Dynamics for Aggregating Cardinal Preferences[1]

Matthias Greger

Technical University of Munich

(with Florian Brandl, Felix Brandt, Dominik Peters, Erel Segal-Halevi, Christian Stricker, and Warut Suksompong)

The Second Vienna-Graz Workshop on (Computational) Social Choice
February 9, 2024

[1] Based on *Funding Public Projects: A Case for the Nash Product Rule* (Brandl et al., WINE/JME 2021) and *Balanced Donor Coordination* (Brandt et al., EC 2023).

General Model

- Set of alternatives $A = \{a, b, c, \dots\}$ of size m.
  - Consider lotteries $\delta \in \Delta(1)$ over the alternatives as outcomes.
- Set of agents $N = \{1, \dots, n\}$ with utility functions $u_i : \Delta(1) \to \mathbb{R}$.
- Distribution rule f mapping each *profile* $(u_i)_{i \in N}$ to a lottery δ.

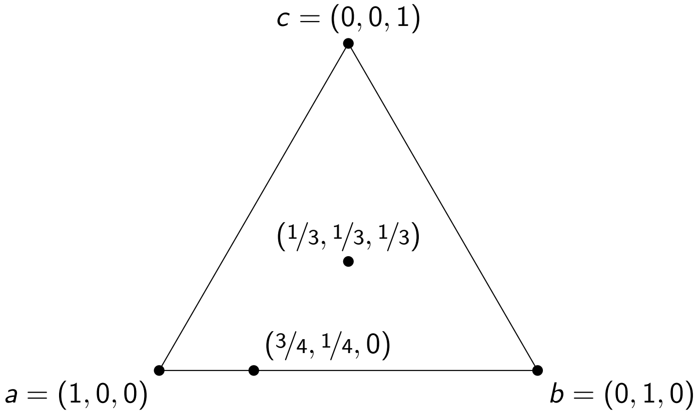

Applications

Participatory Budgeting (Aziz and Shah, 2021; Rey and Maly, 2023)

- Residents decide how to distribute a budget (provided by *the city*) on a set of public projects.
- In this talk, projects do not have fixed costs but rather profit from any amount of money they receive; sometimes called *portioning*.

Donor Coordination (Brandl et al., 2022; Brandt et al., 2023)

- Donors decide how to distribute a budget (provided by *themselves*) on a set of public projects.
- Participation incentives become even more important.

Utility Functions - Substitutes versus Complements

Denote by $v_{i,x} \geq 0$ agent i's valuation for project x. If $v_{i,x} \in \{0, 1\}$ for all agents and projects, define agent i's set of approved projects as $A_i := \{x \in A : v_{i,x} = 1\}$.

Perfect substitutes: , ,
(Bogomolnaia et al., 2005)

- "I like reading books and playing football . Supporting any of the two is fine."
- Dichotomous utilities $u_i(\delta) = \sum_{x \in A_i} \delta(x)$.
- More general: Linear utilities $u_i(\delta) = \sum_{x \in A} v_{i,x} \cdot \delta(x)$.

Perfect complements: , ,
(Brandt et al., 2023)

- "I approve charities and . Both should receive some money."
- Binary Leontief utilities $u_i(\delta) = \min_{x \in A_i} \delta(x)$.
- More general: Leontief utilities $u_i(\delta) = \min_{x \in A : v_{i,x} > 0} \delta(x)/v_{i,x}$.

In the following, restriction to approvals but results carry over to the more general utility functions.

All Power to the Agents

Standard approach:

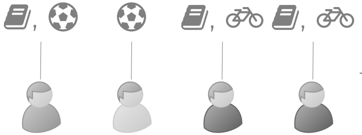

$\xrightarrow{P}$ distribution rule f $\xrightarrow{f(P)}$ returned distribution δ

Dynamical approach:
Each agent i receives $1/n$ of the total budget or "decision power" and distributes it via $\delta_i \in \Delta(1/n)$. Then, $\delta = \sum_{i \in N} \delta_i$.
Observing the (overall) distribution δ, agents are allowed to update δ_i.

$$\delta^0 \rightarrow \delta^1 \rightarrow \cdots \rightarrow \delta^t \rightarrow \delta^{t+1} \rightarrow \ldots \xrightarrow{?} \delta$$

$\delta_1^0, \delta_1^1, \ldots$ $\delta_2^0, \delta_2^1, \ldots$ $\delta_3^0, \delta_3^1, \ldots$ $\delta_4^0, \delta_4^1, \ldots$

Related ideas can be found in various areas, e.g., *fair division* (Zhang, 2011).

Perfect Substitutes: $u_i(\delta) = \sum_{x \in A_i} \delta(x)$

Consider the following dynamics with the sequence of distributions $(\delta^t)_{t=0,1,2,\ldots}$:

				$u_i(\delta^0)$
Agent 1	1/8	1/8	·	6/8
Agent 2	·	2/8	·	3/8
Agent 3	1/8	·	1/8	5/8
Agent 4	1/8	·	1/8	5/8
δ^0	3/8	3/8	2/8	

$t = 0$: Each agent i distributes uniformly over A_i.

$t \geq 1$: Each agent i updates her individual distribution δ_i via $n \cdot \delta_i^t(x) = \delta^{t-1}(x)/u_i(\delta^{t-1})$ for every $x \in A$ ("fractional gain").

Perfect Substitutes: $u_i(\delta) = \sum_{x \in A_i} \delta(x)$

Consider the following dynamics with the sequence of distributions $(\delta^t)_{t=0,1,2,\ldots}$:

				$u_i(\delta^1)$
Agent 1	$1/8$	$1/8$	$\cdot$	$4/5$
Agent 2	$\cdot$	$2/8$	$\cdot$	$3/8$
Agent 3	$3/20$	$\cdot$	$2/20$	$5/8$
Agent 4	$3/20$	$\cdot$	$2/20$	$5/8$
δ^1	$17/40$	$15/40$	$8/40$	

$t = 0$: Each agent i distributes uniformly over A_i.

$t \geq 1$: Each agent i updates her individual distribution δ_i via $n \cdot \delta_i^t(x) = \delta^{t-1}(x)/u_i(\delta^{t-1})$ for every $x \in A$ ("fractional gain").

Perfect Substitutes: $u_i(\delta) = \sum_{x \in A_i} \delta(x)$

Consider the following dynamics with the sequence of distributions $(\delta^t)_{t=0,1,2,\dots}$:

				$u_i(\delta^2)$
Agent 1	$17/128$	$15/128$	$\cdot$	$84/100$
Agent 2	$\cdot$	$2/8$	$\cdot$	$49/128$
Agent 3	$17/100$	$\cdot$	$8/100$	$79/128$
Agent 4	$17/100$	$\cdot$	$8/100$	$79/128$
δ^2	≈ 0.4728	$47/128$	$16/100$	

$t = 0$: Each agent i distributes uniformly over A_i.
$t \geq 1$: Each agent i updates her individual distribution δ_i via $n \cdot \delta_i^t(x) = \delta^{t-1}(x)/u_i(\delta^{t-1})$ for every $x \in A$ ("fractional gain").

Perfect Substitutes: $u_i(\delta) = \sum_{x \in A_i} \delta(x)$

Consider the following dynamics with the sequence of distributions $(\delta^t)_{t=0,1,2,\dots}$:

				$u_i(\delta)$
Agent 1	$2/12$	$1/12$	$\cdot$	1
Agent 2	$\cdot$	$1/4$	$\cdot$	$1/3$
Agent 3	$1/4$	$\cdot$	$\cdot$	$2/3$
Agent 4	$1/4$	$\cdot$	$\cdot$	$2/3$
δ	$2/3$	$1/3$	0	

$t = 0$: Each agent i distributes uniformly over A_i.

$t \geq 1$: Each agent i updates her individual distribution δ_i via $n \cdot \delta_i^t(x) = \delta^{t-1}(x)/u_i(\delta^{t-1})$ for every $x \in A$ ("fractional gain").

$\rightarrow$ Convergence to δ.

Perfect Substitutes: $u_i(\delta) = \sum_{x \in A_i} \delta(x)$

Dynamics:

$t = 0$: Each agent i distributes uniformly on A_i.

$t \geq 1$: Each agent i updates her individual distribution δ_i via $n \cdot \delta_i^t(x) = \delta^{t-1}(x)/u_i(\delta^{t-1})$ for every $x \in A$ ("fractional gain").

Definition

The *Nash welfare* of a distribution δ is defined as $Nash(\delta) = \prod_{i \in N} u_i(\delta)$.

Theorem (Cover, 1984; Brandl et al., 2022)

For any profile, $(Nash(\delta^t))_{t \in \mathbb{N}}$ converges to the optimum Nash product. If Nash welfare is maximized by a unique distribution, the dynamics converges to it.

Open question: In case of multiple Nash maximizers, does the dynamics converge to a specific distribution?

Perfect Complements: $u_i(\delta) = \min_{x \in A_i} \delta(x)$

Consider the following dynamics with the sequence of distributions $(\delta^t)_{t=0,1,2,\dots}$:

				$u_i(\delta^0)$
Agent 1	$\cdot$	$\cdot$	$\cdot$	0
Agent 2	$\cdot$	$\cdot$	$\cdot$	0
Agent 3	$\cdot$	$\cdot$	$\cdot$	0
Agent 4	$\cdot$	$\cdot$	$\cdot$	0
δ^0	0	0	0	

Consider an infinite, arbitrary sequence $\mathcal{S} = (i_t)_{t \in \mathbb{N}}$ of agents from N.
At each time step t, agent i_t (re-)distributes δ_{i_t} via
$\delta_{i_t}^{t+1} = \arg\max_{\delta_{i_t}} u_{i_t}(\delta^t - \delta_{i_t}^t + \delta_{i_t})$ ("best response") and
$\delta^{t+1} = \delta^t - \delta_{i_t}^t + \delta_{i_t}^{t+1}$.
Example: Let $\mathcal{S} = (1, 2, 3, 4, 1, 2, 3, 4, \dots)$ be a round-robin sequence.

Perfect Complements: $u_i(\delta) = \min_{x \in A_i} \delta(x)$

Consider the following dynamics with the sequence of distributions $(\delta^t)_{t=0,1,2,\ldots}$:

				$u_i(\delta^1)$
Agent 1	1/8	1/8	·	1/8
Agent 2	·	·	·	1/8
Agent 3	·	·	·	0
Agent 4	·	·	·	0
δ^1	1/8	1/8	0	

Consider an infinite, arbitrary sequence $\mathcal{S} = (i_t)_{t \in \mathbb{N}}$ of agents from N.
At each time step t, agent i_t (re-)distributes δ_{i_t} via
$\delta_{i_t}^{t+1} = \arg\max_{\delta_{i_t}} u_{i_t}(\delta^t - \delta_{i_t}^t + \delta_{i_t})$ ("best response") and
$\delta^{t+1} = \delta^t - \delta_{i_t}^t + \delta_{i_t}^{t+1}$.
Example: Let $\mathcal{S} = (\mathbf{1}, 2, 3, 4, 1, 2, 3, 4, \ldots)$ be a round-robin sequence.

Perfect Complements: $u_i(\delta) = \min_{x \in A_i} \delta(x)$

Consider the following dynamics with the sequence of distributions $(\delta^t)_{t=0,1,2,\ldots}$:

				$u_i(\delta^2)$
Agent 1	1/8	1/8	·	1/8
Agent 2	·	1/4	·	3/8
Agent 3	·	·	·	0
Agent 4	·	·	·	0
δ^2	1/8	3/8	0	

Consider an infinite, arbitrary sequence $\mathcal{S} = (i_t)_{t \in \mathbb{N}}$ of agents from N.
At each time step t, agent i_t (re-)distributes δ_{i_t} via
$\delta_{i_t}^{t+1} = \arg\max_{\delta_{i_t}} u_{i_t}(\delta^t - \delta_{i_t}^t + \delta_{i_t})$ ("best response") and
$\delta^{t+1} = \delta^t - \delta_{i_t}^t + \delta_{i_t}^{t+1}$.
Example: Let $\mathcal{S} = (1, \mathbf{2}, 3, 4, 1, 2, 3, 4, \ldots)$ be a round-robin sequence.

Perfect Complements: $u_i(\delta) = \min_{x \in A_i} \delta(x)$

Consider the following dynamics with the sequence of distributions $(\delta^t)_{t=0,1,2,\ldots}$:

				$u_i(\delta^3)$
Agent 1	1/8	1/8	·	3/16
Agent 2	·	1/4	·	6/16
Agent 3	1/16	·	3/16	3/16
Agent 4	·	·	·	3/16
δ^3	3/16	6/16	3/16	

Consider an infinite, arbitrary sequence $\mathcal{S} = (i_t)_{t \in \mathbb{N}}$ of agents from N.
At each time step t, agent i_t (re-)distributes δ_{i_t} via
$\delta_{i_t}^{t+1} = \arg\max_{\delta_{i_t}} u_{i_t}(\delta^t - \delta_{i_t}^t + \delta_{i_t})$ ("best response") and
$\delta^{t+1} = \delta^t - \delta_{i_t}^t + \delta_{i_t}^{t+1}$.

Example: Let $\mathcal{S} = (1, 2, \mathbf{3}, 4, 1, 2, 3, 4, \ldots)$ be a round-robin sequence.

Perfect Complements: $u_i(\delta) = \min_{x \in A_i} \delta(x)$

Consider the following dynamics with the sequence of distributions $(\delta^t)_{t=0,1,2,\dots}$:

				$u_i(\delta^4)$
Agent 1	1/8	1/8	·	5/16
Agent 2	·	1/4	·	6/16
Agent 3	1/16	·	3/16	5/16
Agent 4	1/8	·	1/8	5/16
δ^4	5/16	6/16	5/16	

Consider an infinite, arbitrary sequence $\mathcal{S} = (i_t)_{t \in \mathbb{N}}$ of agents from N.
At each time step t, agent i_t (re-)distributes δ_{i_t} via
$\delta_{i_t}^{t+1} = \arg\max_{\delta_{i_t}} u_{i_t}(\delta^t - \delta_{i_t}^t + \delta_{i_t})$ ("best response") and
$\delta^{t+1} = \delta^t - \delta_{i_t}^t + \delta_{i_t}^{t+1}$.
Example: Let $\mathcal{S} = (1, 2, 3, \mathbf{4}, 1, 2, 3, 4, \dots)$ be a round-robin sequence.

Perfect Complements: $u_i(\delta) = \min_{x \in A_i} \delta(x)$

Consider the following dynamics with the sequence of distributions $(\delta^t)_{t=0,1,2,\ldots}$:

				$u_i(\delta^5)$
Agent 1	5/32	3/32	·	11/32
Agent 2	·	1/4	·	11/32
Agent 3	1/16	·	3/16	10/32
Agent 4	1/8	·	1/8	10/32
δ^5	11/32	11/32	10/32	

Consider an infinite, arbitrary sequence $\mathcal{S} = (i_t)_{t\in\mathbb{N}}$ of agents from N.
At each time step t, agent i_t (re-)distributes δ_{i_t} via
$\delta_{i_t}^{t+1} = \arg\max_{\delta_{i_t}} u_{i_t}(\delta^t - \delta_{i_t}^t + \delta_{i_t})$ ("best response") and
$\delta^{t+1} = \delta^t - \delta_{i_t}^t + \delta_{i_t}^{t+1}$.
Example: Let $\mathcal{S} = (1, 2, 3, 4, \mathbf{1}, 2, 3, 4, \ldots)$ be a round-robin sequence.

Perfect Complements: $u_i(\delta) = \min_{x \in A_i} \delta(x)$

Consider the following dynamics with the sequence of distributions $(\delta^t)_{t=0,1,2,\dots}$:

				$u_i(\delta^6)$
Agent 1	5/32	3/32	·	11/32
Agent 2	·	1/4	·	11/32
Agent 3	1/16	·	3/16	10/32
Agent 4	1/8	·	1/8	10/32
δ^6	11/32	11/32	10/32	

Consider an infinite, arbitrary sequence $\mathcal{S} = (i_t)_{t \in \mathbb{N}}$ of agents from N.
At each time step t, agent i_t (re-)distributes δ_{i_t} via
$\delta_{i_t}^{t+1} = \arg\max_{\delta_{i_t}} u_{i_t}(\delta^t - \delta_{i_t}^t + \delta_{i_t})$ ("best response") and
$\delta^{t+1} = \delta^t - \delta_{i_t}^t + \delta_{i_t}^{t+1}$.
Example: Let $\mathcal{S} = (1, 2, 3, 4, 1, \mathbf{2}, 3, 4, \dots)$ be a round-robin sequence.

Perfect Complements: $u_i(\delta) = \min_{x \in A_i} \delta(x)$

Consider the following dynamics with the sequence of distributions $(\delta^t)_{t=0,1,2,\ldots}$:

				$u_i(\delta^7)$
Agent 1	5/32	3/32	·	21/64
Agent 2	·	1/4	·	22/64
Agent 3	3/64	·	13/64	21/64
Agent 4	1/8	·	1/8	21/64
δ^7	21/64	22/64	21/64	

Consider an infinite, arbitrary sequence $\mathcal{S} = (i_t)_{t\in\mathbb{N}}$ of agents from N.
At each time step t, agent i_t (re-)distributes δ_{i_t} via
$\delta_{i_t}^{t+1} = \arg\max_{\delta_{i_t}} u_{i_t}(\delta^t - \delta_{i_t}^t + \delta_{i_t})$ ("best response") and
$\delta^{t+1} = \delta^t - \delta_{i_t}^t + \delta_{i_t}^{t+1}$.
Example: Let $\mathcal{S} = (1, 2, 3, 4, 1, 2, \mathbf{3}, 4, \ldots)$ be a round-robin sequence.

Consider the following dynamics with the sequence of distributions $(\delta^t)_{t=0,1,2,\dots}$:

				$u_i(\delta)$
Agent 1	2/12	1/12	·	1/3
Agent 2	·	1/4	·	1/3
Agent 3	1/24	·	5/24	1/3
Agent 4	1/8	·	1/8	1/3
δ	1/3	1/3	1/3	

Consider an infinite, arbitrary sequence $\mathcal{S} = (i_t)_{t \in \mathbb{N}}$ of agents from N.
At each time step t, agent i_t (re-)distributes δ_{i_t} via
$\delta_{i_t}^{t+1} = \arg\max_{\delta_{i_t}} u_{i_t}(\delta^t - \delta_{i_t}^t + \delta_{i_t})$ ("best response") and
$\delta^{t+1} = \delta^t - \delta_{i_t}^t + \delta_{i_t}^{t+1}$.
Example: Let $\mathcal{S} = (1, 2, 3, 4, 1, 2, 3, 4, \dots)$ be a round-robin sequence.
$\rightarrow$ Convergence to δ.

Perfect Complements: $u_i(\delta) = \min_{x \in A_i} \delta(x)$

Dynamics:

Consider an infinite, arbitrary sequence $\mathcal{S} = (i_t)_{t \in \mathbb{N}}$ of agents from N.
At each time step t, agent i_t (re-)distributes δ_{i_t} via
$\delta_{i_t}^{t+1} = \arg\max_{\delta_{i_t}} u_{i_t}(\delta^t - \delta_{i_t}^t + \delta_{i_t})$ ("best response") and
$\delta^{t+1} = \delta^t - \delta_{i_t}^t + \delta_{i_t}^{t+1}$.

Theorem (Brandt et al., 2023)

Nash welfare is maximized by a unique distribution.

Perfect Complements: $u_i(\delta) = \min_{x \in A_i} \delta(x)$

Dynamics:
Consider an infinite, arbitrary sequence $\mathcal{S} = (i_t)_{t\in\mathbb{N}}$ of agents from N.
At each time step t, agent i_t (re-)distributes δ_{i_t} via
$\delta_{i_t}^{t+1} = \arg\max_{\delta_{i_t}} u_{i_t}(\delta^t - \delta_{i_t}^t + \delta_{i_t})$ ("best response") and
$\delta^{t+1} = \delta^t - \delta_{i_t}^t + \delta_{i_t}^{t+1}$.

Theorem (Brandt et al., 2023)

For any profile and any sequence $\mathcal{S}$ where each agents appears infinitely often, the dynamics converges to the Nash welfare maximizer.

Remark

For Leontief utilities beyond the binary case, we need to make an additional technical assumption on $\mathcal{S}$.

Open question: Is the additional assumption required?

Proof Ideas for Convergence of Dynamics

1. Find fixed points of $(\delta^t)_{t \geq 0}$.
2. Find potential function F with $F(\delta^{t+1}) > F(\delta^t)$ that is bounded on $\Delta(1)$.
   $\Rightarrow (F(\delta^t))_{t \geq 0}$ converges.
3. Characterize limit distribution(s).

From Dynamics to Equilibrium

Dynamics:
Consider an infinite, arbitrary sequence $\mathcal{S} = (i_t)_{t\in\mathbb{N}}$ of agents from N.
At each time step t, agent i_t (re-)distributes δ_{i_t} via
$\delta_{i_t}^{t+1} = \arg\max_{\delta_{i_t}} u_{i_t}(\delta^t - \delta_{i_t}^t + \delta_{i_t})$ ("best response") and
$\delta^{t+1} = \delta^t - \delta_{i_t}^t + \delta_{i_t}^{t+1}$.

Definition

A distribution $\delta \in \Delta(1)$ is in equilibrium iff it admits a decomposition $(\delta_i)_{i\in N}$ such that $u_i(\delta) \geq u_i(\delta - \delta_i + \delta_i')$ for all $i \in N$ and $\delta_i' \in \Delta(1/n)$.

Theorem (Brandt et al., 2023)

For Leontief utilities, the unique equilibrium distribution coincides with the Nash welfare maximizer.

Beyond Leontief utilities

Theorem (Brandt et al., 2023)

For any profile and any sequence $\mathcal{S}$ where each agents appears infinitely often, the dynamics converges to the equilibrium distribution.

The theorem also holds for utility functions other than binary Leontief utilities:

- Separably, additive utility functions $u_i(x) = \sum_{x \in A_i} g_i(\delta(x))$ where $g_i : \Delta(1) \to \mathbb{R}$ is a strictly concave function.
- Linear utility functions $u_i(x) = \sum_{x \in A} v_{i,x} \delta(x)$ (no equilibrium uniqueness).

Discussion

Advantages of dynamical approach:

- Agents' preferences remain private information.
- Agents are able to change their preferences over time.
- Justification for mechanisms arising from such dynamics.

Disadvantages of dynamical approach:

- Requires complete preferences over $\Delta(1)$.
- Limited applicability in certain areas (e.g., reduces to random dictatorship in voting with strict preferences).
- In general, convergence is not guaranteed.

References

H. Aziz and N. Shah. Participatory budgeting: Models and approaches. In T. Rudas and G. Péli, editors, *Pathways Between Social Science and Computational Social Science: Theories, Methods, and Interpretations*, pages 215–236. Springer International Publishing, 2021.

A. Bogomolnaia, H. Moulin, and R. Stong. Collective choice under dichotomous preferences. *Journal of Economic Theory*, 122(2):165–184, 2005.

F. Brandl, F. Brandt, M. Greger, D. Peters, C. Stricker, and W. Suksompong. Funding public projects: A case for the Nash product rule. *Journal of Mathematical Economics*, 99:102585, 2022.

F. Brandt, M. Greger, E. Segal-Halevi, and W. Suksompong. Balanced donor coordination. In *Proceedings of the 24th ACM Conference on Economics and Computation (ACM-EC)*, page 299, 2023.

T. M. Cover. An algorithm for maximizing expected log investment return. *IEEE Transactions on Information Theory*, 30(2):369–373, 1984.

S. Rey and J. Maly. The (computational) social choice take on indivisible participatory budgeting. Technical report, https://arxiv.org/pdf/2303.00621.pdf, 2023.

L. Zhang. Proportional response dynamics in the Fisher market. *Theoretical Computer Science*, 412:2691–2698, 2011.