Partial Preferences for Mediated Bargaining

Piero Bonatti, Marco Faella, Luigi Sauro

Università di Napoli Federico II, Italy

SR 2014, Grenoble, Apr 2014

Basic idea

- Classical Decision Theory assumes that it is always possible to compare two different choices:

$$f \prec g \quad f \succ g \quad f \sim g$$

  i.e. preferences are represented as weak orders (complete and transitive relations)

- here, we take into account a fourth option: incomparability

$$f \not\lessgtr g$$

  which means to relinquish completeness and represent preferences as pre-orders (reflexive and transitive relations)

A meta-question

- Besides some valuable contributions (e.g. Simon 1955, Aumann 1961, Aumann 1962, Dubra et al. 2004, Reffgen 2010), not so many efforts have been deserved to incompleteness, why?
- Moreover, Game and Social Choice Theory are entirely based on classic Decision Theory, does it really worth to shake this imposing cathedral from the foundations?
- Some generic arguments to *not* investigate an idea:
    - bad idea (e.g. to assume not transitive preferences)
    - no good motivations
    - technical difficulties

Are partial preferences a bad idea?

- von Neumann and Morgenstern: *It is conceivable and may even in a way be more realistic to allow for cases where the individual is neither able to state which of two alternatives he prefers nor that they are equally desirable. How real this possibility is, both for individuals and for organizations, seems to be an extremely interesting question, but it is a question of fact. It certainly deserves further study.* [Theory of Games and Economic Behavior]
- Savage: *There is some temptation to explore the possibilities of analyzing preferences among acts as a partial ordering, that is, in effect to replace part 1 of the definition of simple ordering [i.e. completeness] by the very weak proposition $f \preceq f$, admitting that some pairs of act are incomparable. This would seem to give expression to introspective sensations of indecision or vacillation, which we may be reluctant to identify with indifference. My own conjecture is that it would prove a blind alley losing much in power and advancing little, if at all, in realism; but only an enthusiastic exploration could shed real light on the question.* [The Foundations of Statistics]

Motivations

- Motivations sustaining partial preferences that have been proposed so far:
  - unmeasurability [Aumann]
  - lack of knowledge/bounded rationality [Simon]
  - trembling preferences [Savage]
  - multi person utilities [Reffgen]
- Our claim:
  - good ingredients, but not sufficient to motivate incomparability!
  - stronger motivations come from Computer Science, not Economics!

Nennella's example

- Nennella is a typical trattoria in Naples
- Staff does not speak English (not even Italian) and the menu is entirely in dialect (No Google translate)
- 4 dishes: scammaro, pignatiello, purpetiell'a'Luciana, gattò

Nennella’s example

- Alice considers scammaro better than pignatiello whereas, according to Bob, gattò is better than purpetiell’a’Luciana
- You generally agree with them.
- What you get is not a total order!

- Due to the lack of information, any linearization seems to be discretional. If someone ask to you: *what do you like the most, scammaro or gattò?* probably your answer would be: *I have no idea! When I will have tried both I will know*.

Nennella's example

- However, a (classical) decision theorist would rise an objection: Decision Theory is not psychology. We do not look inside your mind and see what you like the most. We only aim at justifying how an *agent* behaves *at the moment* she is making a decision among a collection of choices (and not other option is possible).
- Going back to the example, pignatiello and purpetiell'a'Luciana are dominated and hence discarded. Then, since you have no information about scammaro and gattò, at the moment you are making a decision you will probably flip a coin. In any case, if someone ask you: "*Do you have any reason to systematically choose one of them?* ", your answer will be "*No*". This is exactly what we call *equivalence* or *indifference*.
- In other words, completeness means that a decision has to be made [Gilboa]

Motivations (Revisited)

- Missing ingredient: what if the agent who makes the decision operates in behalf of another agent?
  - Economics: a businessman makes decisions by his own
  - Computer Science: we want to design software agents that operates in behalf of real users
- Problems with delegation:
  - preference injection can be time consuming
  - arbitrary disambiguation of trembling preferences
  - all possible information has to be acquired in advance
- Goal: to discard the majority of the dominated choices, not to return a single one
- Solution: the user instructs the software agent with a partial preference which approximates his actual preference
- Behavioral definition of incomparability [Aumann, Gilboa]:
  - indifference: flip a coin
  - incomparability: refuse to decide and report back to the user

Incomplete preferences in Computer Science

Applicative scenario:

- User: a software house willing to deploy a new application on the cloud
- Choices: different offers from cloud-service providers

Features:

- dozens (or even hundreds) of different offers [software agent]
- different relevant aspects (security, cost, reliability, scalability etc.) [unmeasurability]
- the software house consists of several people [multiple utilities]
- new technologies [lack of knowledge]
- similar offers [trembling preferences]

Classical Decision Theory

Preliminaries:

- $\mathcal{A}$: a set of alternatives
- f: a probability distribution over $\mathcal{A}$ with finite support $supp(f)$
  - where $supp(f) = \{a \in \mathcal{A} \mid f(a) > 0\}$

Classical axioms:

- $f \preceq g$ or $g \preceq f$ (completeness)
- $f \preceq g$ and $g \preceq h \implies f \preceq h$ (transitivity)
- $f \preceq g \implies \alpha f + (1-\alpha)h \preceq \alpha g + (1-\alpha)h, \forall \alpha \in [0,1]$ (independence)
- $f_n \to f$, $g_n \to g$ and $\forall n\, f_n \preceq g_n \implies f \preceq g$ (continuity)

Weak Decision Theory with Incompleteness (WDTI) [Dubra et al., 2004]

Replace completeness with reflexivity:

- $f \preceq f$ (reflexivity)
- $f \preceq g$ and $g \preceq h \implies f \preceq h$ (transitivity)
- $f \preceq g \implies \alpha f + (1-\alpha)h \preceq \alpha g + (1-\alpha)h, \forall \alpha \in [0,1]$ (independence)
- $f_n \preceq g_n, \forall n \implies f \preceq g$ where $f_n \to f$ and $g_n \to g$ (continuity)

Myerson's axioms:

- $f \prec g$ and $0 \leq \alpha < \beta \leq 1 \implies \alpha g + (1-\alpha)f \prec \beta g + (1-\beta)f$
- $f_1 \preceq g_1,\ f_2 \preceq g_2 \implies \alpha f_1 + (1-\alpha)f_2 \preceq \alpha g_1 + (1-\alpha)g_2$
- $f_1 \prec g_1,\ f_2 \preceq g_2 \implies \alpha f_1 + (1-\alpha)f_2 \prec \alpha g_1 + (1-\alpha)g_2$
- $f \preceq g \preceq h \implies \exists \alpha \quad \alpha h + (1-\alpha)f \sim g$ (continuity)

Properties of WDTI

Theorem

WDTI implies

- $f \prec g$ *and* $0 \leq \alpha < \beta \leq 1 \implies \alpha g + (1-\alpha) f \prec \beta g + (1-\beta) f$
- $f_1 \preceq g_1,\ f_2 \preceq g_2 \implies \alpha f_1 + (1-\alpha) f_2 \preceq \alpha g_1 + (1-\alpha) g_2$
- $f_1 \prec g_1,\ f_2 \preceq g_2 \implies \alpha f_1 + (1-\alpha) f_2 \preceq \alpha g_1 + (1-\alpha) g_2$

Definition

Given two distributions f, g, we write $f \to g$ in case there exist $\epsilon > 0$ and two alternatives a_1 and a_2 such that *(i)* $[a_1] \prec [a_2]$, *(ii)* $g(a_1) = f(a_1) - \epsilon$, $g(a_2) = f(a_2) + \epsilon$, and *(iii)* for all $a \neq a_1, a_2$, $g(a) = f(a)$.

Theorem

$f \xrightarrow{*} g$ *implies* $f \prec g$.

is WDTI too weak?

- A theory can be seen as rules that allow to foresee how agents behave and consequently how mechanisms should be designed
- A typical use could be: given a partial preference over degenerate choices, which preferences over lotteries can I assume?

$$\frac{1}{2}[\text{scammaro}] + \frac{1}{2}[\text{gattò}] \quad \bowtie \quad [\text{scammaro}]$$

- In WDTI $\bowtie = \sim, \prec, \succ, \not\lessgtr$

$$\frac{1}{2}[\text{scammaro}] + \frac{1}{2}[\text{gattò}] \quad \prec \quad [\text{scammaro}]$$

Justification Principle

- Let $f = \alpha f_1 + (1-\alpha) f_2$ and $g = \beta g_1 + (1-\beta) g_2$
- In classic Decision Theory it holds that

$$f \prec g \implies \exists j, k \in \{1, 2\}\ f_j \prec g_k$$

which roughly means that the decision maker is able to justify the preference over combined lotteries in terms of their components

Intuition: Comparing $f = \alpha f_1 + (1-\alpha) f_2$ and $g = \alpha g_1 + (1-\alpha) g_2$ encompasses comparing four possible draws: (f_1, g_1), (f_1, g_2), (f_2, g_1), and (f_2, g_2). If in each draw the delegated agent

- favors f
- chooses randomly
- reports back

then it cannot systematically choose g.

Decision Theory with Incompleteness

- $f \preceq f$
- $f \preceq g$ and $g \preceq h \implies f \preceq h$
- $f \prec g$ and $0 \leq \alpha < \beta \leq 1 \implies \alpha g + (1-\alpha) f \prec \beta g + (1-\beta) f$
- $f_1 \preceq g_1,\ f_2 \preceq g_2 \implies \alpha f_1 + (1-\alpha) f_2 \preceq \alpha g_1 + (1-\alpha) g_2$
- $f_1 \prec g_1,\ f_2 \preceq g_2 \implies \alpha f_1 + (1-\alpha) f_2 \preceq \alpha g_1 + (1-\alpha) g_2$
- $\alpha f_1 + (1-\alpha) f_2 \prec \beta g_1 + (1-\beta) g_2 \implies \exists j, k \in \{1, 2\}\ f_j \prec g_k$

Properties of Weak Partial Decision Theory

Theorem

Let f, g be such that, for all $a \in supp(f)$ and $a' \in supp(g)$, either $a \sim a'$ or $a \lessgtr\!\!\!\!/\; a'$. Then, $f \lessgtr\!\!\!\!/\; g$ or $f \sim g$.

Notice that

- g is maximally preferred if $\forall f \; g \not\prec f$ (i.e. $f \preceq g$ or $f \lessgtr\!\!\!\!/\; g$)
- without the previous theorem maximally preferred lotteries can be hardly inferred $\Longrightarrow$ no Best Responses $\Longrightarrow$ no Nash Equilibrium $\Longrightarrow$ no Dominant Equilibrium $\Longrightarrow$ no Faithfulness $\Longrightarrow$ etc.

Future Works

- Different notions of continuity:
  - C1: $f_n \preceq g_n, \forall n \implies f \preceq g$ where $f_n \to f$ and $g_n \to g$
  - C2: $f \preceq g \preceq h \implies \exists\alpha \quad \alpha h + (1-\alpha)f \sim g$
  - C3: $\alpha f + (1-\alpha)g \prec h, : \forall\alpha > 0 \implies h \not\prec g$

  C3 is implied by C1. Whereas, in order to obtain C2, WDTI seems to require at least

  $$\alpha f_1 + (1-\alpha)f_2 \not\lesssim \beta g_1 + (1-\beta)g_2 \implies \exists j,k \in \{1,2\}\ f_j \not\lesssim g_k$$

  does it make sense?

Conclusions

- Partial preferences have been considered an interesting issue from the very beginning, even if
  - the main objection has been that *at the moment* an agent is making a decision it *collapses* into a classical decision maker
  - simply replacing completeness with reflexivity results in a weak theory
- Finding new motivations: if the agent operates in behalf of another one, different rules of engagement may determinate whether it can flip a coin or it is not allowed to make a decision by its own
- Adjusting the theory: maintaining preferences over lotteries justifiable
- Not a final version, but it seems powerful enough to be used in mechanism design