterness model $\mathcal{M}$, the *truth definition* for formulas is as before, but with one new key clause for the action modality:

$(\mathcal{M}, s) \models [\sharp\varphi]\psi \quad \text{iff} \quad \mathcal{M}_{\sharp\varphi}, s \models \psi.$

Recursion axiom for suggestions $\sharp\varphi$

$$\langle\sharp\varphi\rangle\langle R\rangle\psi \leftrightarrow (\neg\varphi \wedge \langle R\rangle\langle\sharp\varphi\rangle\psi) \vee \langle R\rangle(\varphi \wedge \langle\sharp\varphi\rangle\psi).$$

A general view: PDL-programs

Definition

Betterness change programs are built from tests for modal betterness formulas, weak and strict basic order relations R, $R^<$, and the universal relation $\top$, using arbitrary unions and sequential compositions:

$$\pi := ?\varphi \mid R \mid R^< \mid \top \mid ; \mid \cup$$

These are interpreted as the standard *PDL* program operations of *test* $?\varphi$, *sequential composition* ; and *choice* $\cup$. Many further relation transformers can be defined in *PDL* format.

Example: radical revision

Definition

Given any modal betterness model $(\mathcal{M}, s)$ and formula φ, the *radical revision* $(\mathcal{M}_{\Uparrow\varphi}, s)$ is the model with relations defined as follows in *PDL*-format:

$$\Uparrow\varphi(R) := (?\varphi; R; ?\varphi) \cup (?\neg\varphi; R; ?\neg\varphi) \cup (?\neg\varphi; \top; ?\varphi).$$

Here $\top$ denotes the universal relation. Under this transformation, all φ-worlds become better than all $\neg\varphi$-worlds, whether or not they were better before, and within these two zones, the old ordering remains.

One can think of this as obeying a command.

Completeness

[van Benthem and Liu, 2007] prove that program expressions drive an automatic formulation of completeness theorems.

Theorem

Let π be a betterness change program as defined above. There is a complete dynamic betterness logic for π, and its recursion axioms for the weak and strict preference modalities can be computed automatically.

Summary

By now, we have an abstract model for betterness and *intrinsic* preference, both weak and strong, and we can deal with dynamic changes to this structure. We now move to *extrinsic* preference and its logical analysis.

Why does Alice prefer that house?

Example (Liu 2008)

Alice is going to buy a house. She considers three things: cost, quality and neighborhood, strictly in that order. All these are clear-cut: e.g., the cost is good if inside her budget, bad otherwise. Her decision is then determined by the information whether the alternatives have the desirable properties, and also by the given order of importance of the properties.

Why do we like some cake in particular?

Example (Dietrich and List 2013)

An agent faces a choice between four different cakes:

S&H: a sweet and healthy cake,
nS&H: a non-sweet and healthy cake,
S&nH: a sweet and unhealthy cake,
nS&nH: a non-sweet and unhealthy cake.

The agent's preference relies on two 'salient properties':

S: The cake is sweet. H: The cake is healthy.

Main ideas

- We think of extrinsic preference as given by priority orders of propositions that encode relevant criteria for comparing worlds or objects.
- Simple format: [de Jongh and Liu, 2009], in which betterness order is derived from linearly ordered priorities.
- We take a more general approach here, using strict partial orders of propositions inducing pre-orders of worlds, following [Andréka et al., 2002].

Priorities and dynamics

Definition

A *priority graph* $\mathscr{G} = \langle \mathbf{P}, < \rangle$ is a strictly partially ordered set of atomic propositions in propositional language L.

Definition

Let $\mathscr{G} = \langle \mathbf{P}, < \rangle$ be a priority graph, and $\mathcal{M}$ a model in which the language L defines properties of objects. The *induced betterness relation* $\preceq_{\mathscr{G}}$ is defined as follows:

$y \preceq_{\mathscr{G}} x \ := \ \forall P \in \mathbf{P}((Py \to Px) \vee \exists P' \ P < P'(P'x \wedge \neg P'y))$.

Lexicographic ordering

For *total orders* $\mathscr{G}$, this reduces to lexicographic ordering:

$$y \preceq^{lin}_{\mathscr{G}} x \ := \ \forall P \in \mathbf{P}(Px \leftrightarrow Py) \lor \exists P' \in \mathscr{G}(\forall P < P'(Px \leftrightarrow Py) \land (P'x \land \neg P'y)).$$

We refer to [Andréka et al., 2002], [Liu, 2008] for mathematical theory of priority graphs.

Representation theorem

Theorem

Let $\mathcal{M} = (W, \preceq, V)$ be any modal model, without constraints on its relation. The following two statements are equivalent:

(a) *The relation $\preceq$ is a reflexive and transitive order,*

(b) *There is a priority graph $\mathscr{G} = (\mathbf{P}, <)$ such that, for all worlds $x, y \in W$, $y \preceq x$ iff $y \preceq_{\mathscr{G}} x$.*

If we start with a given betterness order, we can always find a priority graph that derives it. In other words, a given betterness order can always be rationalized.

Basic operations on priority graphs

[Andréka et al., 2002] has two basic operations:

- the *sequential composition* $\mathscr{G};\mathscr{G}'$ adds the graph $\mathscr{G}$ on top of $\mathscr{G}'$ in the order: all nodes in the first come before all those in the second,
- the *parallel composition* $\mathscr{G}||\mathscr{G}'$ is the disjoint union of the graphs $\mathscr{G}$ and $\mathscr{G}'$, without any order links between them.

Basic graph updates

Definition

Let 'A' be the priority graph with one single node A. The set $\alpha(\mathscr{G}, A)$ of *basic graph updates* is defined by:

$$\alpha(\mathscr{G}, A) := A \mid \mathscr{G}_1 ; \mathscr{G}_2 \mid \mathscr{G}_1 \| \mathscr{G}_2.$$

Definition

The *top deletion* of a (non-empty) priority graph $\mathscr{G}$ deletes all propositions that are not dominated by another in the graph order, leaving the rest in their old order to produce the new graph $del(\mathscr{G})$.

Valid algebraic equations

Fact

1. $\mathscr{G} ; \mathscr{G} \equiv \mathscr{G}$.
2. $\mathscr{G} \parallel \mathscr{G} \equiv \mathscr{G}$.
3. $\mathscr{G}_1 \parallel \mathscr{G}_2 \equiv \mathscr{G}_2 \parallel \mathscr{G}_1$.
4. $(\mathscr{G}_1 \parallel \mathscr{G}_2)^{<} \equiv (\mathscr{G}_1^{<} \parallel \mathscr{G}_2) \cup (\mathscr{G}_1 \parallel \mathscr{G}_2^{<})$.
5. $(\mathscr{G}_1 ; \mathscr{G}_2)^{<} \equiv (\mathscr{G}_1^{<} \cup (\mathscr{G}_1 \parallel \mathscr{G}_2^{<}))$.

This algebra is decidable and in fact, translatable into the two-variable fragment of FOL.

Modal logic of graph-induced relations

Definition

Consider a set $\mathtt{Prop}$ of propositional variables p, and a set $\mathtt{Nom}$ of nominals n. Let $\mathbb{G}$ be a set of priority graphs $\mathcal{G}$. The *modal graph language* is defined by the following syntax rule:

$$\varphi \;:=\; n \mid p \mid \neg\varphi \mid \psi \wedge \varphi \mid \langle\mathcal{G}\rangle^{\leq}\varphi \mid \langle\mathcal{G}\rangle^{<}\varphi \mid E\varphi.$$

$$\mathcal{G} \;:=\; \mathcal{G}_1;\mathcal{G}_2 \mid \mathcal{G}_1 \parallel \mathcal{G}_2.$$

[Girard, 2008] axiomatizes this modal graph logic.

Summary

The modal graph language can be seen as an extension of the modal betterness language, with consideration of reasons for preference. Dynamic changes can take place naturally in priority graphs, too, and they induce a complete dynamic logic.

Relating betterness and priority dynamics

Definition

Let α: $(\mathscr{G}, A) \to \mathscr{G}'$, with $\mathscr{G}$, $\mathscr{G}'$ priority graphs, and A a new proposition which is not in $\mathscr{G}$. Let σ: $(\preceq, A) \to \preceq'$ be a map with $\preceq$ and $\preceq'$ betterness relations over worlds. We say that α *induces* σ, if:

$$\sigma(\preceq_{\mathscr{G}}, A) \quad = \quad \preceq_{\alpha(\mathscr{G}, A)}$$

We call *the operation* α *PDL-definable* if it induces a relation transformer σ that is *PDL*-definable in the earlier format.

Cases of harmony

Fact

Taking a suggestion A given some betterness relation over worlds is induced by the following basic graph update at the priority level: $\mathcal{G}\|A$. More precisely, the following diagram commutes:

$$\begin{array}{ccc} \langle \mathcal{G}, < \rangle & \xrightarrow{\|A} & \langle (\mathcal{G}\|A), < \rangle \\ \downarrow & & \downarrow \\ \langle W, \preceq \rangle & \xrightarrow{\sharp A} & \langle W, \sharp A(\preceq) \rangle \end{array}$$

Cases of harmony

Fact

Prefixing a new proposition A to a priority graph $(\mathscr{G}, <)$ induces the radical upgrade operation $\Uparrow A$ on possible worlds models. More precisely, the following diagram commutes:

$$
\begin{array}{ccc}
\langle \mathscr{G}, < \rangle & \xrightarrow{A;\mathscr{G}} & \langle (A;\mathscr{G}), < \rangle \\
\downarrow & & \downarrow \\
\langle W, \preceq \rangle & \xrightarrow{\Uparrow A} & \langle W, \Uparrow A(\preceq) \rangle
\end{array}
$$

Basic graph updates

Theorem

Basic graph updates induce PDL-betterness transformers.

Lemma

All basic graph updates reduce to a finite set of cases.

Up to graph equivalences, all basic graph updates reduce to the five cases A, $\mathscr{G}$, $A;\mathscr{G}$, $\mathscr{G};A$, and $A \parallel \mathscr{G}$. They are closed under operations ; and $\parallel$. All these operations indeed induce *PDL*-definable betterness transformers.

Obstacles to a complete match

Fact

The deletion operation del($\mathscr{G}$) is not PDL-definable.

Fact

Not all PDL-definable operations are graph-definable.

Proof.

Here is a counter-example. Not all betterness transformers preserve the base properties of reflexivity and transitivity. To see this, consider $?A; R$, that is: 'keep the old relation only when A is true'. This does not preserve reflexivity, as $\neg A$-worlds have no relations any more. So this relation-transformer cannot be defined using a partial priority graph. □

Two-levels needed

In general then, both levels, betterness order of worlds and priority order of criteria, have their own intrinsic dynamics of preference change, supporting separate logical systems.

For a recent more general development of the two level view: see van Benthem 2015 on ‘tracking information’.

Back to Von Wright

- Ceteris paribus preference: Von Wright discussed it extensively in his book, see also [van Benthem, Girard and Roy 2010]
- Two basic topics that were not mentioned by Von Wright:

  (a) Entanglement of preference and belief

  (b) Group behavior

Example: three options to define preference

In discussing the options, we use priority sequence:

Low cost (C) > high quality (Q) > good neighborhood (N).

Definition (1)

Given a priority sequence of length n, **Pref(x,y)** is defined as follows:

$$Pref_1(x,y) ::= BC_1(x) \wedge \neg BC_1(y),$$
$$Pref_{k+1}(x,y) ::=$$
$$Pref_k(x,y) \vee (Eq_k(x,y) \wedge BC_{k+1}(x) \wedge \neg BC_{k+1}(y)), k < n,$$
$$Pref(x,y) ::= Pref_n(x,y),$$

where $Eq_k(x,y)$ stands for $(BC_1(x) \leftrightarrow BC_1(y)) \wedge \cdots \wedge (BC_k(x) \leftrightarrow BC_k(y))$.

Alice favors (1): She looks at what information she can get, she reads that d_1 has a low cost, about d_2 there is no information. This immediately makes her decide for d_1. She will stick to this decision no matter what she may hear about quality or neighborhood.

In formulas: From $B_aC(d_1)$ and $\neg B_aC(d_2)$, we get $Pref_a(d_1, d_2)$.

Definition (2)

Given a priority sequence of length n, **Pref**(**x,y**) is defined below:

$$Pref_1(x,y) ::= BC_1(x) \wedge B\neg C_1(y),$$
$$Pref_{k+1}(x,y) ::= Pref_k(x,y) \vee (Eq_k(x,y) \wedge BC_{k+1}(x) \wedge B\neg C_{k+1}(y)), k < n,$$
$$Pref(x,y) ::= Pref_n(x,y)$$

where $Eq_k(x,y)$ stands for $(BC_1(x) \leftrightarrow BC_1(y)) \wedge (B\neg C_1(x) \leftrightarrow B\neg C_1(y)) \wedge \cdots \wedge (BC_k(x) \leftrightarrow BC_k(y)) \wedge (B\neg C_k(x) \leftrightarrow B\neg C_k(y))$.

Bob favors (2): He gets the same information. But he has no preference, and that will remain so as long as he hears nothing about the cost of d_2, no matter what he hears about quality or neighborhood.

In formulas: $B_bC(d_1)$, and $\neg B_bC(d_2)$, BUT not $B_b\neg C(d_2)$. Bob cannot make his decision.

Definition (3)

Given a priority sequence of length n, **Pref(x,y)** is defined below:

$$Better_1(x,y) ::= C_1(x) \wedge \neg C_1(y),$$
$$Better_{k+1}(x,y) ::=$$
$$Better_k(x,y) \vee (Eq_k(x,y) \wedge C_{k+1}(x) \wedge \neg C_{k+1}(y)), k < n,$$
$$Better(x,y) ::= Better_n(x,y),$$
$$Pref(x,y) ::= B(Better(x,y)),$$

where $Eq_k(x,y)$ stands for
$(C_1(x) \leftrightarrow C_1(y)) \wedge \cdots \wedge (C_k(x) \leftrightarrow C_k(y))$.

Carol favors (3): She also has the same information. On that basis Carol cannot decide. But some additional information about quality and neighborhood helps her. For instance, when she hears that d_1 is of good quality or is in a good neighborhood, and d_2 is not of good quality and not in a good neighborhood. Then Carol believes that, no matter what, d_1 is superior, so d_1 is her preference. Note that this additional information would not help Bob.

In formulas: From $C(d_1)$, $Q(d_1) \vee N(d_1), \neg Q(d_2) \wedge \neg N(d_2)$, we get $Supe(d_1, d_2)$, so $Pref_c(d_1, d_2)$.

Some comments

- Three different “procedures" to obtain preference;
- Among the above three definitions, we have (2)$\rightarrow$ (1) and (2)$\rightarrow$ (3), but (1) and (3) are incomparable.

Entanglement of preference and belief

- Reasoning about preference under uncertainties: various ways to introduce beliefs into definition of preference
- Connect to classical decision theory and game theory
- Alternatives for standard optimization views (‘satisficing’, ‘good enough’) supported by logical analysis
- Fine-structure of deliberation and decision as a dynamic logical activity

From individual to group behavior

- Social choice theory, preference aggregation
- Modelling social relations: current logics of social networks (e.g. epistemic friendship logic): [Girard, Liu and Seligman 2013]
- Preference/belief change due to peer pressure: assuming that you prefer $\neg p$, if all your friends prefer p, how would you do? [Girard, Liu and Seligman 2014]
- Reasons for preference or belief change. [Baltag, Liu and Smets 2015]
- Long-term information and preference evolution in networks

Conclusions

- We have explored the idea that preference be represented at two levels, as betterness among worlds and as priority among propositions.
- We showed that both levels support significant logical structure, and that the two are connected, though not reducible, in interesting ways.
- The resulting picture offers a realistic view of how preferences are structured and can be changed, that applies to many areas, for instance, deontic logic.
- We pointed out a few further currently active research lines.

Andréka, H., Ryan, M., and Schobbens, P.-Y. (2002).
Operators and laws for combining preferential relations.
Journal of Logic and Computation, 12:12–53.

de Jongh, D. and Liu, F. (2009).
Preference, priorities and belief.
In T.Grüne-Yanoff and Hansson, S., editors, *Preference Change: Approaches from Philosophy, Economics and Psychology*, pages 85–108. Springer.

Doyle, J., Shoham, Y., and Wellman, M. (1991).
A logic of relative desire.
In *Proceedings of 6th International Symposium on Methodologies for Intelligence Systems*, pages 16–31.

Doyle, J. and Wellman, M. (1994).

Representing preferences as Ceteris Paribus comparatives.
Working Notes of the AAAL Symposium on Decision-Theoretic Planning.

Girard, P. (2008).
Modal Logics for Belief and Preference Change.
PhD thesis, Stanford University.

Liu, F. (2008).
Changing for the Better: Preference Dynamics and Agent Diversity.
PhD thesis, ILLC, University of Amsterdam.

Rao, A. and Georgeff, M. (1991).
Modeling rational agents within a BDI-architecture.

In Allen, J., Fikes, R., and Sandewall, E., editors, *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning*, pages 473–484. Morgan Kaufmann publishers Inc.: San Mateo, CA, USA.

van Benthem, J. and Liu, F. (2007).
Dynamic logic of preference upgrade.
Journal of Applied Non-Classical Logic, 17:157–182.

von Wright, G. (1963).
The Logic of Preference.
Edinburgh University Press.