Choosing the right concept of “right choice”

Frederik Van De Putte
Centre for Logic and Philosophy of Science
Ghent University
Blandijnberg 2, 9000 Gent, Belgium
`frvdeput.vandeputte@UGent.be`

March 16, 2018

This is a technical report; to be replaced by a full-fledged paper in due course. It is perhaps preliminary, sketchy, incomplete.

The note concerns the well-known topic of the relation between the rightness of actions for a group and the rightness of actions for that group's members. We focus on one direction, asking when every member of the group doing a right action suffices for the group to do a right action. We moreover approach the matter from a somewhat more abstract viewpoint than Horty's in his [1]. That is, we look at various ways one can define “right actions”, and see under what conditions the definitions yield certain bridge principles, from member to group rightness.

Thanks to Jeff Horty, Eric Pacuit, and others at UMD for discussions on this topic. Also thanks to Hein Duijf and Allard Tamminga for comments on a previous version.

1 Definition of Models

Definition 1 (Models) *A* deontic free choice model *is a tuple*

$$M = \langle W, \langle \sim_i \rangle_{i \in N}, P, V \rangle$$

where

(C1) $W \neq \emptyset$
(C2) Every $\sim_i \subseteq W \times W$ *is an equivalence relation*
(C3) $\emptyset \neq P \subseteq W$
(C4) $V : \mathfrak{P} \to \wp(W)$ *is a valuation function*
(IOA) for all $w_1, \ldots, w_n \in W$*, there is a* w *such that* $w_1 \sim_1 w$*, …, and* $w_n \sim_n w$

We call W *the* domain *of* M*.*

Where $M = \langle W, \langle \sim_i \rangle_{i \in N}, P, V \rangle$ is given, we define the relations $\sim_G$ for $G \subseteq N$ as the intersection of the relations $\sim_i$ for all members i of G: $\sim_G = \bigcap_{i \in N} \sim_i$. We define the relation $\sim_\emptyset$ as the total relation over W: $\sim_\emptyset = W \times W$.

Where $G \subseteq N$, let $Choice_G(M) = \{\{v \mid w \sim_G v\} \mid w \in W\}$. We refer to $Choice_G(M)$ as the set of choices *of G in M.*

M is deterministic *iff for all $w, w' \in W$, $w \sim_N w'$ iff $w = w'$.*

In the remainder we use $G, G', \ldots$ as metavariables for subsets of N.

When working with only two agents, a finite deontic free choice frame, i.e. a deontic free choice model without the valuation, can be represented by means of a matrix, where the rows represent choices of one agent, columns represent choices of the other agent, and each occurrence of a 1 or 0 represents a single possible world; 1 indicates that the world in question is permitted, 0 indicates that it is not permitted. See Figure 1 for a simple example of a deterministic deontic free choice model.

2 Impossible to have one's cake and eat it too

Consider an arbitrary property $\mathbb{P}$ that can be attributed to choices of (group) agents in any given model M. So formally, $\mathbb{P}(M)$ is a set of choices of various agents, $\mathbb{P}(M) \subseteq \bigcup_{G \subseteq N} Choice_G(M)$. We say that $\mathbb{P}$ is *total* iff, for every deontic free choice model $M = \langle W, \langle \sim_i \rangle_{i \in N}, P, V \rangle$, and for every group G, there is at least one $X \in Choice_G(M) \cap \mathbb{P}(M)$. $\mathbb{P}$ is *safe for aggregation* iff, for every deontic free choice model $M = \langle W, \langle \sim_i \rangle_{i \in N}, P, V \rangle$ and all groups G, G' (with $G \cap G' = \emptyset$): if $X \in Choice_G(M) \cap \mathbb{P}(M)$ and $X' \in Choice_{G'}(M) \cap \mathbb{P}(M)$, then $X \cap X' \cap P \neq \emptyset$.

$\mathbb{P}$ is *aggregative* iff, for every deontic free choice model $M = \langle W, \langle \sim_i \rangle_{i \in N}, P, V \rangle$ and all groups G, G' (with $G \cap G' = \emptyset$): if $X \in Choice_G(M) \cap \mathbb{P}(M)$ and $X' \in Choice_{G'}(M) \cap \mathbb{P}(M)$, then $X \cap X' \in \mathbb{P}(M)$.

Finally, $\mathbb{P}$ is *label-independent* iff whether a given choice of a given group G has property $\mathbb{P}$ in a model M depends neither on the labels we attach to the worlds making up that choice, nor on the labels we attach to the members of G. Part of this property is sometimes called "anonymity" in the game theoretic literature; spelling it out in full detail is a bit tedious but feasible (see Appendix A).

Now one may ask (Jeff Horty, in a discussion at UMD): could there ever be some property $\mathbb{P}$ that satisfies all three of the above defined conditions, hence, a $\mathbb{P}$ that is total, safe for aggregation, and label-independent?

1	0
0	1

Figure 1: A coordination problem for two agents.

The answer is simply negative. It suffices to take a look at Figure 1 to see why. Any label-independent property $\mathbb{P}$ will either hold for all the actions of the two agents in question, or for neither (since the frame is completely symmetric, both in the agents and in the choices they have). If such a $\mathbb{P}$ is moreover total, it follows that all actions have property $\mathbb{P}$. But then $\mathbb{P}$ is not safe for aggregation.

3 Preferences and choices

3.1 Preference Relations on Sets and Choices

Fix a deontic free choice model $M = \langle W, \langle \sim_i \rangle_{i \in N}, P, V \rangle$. Let $X, Y \subseteq W$. Then[1]

$$X \preceq Y \text{ iff } X = Y \text{ or } X \subseteq P \text{ or } Y \cap P = \emptyset$$

[Another way to define $\preceq$ is as follows. Define $d : W \to \{0, 1\}$ as follows: $d(w) = 1$ if $w \in P$ and $d(w) = 0$ otherwise. Then $X \preceq Y$ iff $X = Y$ or $\forall x \in X, \forall y \in Y$, $d(x) \geq d(y)$.]

Note: Intuitively, $X \preceq Y$ means that X is "better" than Y. I will stick to this way of representing ranking relations between sets of worlds and actions in the remainder (so the element that comes before the relational symbol is always "better").

Obviously, $\preceq$ is reflexive. It is also transitive when defined over a set of non-empty sets (so e.g., it is transitive over all sets of the type $Choice_G(M)$). This implies that for finite models (where W is finite, and hence also $\wp(W)$ is finite), it will also be smooth (every chain of ever better states has a last/best element).

Let now $X, Y \in Choice_G(M)$. Then we define:

$$X \sqsubseteq Y \text{ iff } \forall Z \in Choice_{N \setminus G}(M), X \cap Z \preceq Y \cap Z$$

$$X \sqsubset Y \text{ iff } X \sqsubseteq Y \text{ and } Y \not\sqsubseteq X$$

3.2 Various concepts of "right" choice

Translated to the present minimalistic setting, Horty's original account of "dominance act utilitarianism" (Chapter 4, Section 2 of his 2001 book) is based on the idea that we should consider those actions right that are not strongly dominated by any other action. This gives us the following definition:[2]

Definition 2 (Maximal actions) *Let $X \in Choice_G(M)$. Then X is* maximal *for G in M iff there is no $Y \in Choice_G(M)$ such that $Y \sqsubset X$.*

[1] In principle, all the preference relations defined below should be indexed with M. I omit this index and assume a fixed M that is given by the context.

[2] I follow terminology by Xavier Parent, who in his turn follows Sen, in distinguishing between "optimality" and "maximality". Mind that what Horty calls "optimal" in his 2001 book is what we call "maximal" here; what we call "optimal" is not considered in his book.

As a property of choices, "maximal" is total (at least when we restrict the focus to finite models), but it is not safe for aggregation and it is also not aggregative. This is a well-known fact, cf. [1].

We can however also decide to consider only those actions right that weakly dominate every other action of the group in question. Formally:

Definition 3 (Optimal actions) *Let $X \in Choice_G(M)$. Then X is* optimal *for G in M iff for all $Y \in Choice_G(M)$, $X \sqsubseteq Y$.*

"Optimal" is aggregative and safe for aggregation (see Section 6), but it is not total – there is not always an optimal action for a given group G. An example is the simple coordination problem for two agents given by Figure 1.

Finally, we can also call "right" those choices X of G such that, whatever the agents in $N \setminus G$ do, if the combination Y of what they do is weakly permitted, then the combination of Y with X is also weakly permitted. These are the choices that I called "freely permitted choices" in previous unpublished work; following Jeff's remarks I rename them to "safe" actions. Formally:

Definition 4 (Safe actions) *$X \in Choice_G(M)$ is* safe *for G in M iff for all $Z \in Choice_{N \setminus G}(M)$, if $X \cap Z \cap P = \emptyset$, then $Z \cap P = \emptyset$.*

"Safe" is aggregative and safe for aggregation. Consequently, "safe" is not total; again Figure 1 serves as an illustration.

4 Alternative Characterization of "safe" actions

Let

$$X \preceq' Y \text{ iff } X = Y \text{ or } X \cap P \neq \emptyset \text{ or } Y \cap P = \emptyset$$

Note that $\preceq'$ is reflexive and transitive. Also, where $X \neq \emptyset$: if $X \preceq Y$ then $X \preceq' Y$ (†). Where X is a singleton, ($\star$) $X \preceq Y$ iff $X \preceq' Y$.

Let $X \sqsubseteq' Y$ iff $\forall Z \in Choice_{N \setminus G}(M), X \cap Z \preceq' Y \cap Z$.

We can now define the set of safe actions equivalently in terms of the preference relation $\sqsubseteq'$ over the actions of G:

Theorem 1 *Let $X \in Choice_G(M)$. Then X is safe for G in M iff for all $Y \in Choice_G(M)$, $X \sqsubseteq' Y$.*

Note that this alternative characterization is obtained by replacing $\sqsubseteq$ with $\sqsubseteq'$ in the definition of optimal actions. In Section 6 we investigate a property of the class of all conceptions of "right" that can be defined in this way.

1 1	1 0
0 0	0 0

Table 1: A non-deterministic deontic (free) choice frame for two agents, where the choice "right" is safe but not maximal, and hence also not optimal for the column-choosing agent.

5 Relations between the various concepts

The following can easily be checked (and it relies on no specific properties of the relations $\preceq$ and $\sqsubseteq$):

Theorem 2 *If G has any optimal actions in M, then X is an optimal action for G in M iff X is a maximal action for G in M.*

By Theorem 1 and (†) from the previous section, we have:

Theorem 3 *If X is an optimal action for G in M, then X is a safe action for G in M.*

Relying on Theorem 1 and ($\star$) from the previous section, we can also derive the following:

Theorem 4 *If M is deterministic, then X is a safe action for G in M iff X is an optimal action for G in M.*

In general, not every safe action for G in M is also maximal for G in M. We illustrate this by means of Table 1. In this table, we use 0 and 1 to refer to worlds in a choice cell; so if e.g. there is a 1 and a 0 in one cell, this means that cell contains both an acceptable world and one that is not acceptable (deontically speaking). In Table 1, both "left (column)" and "right (column)" are safe, but only "left" is maximal and optimal. "Right" is neither maximal nor optimal for the column-choosing agent.

In view of the preceding, one might expect that the optimal actions are exactly the ones that are both maximal and safe. But this is also not true, in view of the example from Table 2. In that example, all actions of all agents are both maximal and safe, but none of those actions are optimal.

However, for deterministic models it can easily be shown that the optimal actions are exactly the ones that are both maximal and safe. This is a simple corollary of Theorems 4 and 2.

Next, one may wonder whether it can ever be the case that a model has safe actions for a given group G, but that nevertheless, none of its safe actions are maximal. This turns out to be impossible, at least for finite models:

1, 1	1, 0
1, 0	1, 1

Table 2: A non-deterministic deontic (free) choice frame for two agents, where all actions are maximal and safe, but no actions are optimal.

1, 1	1, 0
0, 0	1, 0

Table 3: A non-deterministic deontic (free) choice frame for two agents. Here, "left column" is maximal for the colmun-choosing agent, but not safe, and yet there are safe (and maximal) actions for that agent, viz. "right column".

Theorem 5 *Let M be finite. If there are safe actions for G in M, then there are also actions that are both safe and maximal for G in M.*

Proof. Let X be a safe action for G in M. Since G has only finitely many actions, we can easily show that there is a maximal action $X' \in Choice_G(M)$ such that $X' \sqsubset X$. Consider now an arbitrary $Y \in Choice_{N \setminus G}$ such that $Y \cap P \neq \emptyset$. By the supposition, $Y \cap X \cap P \neq \emptyset$ (‡). Note that, since $X' \sqsubset X$, $X' \cap Y = X \cap Y$ or $X' \cap Y \subseteq P$ or $X \cap Y = \emptyset$. The first disjunct is excluded since X and X' must be disjoint in order to be distinct, and since by (IOA), $X' \cap Y \neq \emptyset$ and $X \cap Y \neq \emptyset$. The third disjunct is excluded in view of (‡). From this we can infer by (IOA) that $X' \cap Y \cap P \neq \emptyset$. Since Y was arbitrary, it follows that also X' is safe for G in M. ■

A final question that remains is: if there are safe actions (and hence there are also actions that are both safe and maximal), does it follow that every maximal action is also safe? The answer is negative – see Table 3.

6 "optimal" is aggregative

Fix an arbitrary preference relation $\ll \subseteq Choice_N(M) \times Choice_N(M)$ that is transitive:

$$\text{for all } X, Y, Z \in Choice_N(M), \text{ if } X \ll Y \text{ and } Y \ll Z, \text{ then } X \ll Z$$

We generalize Definition 3 as follows:

Definition 5 ($\ll$-Optimal actions) *Let $X \in Choice_G(M)$. Then X is* $\ll$-optimal *for G in M iff for all $Y \in Choice_G(M)$ and for all $Z \in Choice_{N \setminus G}$, $X \cap Z \ll Y \cap Z$.*

Theorem 6 *"$\ll$-optimal" is aggregative. That is, where $G \cap G' = \emptyset$: if $X \in Choice_G(M)$ is $\ll$-optimal for G in M and $X' \in Choice_G(M)$ is $\ll$-optimal for G' in M, then $X \cap X'$ is $\ll$-optimal for $G \cup G'$ in M.*

Proof. Suppose the antecedent holds. Let $Y \in Choice_{G \cup G'}(M)$ and $Z \in Choice_{N \setminus (G \cup G')}(M)$ be arbitrary. Let Y_G be the unique element of $Choice_G(M)$ that is a superset of Y; let $Y_{G'}$ be the unique element of $Choice_{G'}(M)$ that is a superset of Y.

Note that $X' \cap Z \in Choice_{N \setminus G}$. So by the $\ll$-optimality of X for G,

$$X \cap (X' \cap Z) \ll Y_G \cap (X' \cap Z) \tag{1}$$

Note that $Y_G \cap Z \in Choice_{N \setminus G'}$. So by the $\ll$-optimality of X' for G',

$$X' \cap (Y_G \cap Z) \ll Y_{G'} \cap (Y_G \cap Z) \tag{2}$$

Note that the right hand side of (1) and the right hand side of (2) refer to the same set. Since $\ll$ is transitive over $Choice_N(M)$ and since $Y = Y_G \cap Y_{G'}$, we obtain that $(X \cap X') \cap Z \preceq Y \cap Z$. Since Y and Z were arbitrary, it follows that $X \cap X'$ is $\ll$-optimal for $G \cup G'$. ■

As an immediate corollary of this theorem, it can be inferred that if every agent in a group does an $\ll$-optimal action, then the group as a whole is doing an $\ll$-optimal action.

Note that the above proof relies only on the transitivity of $\ll$. So it applies to our original notion of "optimal", to our notion of "safe" (using Theorem 1) and to a great many other notions of "right (action)". In particular, one can easily start from a richer semantics (e.g. with deontic preference relations over worlds), define a transitive $\ll$ on $Choice_N(M)$, giving us a notion of "right" that is aggregative.

Some readers may wonder whether the possibility of redefining a given notion "right" in terms of optimality for *some* transitive preference relation $\ll$ is also a necessary condition for that notion to be aggregative (perhaps, under certain conditions such as determinism). The answer is negative. Consider the two-player game in Figure 2. To feed intuitions, one may interpret the numbers in this game as payoffs, and the "rightness" function as a combination of risk-averseness and maximizing: for each group, one first considers only those actions that avoid the worst outcome, and then out of these, one picks the ones that maximize pay-off, whatever the other agents do.

If that is how we reason, then "left column" and "upper row" are the only right actions for the individual agents, but that both "upper row+left column" and "middle row+midle column" are the only right actions for the group of both agents. The resulting property $\mathbb{P}$ will be aggregative (at least in this model; cf. infra), but it cannot be defined from a transitive preference relation on the choices of the grand coalition. That is, such a transitive preference relation would have to satisfy each of the following properties:

"middle row + right column" $\ll$ "middle row + middle column" "middle row + middle column" $\ll X$ for all $X \in Choice_N(M)$; hence by transitivity,

"middle row + right column" $\ll X$ for all $X \in Choice_N(M)$

Making "middle row + right column" right for the grand coalition, counter to what we started with.

2	1	1
1	2	0
1	0	0

Figure 2: A coordination problem for two agents.

(Some may worry that this example fails to show that the rightness condition under consideration is aggregative in general, for any model M. However, formally speaking, we can just enforce this by stipulating that the rightness condition behaves exactly like "optimal" for all models that are not isomorphic to the one under consideration. This is of course an artificial construction, but formally it is well-defined. The question is whether there is some way we can distinguish between "artificial" and other rightness conditions, and whether some non-artificial rightness conditions are not representable in terms of a transitive $\ll$.)

7 A loose end

One can also ask whether the availability of $\preceq$-optimal actions is not just a sufficient, but also a necessary condition for ensuring group maximality by means of member maximality. The results I have so far seem to be mixed:

(i) with only two agents and determinism, only one agent having an optimal action, and the other agent doing something that is weakly permitted (i.e., some $X \in Choice_j(M)$ with $X \cap P \neq \emptyset$), is already sufficient to end up in a permitted state.

(ii) with three agents, we can easily construct a deterministic model where α does an optimal action, β and γ each do a maximal action, but we still end up in an impermissible state (just put the standard coordination problem left, and a game with only 0s right; α is the agent that chooses between the left and right table; β and γ are agents that choose the rows, resp. columns within the tables)

(iii) without determinism we can construct an example with two agents, one having an optimal action, the other not, where, if they both take a maximal action, the group ends up not doing a maximal action: $row_1 = \{\{0,1\},\{1\}\}, row_2 = \{\{0,1\},\{0\}\}$. The columns are both maximal but not optimal, (only) the upper row is optimal.

8 Concluding remarks

This final section is a bit more tentative. In view of Theorem 2, we can conclude that in the lucky case where you happen to have optimal actions, those actions will in fact be the only maximal ones. Hence, in such cases the best advice one can give is: take one of your maximal actions. A further conclusion would be: the concept of an optimal action is actually quite useless, since the only cases where it may be useful are already covered by the old concept of a maximal action. However, it is useful as a way to point out in which situations exactly one can ensure that group maximality is inherited from each member doing an individually maximal action.

If there are no optimal actions, there are two possibilities.

First, there are no safe actions, only maximal ones. In that case the only sensible advice one can give (without changing the rules of the game) is: do a maximal action, fully aware that if everyone does so we might still end up in disaster. This is just the best to do given the unlucky circumstances.

Second, there are safe actions. In view of Theorem 5, it follows that some of these safe actions are also maximal (assuming that the model is finite). Then it seems the best thing to do is to pick actions that are both maximal and safe. Note that these need not be optimal (see again Table 1). However, they are among the best actions each one can do individually (they are not strongly dominated), and they are moreover such that if we combine them with safe actions of the other agents, we can be sure that at least the possibility of permissibility is left open.

However, we might also be in a situation like in Table 3. In such a situation, there seems to be a genuine conflict of intuitions. Should we go for an option where we know that, whatever the other agents do, the grand coalition at least leaves open the possibility that we end up in a permissible state? Or should we rather go for an option where we know that, if we luckily happen to coordinate with all other agents, we in fact guarantee ending up in a permissible state? Who is to blame for what, in such cases? Or can't we blame anybody?

One might think the problem is solved by noting that in the example from Table 3, the row choosing agent should just go for "top", and hence this is a reason for the column choosing agent to pick "left". But that is not a satisfying answer. Table 4 shows that sometimes even this will not be of much help.

References

[1] John F. Horty. *Agency and Deontic Logic.* Oxford University Press, New York, 2001.

A Label-independence

Fix a function $f : W \to W, f : N \to N$ that is one-one. Let M be a deontic free choice model. Define the model M^f from M and f as follows:

1 1	0 0	1 0
0 0	1 1	1 0

Table 4: A non-deterministic deontic (free) choice frame for two agents. Here, "left column" is maximal for the colmun-choosing agent, but not safe, and yet there are safe (and maximal) actions for that agent, viz. "right column". Moreover, both "top" and "bottom" are maximal for the row-choosing agent, whereas neither of these two are safe.

$$M^f = \langle W, \langle \sim_i^f \rangle_{i \in N}, P^f, V^f \rangle$$

where for all $i \in N$, $\sim_i^f = \{(f(w), f(w')) \mid (w, w') \in \sim_{f(i)}\}$, $P^f = \{f(w) \mid w \in P\}$ and $V^f(p) = \{f(w) \mid w \in V(p)\}$. So M^f is structurally isomporhic to M. Let $f(G) =_{\mathsf{df}} \{f(i) \mid i \in G\}$.

$\mathbb{P}$ is label-independent iff for any deontic free choice model $M = \langle W, \langle \sim_i \rangle_{i \in N}, P, V \rangle$ and any isomorphism $f : W \to W, N \to N$, the following holds: for all $G \subseteq N$, $X \in Choice_G(M) \cap \mathbb{P}(M)$ iff $\{f(w) \mid w \in X\} \in Choice_{f(G)}(M^f) \cap \mathbb{P}(M^f)$.