Single-peakedness of preferences via deliberation: A formal study

(An extended abstract)

Sujata Ghosh
Indian Statistical Institute, Chennai, India
sujata@isichennai.res.in

1 Introduction

There are two important aspects of any democratic decision: aggregation of preferences and deliberation about preferences. They are essential and complementary components of any decision making process. While the well-studied process of aggregation focuses on accumulating individual preferences without discussing their origin [4], deliberation can be seen as a conversation through which individuals justify their preferences, a process that might lead to changes in their opinions as they get influenced by one another. Till now, there has been a lot of work on the 'aggregation' aspect (e.g., [12, 14, 6]). However, some recent work has focussed on the deliberation aspect as well [8, 9, 10, 15].

Sometimes, deliberation does not lead to unanimity in preferences, but the discussion can lead to some 'preference uniformity' (see how deliberation can help in bypassing social choice theory's impossibility results in [5]), which might facilitate their eventual aggregation. In addition, the combination of both processes provides a more realistic model for decision making scenarios. In light of this status quo, our focus is on the formal study of achieving such preference uniformities, e.g., single-peaked, single-caved, single-crossing, value-restricted, best-restricted, worst-restricted, medium-restricted, or group-separable profiles. In this short abstract we provide our preliminary ideas towards achieving single-peakedness of preference profiles via deliberation.

In what follows, we define two preference upgrade operators based on [8, 9] and provide a preliminary discussion on how single-peaked preference profiles can be achieved through such operations. We will delve into the details of the logical language in the main paper.

2 Basic concepts

The focus of this work is public deliberation, so let Ag be a *finite non-empty* set of agents with $|Ag| = n \geq 2$ (if $n = 1$, there is no scope for joint discussion). Below we present the most important definitions of this framework.

Definition 1 (*PR* frame)**.** *A* preference and reliability (*PR*) frame F *is a tuple* $\langle W, \{\leq_i, \preccurlyeq_i\}_{i \in Ag} \rangle$ *where*

- W *is a finite non-empty set of worlds;*
- $\leq_i \subseteq (W \times W)$ *is a total preorder (a total, reflexive and transitive relation), agent* i*'s* preference relation *over worlds in* W *(*$u \leq_i v$ *is read as "world* v *is at least as preferable as world* u *for agent* i*");*
- $\preccurlyeq_i \subseteq (Ag \times Ag)$ *is a total order (a total, reflexive, transitive and antisymmetric relation), agent* i*'s* reliability relation *over agents in* Ag *(*$j \preccurlyeq_i j'$ *is read as "agent* j' *is at least as reliable as agent* j *for agent* i*").*

Some further useful definitions are given below.

Definition 2. *Let* $F = \langle W, \{\leq_i, \preccurlyeq_i\}_{i \in Ag} \rangle$ *be a PR frame.*

- $u <_i v$ *("*u *is less preferred than* v *for agent* i*") iff*$_{def}$ $u \leq_i v$ *and* $v \not\leq_i u$.
- $u \simeq_i v$ *("*u *and* v *are equally preferred for agent* i*") iff*$_{def}$ $u \leq_i v$ *and* $v \leq_i u$.
- $j \prec_i j'$ *("*j *is less reliable than* j' *for agent* i*") iff*$_{def}$ $j \preccurlyeq_i j'$ *and* $j' \not\preccurlyeq_i j$.
- $\mathrm{mr}(i) = j$ *(*j *is agent* i*'s most reliable agent) iff*$_{def}$ $j' \preccurlyeq_i j$ *for every* $j' \in Ag$.
- $\mathrm{Max}_{\leq_i}(U)$*, the set containing agent* i*'s most preferred worlds among those in* $U \subseteq W$*, is formally defined as* $\{v \in U \mid u \leq_i v \text{ for every } u \in U\}$.

3 Preference dynamics: lexicographic upgrade

Intuitively, a public announcement of the agents' individual preferences might induce an agent i to adjust her own preferences according to what has been announced and the reliability she assigns to the set of agents.[1] Thus, agent i's preference ordering *after* such announcement, $\leq'_i$, can be defined in terms of the just announced preferences (the agents' preferences *before* the announcement, $\leq_1, \ldots, \leq_n$) and how much i relied on each agent (i's reliability *before* the announcement, $\preccurlyeq_i$): $\leq'_i := f(\leq_1, \ldots, \leq_n, \preccurlyeq_i)$ for some function f. Below, we define a general upgrade operation based on agent reliabilities from [8].

[1] Note that we do not study the formal representation of such announcement, but rather the representation of its effects.

Definition 3 (General lexicographic upgrade)**.** *A* lexicographic list $\mathcal{R}$ *over* W *is a finite non-empty list whose elements are indices of preference orderings over* W*, with* $|\mathcal{R}|$ *the list's length and* $\mathcal{R}[k]$ *its* k*th element* ($1 \leq k \leq |\mathcal{R}|$)*. Intuitively,* $\mathcal{R}$ *is a priority list of preference orderings, with* $\leq_{\mathcal{R}[1]}$ *the one with the highest priority. Given* $\mathcal{R}$*, the preference ordering* $\leq_{\mathcal{R}} \subseteq (W \times W)$ *is defined as*

$$u \leq_{\mathcal{R}} v \quad \textit{iff}_{def} \quad \underbrace{\Big(u \leq_{\mathcal{R}[\,|\mathcal{R}|\,]} v \;\wedge\; \bigwedge_{k=1}^{|\mathcal{R}|-1} u \simeq_{\mathcal{R}[k]} v\Big)}_{1} \;\vee\; \underbrace{\bigvee_{k=1}^{|\mathcal{R}|-1} \Big(u <_{\mathcal{R}[k]} v \;\wedge\; \bigwedge_{l=1}^{k-1} u \simeq_{\mathcal{R}[l]} v\Big)}_{2}$$

Thus, $u \leq_{\mathcal{R}} v$ *holds if this agrees with the least prioritised ordering* ($\leq_{\mathcal{R}[|\mathcal{R}|]}$) *and for the rest of them* u *and* v *are equally preferred (part 1), or if there is an ordering* $\leq_{\mathcal{R}[k]}$ *with a strict preference for* v *over* u *and all orderings with higher priority see* u *and* v *as equally preferred (part 2).*

Proposition 1. *Let* $\mathcal{R}$ *be a lexicographic list over* W*. If* every *ordering* $\mathcal{R}[k]$ ($1 \leq k \leq |\mathcal{R}|$) *is reflexive (transitive, total, respectively), then so is* $\leq_{\mathcal{R}}$*.*

As a consequence of this proposition, the general lexicographic upgrade preserves total preorders (and thus our class of semantic models) *when every preference ordering in* $\mathcal{R}$ *satisfies the requirements.*

Even though the general lexicographic upgrade covers many natural upgrades [8], there are also 'reasonable' policies that fall outside its scope. Sometimes we are not interested in considering the complete order among the choices of the most reliable agent, but only her most preferred choices. To model such upgrades, as mentioned in [9] we provide the following preference upgrade definition.

Definition 4 (General layered upgrade)**.** *A* layered list $\mathcal{S}$ *over* W *is a finite (possibly empty) list of pairwise disjoint subsets of* W *together with a default preference ordering over* W*. The list's length is denoted by* $|\mathcal{S}|$*, its* k*th element is denoted by* $\mathcal{S}[k]$ *(with* $1 \leq k \leq |\mathcal{S}|$*), and* $\leq^{\mathcal{S}}_{\text{def}}$ *is its default preference ordering. Intuitively,* $\mathcal{S}$ *defines layers of elements of* W *in the new preference ordering* $\leq_{\mathcal{S}}$*, with* $\mathcal{S}[1]$ *the set of worlds that will be in the topmost layer and* $\leq^{\mathcal{S}}_{\text{def}}$ *the preference ordering that will be applied to each individual set and to those worlds not in* $\bigcup_{k=1}^{|\mathcal{S}|} \mathcal{S}[k]$*. Formally, given* $\mathcal{S}$*, the ordering* $\leq_{\mathcal{S}} \subseteq (W \times W)$ *is defined as*

$$u \leq_{\mathcal{S}} v \quad \textit{iff}_{def} \quad \underbrace{\Big(u \leq^{\mathcal{S}}_{\text{def}} v \;\wedge\; \big(\{u, v\} \cap \bigcup_{k=1}^{|\mathcal{S}|} \mathcal{S}[k] = \emptyset \vee \bigvee_{k=1}^{|\mathcal{S}|} \{u, v\} \subseteq \mathcal{S}[k]\big)\Big)}_{1} \;\vee\; \underbrace{\bigvee_{k=1}^{|\mathcal{S}|} \Big(v \in \mathcal{S}[k] \;\wedge\; u \notin \bigcup_{l=1}^{k} \mathcal{S}[l]\Big)}_{2}$$

Thus, $u \leq_{\mathcal{S}} v$ holds if this agrees with the default ordering $\leq^{\mathcal{S}}_{\text{def}}$ and either neither u nor v are in any of the specified sets in $\mathcal{S}$ or else both are in the same set (part 1), or if there is a set $\mathcal{S}[k]$ in which v appears and u appears neither in the same set (a case already covered in part 1) nor in one with higher priority (part 2).

Proposition 2. *Let $\mathcal{S}$ be a layered list over W. If $\leq^{\mathcal{S}}_{\text{def}}$ is reflexive (transitive, total, respectively), then so is $\leq_{\mathcal{S}}$.*

Definition 5. *Let $M = \langle W, \{\leq_i, \preccurlyeq_i\}_{i \in Ag}, V\rangle$ be a PR model.*

- *Let $\mathcal{S}$ be a layered list whose default ordering is reflexive, transitive and total; let $j \in Ag$ be an agent. The PR model $\text{gy}^j_{\mathcal{S}}(M) = \langle W, \{\leq'_i, \preccurlyeq_i\}_{i \in Ag}, V\rangle$ is such that, for every agent $i \in Ag$, $\leq'_i := \leq_{\mathcal{S}}$ if $i = j$, and $\leq'_i := \leq_i$ otherwise.*
- *Let $\boldsymbol{\mathcal{S}}$ be a list of $|Ag|$ layered lists whose default ordering are reflexive, transitive and total, with $\boldsymbol{\mathcal{S}}_i$ its ith element. The PR model $\text{gy}_{\boldsymbol{\mathcal{S}}}(M) = \langle W, \{\leq'_i, \preccurlyeq_i\}_{i \in Ag}, V\rangle$ is such that, for every agent $i \in Ag$, $\leq'_i := \leq_{\boldsymbol{\mathcal{S}}_i}$*

We have proposed different preference upgrade operators based on agent reliabilities. Now, the question is under what conditions these upgrade operators may lead to single-peakedness of agent preferences.

4 Deliberating towards single-peakedness

On the one hand we have the general lexicographic upgrade operation which considers a particular list to define the upgraded preferences. On the other hand we have this layered upgrade operation which is based on arbitrary subsets of choices and providing an order between them. There is a whole territory of possible upgrade operators in between these possibilities that is uncharted as of now. We would like to focus on charting the territory with a special emphasis on single-peakedness. We now assume the preference orderings to be asymmetric in addition to being total and transitive. Each agent is endowed with such a preference relation over the worlds.

Definition 6. *A preference profile is single-peaked if there exists a world w_i for each agent i and a linear order L such that $w_i L w' L w''$ or $w'' L w' L w_i$ imply $w' <_i w''$.*

Ballester and Haeringer [2] showed that the following two conditions characterize single-peakedness.

- For any subset of worlds the set of worlds considered as the worst by all agents cannot contains more than 2 elements (known as the worst-restricted condition in the literature).

- There cannot be four worlds w_1, w_2, w_3, w_4 and two agents i, j such that $w_1 <_i w_2 <_i w_3$, $w_3 <_j w_2 <_j w_1$, and $w_4 <_i w_2$, $w_4 <_j w_2$. In other words, two agents cannot disagree on the relative ranking of two alternatives with respect to a third alternative but agree on the (relative) ranking of a fourth one.

Our task is to investigate that under what conditions the given deliberation processes can achieve these properties. The first one should be easy to get: Since the orderings are asymmetric, the lexicographic upgrade policy will be identified with the drastic upgrade policy [8] which would lead to unanimity or oscillation. If unanimity is reached, we have single-peakedness trivially. In case of oscillation, we need to make sure that whichever be the agents included in oscillation for each agent, the least preferred world can only vary between (at most) two of the given worlds. For the layered upgrade ordering we will have a more interesting property of ensuring the weakest layer to contain the same two elements always. The second condition is more tricky, but once again can be broken down into several sub-conditions in the layered case. We leave the formal work for the main paper. We conclude with mentioning the known fact that getting single-peaked preferences via deliberation would pave the way of using aggregation rules which will lead to collective decision making avoiding the impossibility results of Arrow and others.

References

[1] Kenneth J. Arrow, A. K. Sen, and Kotaro Suzumura, editors. *Handbook of Social Choice and Welfare*. Elsevier, 2002. Two volumes.

[2] Miguel A. Ballester and Guillaume Haeringer. A characterisation of the single-peaked domain. *Social 'Choice and Welfare*, 36(2):305–322, 2011.

[3] J. P. Burgess. Basic tense logic. In Dov Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic*, volume II, chapter 2, pages 89–133. Reidel, 1984.

[4] Franz Dietrich and Christian List. Where do preferences come from? *International Journal of Game Theory*, 42(3):613–637, 2013.

[5] John S. Dryzek and Christian List. Social choice theory and deliberative democracy: A reconciliation. *British Journal of Political Science*, 33:1–28, 2003.

[6] Dov M. Gabbay, Odinaldo Rodrigues, and Gabriella Pigozzi. Connections between belief revision, belief merging and social choice. *Journal of Logic and Computation*, 19(3):445–446, 2009.

[7] George Gargov and Solomon Passy. A note on boolean modal logic. In Petio Petrov Petkov, editor, *Mathematical Logic*, pages 299–309. Springer US, 1990.

[8] Sujata Ghosh and Fernando R. Velázquez-Quesada. Agreeing to agree: Reaching unanimity via preference dynamics based on reliable agents. In Rafael Bordini, Edith Elkind, Gerhard Weiss, and Pinar Yolum, editors, *AAMAS 2015*, pages 1491–1499, 2015.

[9] Sujata Ghosh and Fernando R. Velázquez-Quesada. A note on reliability-based preference dynamics. In Wiebe van der Hoek, Wesley H. Holliday, and Wen fang Wang, editors, *LORI 2015*, 2015.

[10] Roosmarijn Goldbach. Modelling democratic deliberation. Master's Thesis, 2015.

[11] R. I. Goldblatt. Semantic analysis of orthologic. *Journal of Philosophical Logic*, 3(1):19–35, March 1974.

[12] Till Grüne-Yanoff and Sven Ove Hansson, editors. *Preference Change*, volume 42 of *Theory and Decision Library*. Springer, 2009.

[13] David Harel, Dexter Kozen, and Jerzy Tiuryn. *Dynamic Logic*. MIT Press, Cambridge, MA, 2000.

[14] Sébastien Konieczny and Ramón Pino Pérez. Logic based merging. *Journal of Philosophical Logic*, 40(2):239–270, 2011.

[15] Fernando R. Velázquez-Quesada. Reliability-based preference dynamics: lexicographic upgrade. *Journal of Logic and Computation*, 27(8):2341–2381, 2017.