Priority Product Update as Social Choice

Expanded version, November 2007

General idea

Product update is social choice between relations in an initial model ***M*** and event model ***E***, where the resulting relation in ***MxE*** is the outcome of either treating the two as equally important, or taking them in some hierarchy. This fits with the provocation in my June note that update/revision over time amounts to integrating signals. 'I' am the social aggregation of my original self plus all signals received (depending on the manner in which I took them).

Abstract setting: ordering products given component orders

Two ordered sets *(A, R)* and *(B, S)* are given, think if you wish of some initial plausibility model ***M*** and an event model ***E***, and we seek *to order the product AxB* via some suitable relation *O(R, S)*. In general, we need to order a subset of the full product where some constraint may rule out some pairs ('think preconditions', but the precise nature is irrelevant). Alexandru and Sonja worked with connected orders, but I rephrase things in terms of *pre-orders*, in line with Andréka, Ryan & Schobbens's 2002 analysis of preference merge. My analysis is close to theirs – which itself is just an instance of the 'standard tinkering' in social choice theory (and generalized quantifier theory, and related areas).

Note This product setting generalizes the usual setting of preference merge, where the given individual preferences are over the same domain. I think a pair, or general vector, phrasing is of more general use anyway, also in social choice. We are really socially constructing *two things*: both (a) a vector space (the 'social space' of all relevant group outcomes) plus (b) the social order over the latter. Conversely, we can view the product case as falling under one with a constant set to be ordered. First lift the component relations to pairs as follows:

$$(a, c)\ R_1\ (b, d) \text{ iff } a\ R_1\ c, \qquad (a, c)\ R_2\ (b, d) \text{ iff } c\ R_2\ d,$$

and then merge in the standard social choice or belief revision style.

Priority Update Rule

Here is what I would propose, following ARS, taking ***E*** to rank 'above' ***M***:

$$(s, e) \leq (t, f) \text{ iff } (s \leq t\ \&\ e \leq f) \quad e < f.$$

This reflects the general idea in *ARS*: in principle, a pair *(x, y)* must satisfy every component relation *R*, but *y* can 'compensate failures' where ¬*Rxy* by doing strictly better than *x* with respect to some relation *R'* with higher priority than *R*. This stipulation also fits with Alexandru & Sonja's priority rule in the special case where the order is connected.

Define the strict order $x<y$ as usual as $x \leq y$ & $\neg y \leq x$. By a simple computation, we see that

Fact $(s, e) < (t, f)$ iff $(s < t \;\&\; e \leq f)$ $e < f$.

Of course, with pre-orders, we have to state our intuitions keeping in mind four distinct cases

$x < y$, $y < x$, $x \sim y$ (indifferent), and $x \# y$ (incomparable).

Modal languages

We can introduce a static modal language here with modalities for both the weak and strict versions of the ordering, plus an existential modality. At least I would find that attractive as a most direct reflection of the underlying ordering structure. The base logic of this language over pre-orders is in my *JPL* paper on preference logic with Olivier & Patrick.

On top of that comes the dynamics as usual, which requires also an existential modality.

Fact There is a complete set of reduction axioms.

E.g., $<\boldsymbol{E}, e> <\leq> \square\ \square\ PRE_e \;\&\; (\ {}_{e \leq f} <\leq> <\boldsymbol{E}, f> \square\ \ {}_{e<f}\ E <\boldsymbol{E}, f> \square)$.

Choice conditions on order 1: invariance and locality

Next comes the usual game in social choice theory: piling up conditions toward uniqueness. What follows was in my previous note, but now with a little twist for the pre-orders. I first choose a very restrictive set to zoom in exclusively on priority product update. Later on I want to relax, as we may want greater freedom in stipulating update rules eventually. [In fact, intuitively, I would like to see epistemic update as another, more democratic 'type' of social choice, even though it may technically be brought under a hierarchical priority heading.]

(a) *Permutation invariance in some obvious sense*

Any pair of permutations of A and of B leads to an obvious invariance condition on O on the models (A, R), (B, S) and their permuted versions, which forces our definitions to be *uniform*. Basically, we have that $O(\square[R], \square[S]) = \square[O(R, S)]$.

This is the usual structural 'logicality' condition in many areas, e.g., in generalized quantifier theory and studies of logical constants, where its effects are well-known.

(b) *Locality*

$O(R, S)((a, b), (a', b'))$ iff $O(R|\{a, a'\}, S|\{b, b'\})((a, b), (a', b'))$.

This is a very strong version of 'Independence from Irrelevant Alternatives'.

Remark Indeed, *(b)* will not hold for belief revision policies like 'conservative update' $\uparrow A$, where we need to check whether some object is 'maximal in A' with respect to the given order, a non-local property. Of course, one could cheat by treating this as radical update $\square(bestA)$ with the artificial 'signal' "*bestA*" defined as $A \;\&\; \neg<<>A$. If we do not go that way, we need to really complicate the social choice analysis, no longer assuming Locality.

Table representation

Together, *(a), (b)* force our *O* to be definable by just its behaviour in a 4x4-Table:

	S on b, b'	□	□	□	#
R on a, a'	□	-	-	-	-
	□	-	-	-	-
	□	-	-	-	-
	#	-	-	-	-

Here the entries stand for the four isomorphism classes on two distinct objects. The relation between the pairs only depends on these. [For the case of singletons, see below. Well, let me say it rightaway: if singletons are reflexive, agents are indifferent between any *x* and *x*, and hence by the Abstentions principle below, we only need to look at the other pair relation.]

Caveat We are going to fill in the same four types of entries in the Table, though strictly speaking, one might just want to put YES/NO in the slots for whether the relation ≤ holds between the pairs. In using the four types, strictly speaking, one should check that all intuitions I am going to state for these hold for Priority Product Update as defined above.

Further choice conditions

Now we rule out all but two combinations by the following conditions:

*(**c**)* *Abstentions*
If a subgroup votes indifferent (□), then the others determine the outcome.

This is a standard condition. The next one is not:

*(**d**)* *Closed Agenda*
The social outcome is always among the opinions of the voters.

Remark This implies the usual condition of *Unanimity*: "if all members of a group agree, then the outcome is their shared outcome". But it is much stronger.

Finally, we state what happens when agents do care about outcomes. An 'over-rule' is a case where voters are not indifferent, and one's opinion wins over the other's.

*(**e**)* *Overruling*
If an agent's opinion overrules that of another, then her opinion *always* does.

This goes against the spirit of democracy and letting everyone win once in a while, but we should not hide the fact that this is what the Priority rule does with its bias toward the last event . [Also, the ***E***-relation may be quite weak, leaving lots of issues open. In social choice terms, an authority without strong views is not much of a brake on the powers of others…]

Characterizing priority update

Theorem 1 A preference aggregation function is a priority update iff it satisfies Permutation Invariance, Locality, Abstentions, Closed Agenda, and Overruling.

Proof We fill in the Table using these conditions. First, the diagonal is clear by Unanimity, and the row and column for the indifference case by Abstentions:

	S on b, b'	□	□	□	#
R on a, a'	□	□	*1*	□	*2*
	□	*3*	□	□	*4*
	□	□	□	□	#
	#	*5*	*6*	#	#

This leaves six slots to be filled, or really three, by permutation invariance: e.g., case □ □ induces case □ □ . Consider Case 1. By Closed Agenda, this must be either □ or □. Without loss of generality, consider the latter: *S* overrules *R*. Then using Overruling to fill the other cases with *S*'s opinion, and applying permutation Invariance, our Table looks like this:

	S on b, b'	□	□	□	#
R on a, a'	□	□	□	□	#
	□	□	□	□	#
	□	□	□	□	#
	#	□	□	#	#

It is easy to see that this final diagram is precisely that for Priority Update as defined above. The other possible case would give preference to the ordering on ***M***.

Remark Both are instances of the 'But' operator of *ARS*, i.e., a Leader/Follower pattern.

This concludes our proof. □

Weaker conditions, more update rules

Now we relax the above conditions to also allow the democratic variant where the opinions of ***M***, ***E*** count equally, a *flat update*. This intersection of relations is the 'And' of *ARS*. Now Closed Agenda fails: e.g., the above clash □ □ ends up in #. Instead, we take principles of

(f) *Unanimity*
If voters all agree, their vote is the social vote.

(g) *Alignment*
'if anyone changes their vote to get closer to the group outcome, the group outcome does not change'.

Theorem 2 A preference aggregation function is either priority update or flat update iff it satisfies Permutation Invariance, Locality, Abstentions, Overruling, Unanimity, and Alignment.

The crucial step in the proof is now that, without Closed Agenda, Slot 6 in our diagram

S on b, b'		□	□	□	#
R on a, a'	□	□	6	□	2
	□	3	□	□	4
	□	□	□	□	#
	#	5	6	#	#

may also have entries ~ or #. However, the former outcome can be ruled out by Alignment. If *S* were to change its opinion to ~, the outcome would still have to be ~, but it is □. So, the outcome must be #. But then, using Alignment once more, for both voters (plus some Permutation Invariance), we see that all remaining slots must be #*:*

S on b, b'		□	□	□	#
R on a, a'	□	□	#	□	#
	□	#	□	□	#
	□	□	□	□	#
	#	#	#	#	#

This is clearly the table for the flat update. □

Discussion

* Clearly, we can get still more update rules if we relax the choice conditions. As said before, Overruling may go against the grain of democracy, and if we drop it, we would create room for 'mixtures' of influence for ***M*** and ***E***, maybe in line with Carnap's inductive logic and learning theories [I made this point before, and see also my paper with Barteld and Jelle].

* Here is an issue from my earlier Note which I am still not clear about. Let's compare priority update with flat update. *Which rule is more general in its dynamic effects?* E.g., Priority Update can never make an established strict preference for *x* over *y* 'indifferent' again, while Democracy can mimick any effect of Priority by interpolating more signals. [Maybe I am confused here. I guess I want to know which relation changes can be achieved, fixing the update rule, on models ***M*** by executing suitable sequences of event models.]

* Extend to *plausibility merge between more relations* for 'belief merge'?