

Justifications of Welfare Guarantees under Normalized Utilities

HARIS AZIZ

It is standard in computational social choice to analyse welfare considerations under the assumptions of normalized utilities. In this note, we summarize some common reasons for this approach. We then mention another justification which is ignored but has solid normative appeal. The central concept used in the ‘new’ justification can also be used more widely as a social objective.

Categories and Subject Descriptors: I.2.11 [**Distributed Artificial Intelligence**]: Multiagent Systems; J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms: Theory, Algorithms, Economics

Additional Key Words and Phrases: Fair allocation, solutions Concepts, multiagent resource allocation, mechanism design without money, computational social choice, approximation algorithms

1. INTRODUCTION

Social welfare under normalized utilities is frequently considered in the computational social choice literature. It is especially the case in multi-agent resource allocation, mechanism design without money and utilitarian voting. The welfare notions considered include utilitarian social welfare (sum of agents’ utilities), egalitarian social welfare (minimum of agents’ utilities) and Nash social welfare (product of agents’ utilities). Most of the focus is on utilitarian welfare which goes back to ideas of Bentham [1789].

For multi-agent resource allocation under ordinal preferences, researchers have considered how well social welfare is approximated by specific mechanisms for truthful preferences or preferences in equilibrium (see e.g., [Filos-Ratsikas et al., 2014, Guo and Conitzer, 2010, Bertsimas et al., 2011]). The standard assumption in these papers is that utilities are normalized i.e., an agent’s sum of utilities for all items is one. The assumption has been termed as *unit-sum*.

Similarly, there is growing literature on *implicit utilitarian voting* (see e.g., [Caragiannis et al., 2016b, Lu and Boutilier, 2011, Boutilier et al., 2012]) where normalized utilities are popular. For these settings, the *distortion* of a voting rule is used as a measure of how good the rule is. Distortion is the worst ratio of the maximum utilitarian social welfare versus the utilitarian social welfare achieved among all problem instances. Many results concerning distortion bounds hinge on the assumption the utilities are normalized.¹

In other multi-agent resource allocation problems, agents are asked to express cardinal utilities over items but these utilities are assumed to be normalized or are processed to be normalized. For example, Bouveret and Lemaître [2015] state that

¹There is another stream of results on distortion that assume that utilities are induced by a metric (see e.g., [Anshelevich and Postl, 2016].)

most collective utility functions only make sense if the utilities are expressed on a common scale or normalized. Even when agent preferences are ordinal, the scoring rules used to get ‘proxy utilities’ satisfy the normalized or equal scale property (see e.g., [Baumeister et al., 2014]).

In short, both within voting and resource allocation, there is a focus on welfare under normalized utilities.² When the ‘real’ utilities of agents are not normalized, it begs the question that why normalized utilities are used. The papers making this assumption either justify it as a standard assumption used in previous work or by mentioning one or two reasons. In this note, we curate and discuss justifications for welfare approaches under normalized utilities.³

First Principle/Philosophical Justifications

- (1) **Scale invariance.** Egalitarian and utilitarian social welfare suffer from being responsive to scale variance.⁴ Imposing normalized utilities is a simple way to regain scale invariance.
- (2) **Guaranteeing a proportion of the ultimate happiness of an individual.** Another way to view scale invariance is that one is concerned less about the amount of utility achieved by an individual agent and more so about the fraction of maximum possible utility that she achieves. This concern is indeed captured when one tries to approximate egalitarian welfare under normalized utilities.
- (3) **Fairness.** In voting, when each voter is allowed to spread a utility of 1 over the alternatives, it is akin to each voter having ‘one’ vote and ‘equal’ say. Similarly, normalized utilities are considered in resource allocation to define rules such as Adjusted Winner [Brams and Taylor, 1996]. If utilities are not normalized, and the rule is welfarist in some sense, then fairness can be undermined. For example, when considering rules that are responsive to utilitarian social welfare, the outcome is favourable to agents with the most magnified utilities. When considering egalitarian welfare, the allocation is aligned with the concerns of the agent with most scaled down utility valuations.

Justification in a common setting

- (4) **Probabilistic perspective.** Another possible justification for using normalized utilities is using the utility as a measure of certainty of liking an item or outcome.

Technical Justifications

- (5) **Strategic.** The reasons for fairness can also be seen as justification regarding strategic issues. Considering normalized utilities circumvents certain trivial

²The welfare notion of utilitarian social welfare under normalized utilities has been referred to as ‘relative utilitarianism’ [Dhillon and Mertens, 1999].

³Welfare maximization under normalization is not without issues. For example, maximizing utilitarian welfare under normalized utilities is incompatible with resource monotonicity [Sziklai and Segal-Halevi, 2015].

⁴Nash welfare on the other hand is scale invariant. The Spliddit website [Goldman and Procaccia, 2014] which provides a convenient interface for fair division algorithms uses a solution based on the Nash welfare [Caragiannis et al., 2016a] for allocation of goods.

manipulation actions of scaling up or down of utilities. When considering utilitarian welfare normalization prevents agents from overshadowing other agents by reporting extremely high utilities. An agent who reports the highest utilities for all items would get all the items in a utilitarian social welfare maximizing solution!

- (6) **Reasonable approximation guarantees for welfare.** Another reason for considering normalized utilities is also technical and somewhat ‘self-serving’. There is little hope of achieving reasonable approximation guarantees of maximum welfare or reasonable low distortion when utilities are unbounded. Therefore the normalization assumption can be seen as a trick of the trade to obtain reasonable approximation guarantees.

We have discussed reasons for considering (approximate) welfare guarantee for normalized utilities. There are at least two possible points of criticism concerning social welfare under normalized utilities. Firstly, one may negatively perceive that the welfare guarantee results only hold under the restrictive assumption of normalizes utilities. Secondly, there can be a more classical objection to a cardinal approach to social welfare itself which involves interpersonal comparison of utilities (see e.g., [Robbins, 1935]). The second issue may appear especially acute in the context of social choice settings such as fair allocation and voting which typically do not involve money. One reason for pursuing social welfare in computer science research has been that ‘without explicit optimization objective that measures the quality of outcomes, approximation cannot play a role’ [Procaccia and Tennenholtz, 2013].

We point out that by simply pursuing justification (6), one can paradoxically get a justification that avoids both criticisms. One can get a similar guarantee for a concept that does not involve interpersonal comparison of utilities and does not require scaling of utilities of any agent. This largely ignored ‘new’ justification is centered around a relaxation of Pareto optimality.

- (7) **Approximate Pareto optimality under any scaling of the utilities.**

We say that utility profile $u = (u_1, \dots, u_n)$ Pareto dominates utility profile $u' = (u'_1, \dots, u'_n)$ if $u_i \geq u'_i$ for all agents i and $u_i > u'_i$ for some agent i . Given any $\alpha \in [0, 1]$, a utility profile u is α -Pareto optimal if there exists no other achievable utility profile u' such that $\alpha \cdot u'$ Pareto dominates u .⁵ The concept is very natural and has been used in the context of routing games [Aumann and Dombb, 2010], probabilistic matchings [Immorlica et al., 2017] and participatory budgeting [Aziz et al., 2017].

It can be proven that any outcome that achieves α fraction of the maximum social welfare under normalized utilities also satisfies α -Pareto optimality under any scaling of the utilities. Suppose agents have normalized utilities. Consider any outcome that achieves α fraction of the maximum social welfare. Suppose the utility profile of the agents is u . Then we first claim that u is α -Pareto optimal. Suppose it is not α -Pareto optimal. Then there exists another achievable utility profile u' such that $\alpha \cdot u'$ Pareto dominates u . But this means that u achieves less

⁵The definition is written for positive utilities but can be adapted for negative or mixed utilities.

than α fraction of social welfare of u' which is a contradiction. We have established that the outcome achieves α -Pareto optimality under normalized utilities. We note that α -Pareto optimality is invariant under scaling of an agent's utility function even when discrete outcomes are considered. The reason is that for an agent i , any two utility functions u_i and u'_i where $u'_i = \alpha u_i$ and for any social outcomes a, b , $u_i(a) \geq u_i(b)$ if and only if $u'_i(a) \geq u'_i(b)$. Hence, the outcome achieves α -Pareto optimality under any scaling of the agents' utilities.

By normalizing the utilities and achieving or establishing some bound on the maximum utilitarian or egalitarian social welfare also implies the same approximation bound on Pareto optimality for the 'real' utilities of the agents. Thus one can use justification (6) for technical ease but then achieve a guarantee that has more wide-spread normative appeal. The 'new' justification can also be used as another motivation for the projects of implicit utilitarian voting and approximate mechanism design without money. It can also be viewed as a source of corollaries for these lines of work. Taking another view, for new settings, one can directly use approximate Pareto optimality rather than particular social welfare as the social objective. A lower bound result for approximate Pareto optimality would imply a similar lower bound for utilitarian welfare.

We conclude by mentioning that the normative/axiomatic approach in traditional social choice and the quantitative welfarist approaches employed in recent computational social choice papers have been termed as distinct from each other. The use of approximate Pareto optimality provides a convenient bridge between the two approaches.

Acknowledgments

The author thanks Felix Brandt, Hu Fu, Xin Huang, Barton Lee, Ariel Procaccia and Matt Weinberg for useful comments.

REFERENCES

- E. Anshelevich and J. Postl. Randomized social choice functions under metric preferences. In *Proc. of 25th IJCAI*, pages 46–59. AAAI Press, 2016.
 - Y. Aumann and Y. Dombb. Pareto efficiency and approximate Pareto efficiency in routing and load balancing games. In *Proc. of 3rd International Symposium on Algorithmic Game Theory (SAGT)*, pages 66–77, 2010.
 - H. Aziz, A. Bogomolnaia, and H. Moulin. Fair mixing: the case of dichotomous preferences. 2017. Working Paper.
 - D. Baumeister, S. Bouveret, J. Lang, N-T. Nguyen, T-T. Nguyen, J. Rothe, and A. Safidine. Axiomatic and computational aspects of scoring allocation rules for indivisible goods. In *Proc. of 5th COMSOC*, 2014.
 - J. Bentham. *An Introduction to the Principles of Morals and Legislation*. 1789.
 - D. Bertsimas, V. F. Farias, and N. Trichakis. The price of fairness. *Operations Research*, 59(1):17–31, 2011.
 - C. Boutilier, I. Caragiannis, S. I. Haber, T. Lu, A. D. Procaccia, and O. Sheffet. Optimal social choice functions: a utilitarian view. In *Proc. of 13th ACM-EC Conference*, pages 197–214. ACM Press, 2012.
- X, Vol. X, No. X, X 20XX.

- S. Bouveret and M. Lemaître. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents and Multi-Agent Systems*, 2015.
- S. J. Brams and A. D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.
- I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang. The Unreasonable Fairness of Maximum Nash Welfare. In *Proc. of 17th ACM-EC Conference*, pages 305–322, 2016a.
- I. Caragiannis, S. Nath, A. D. Procaccia, and N. Shah. Subset selection via implicit utilitarian voting. In *Proc. of 25th IJCAI*, pages 151–157. AAAI Press, 2016b.
- A. Dhillon and J.-F. Mertens. Relative utilitarianism. *Econometrica*, 67(3):471–498, 1999.
- A. Filos-Ratsikas, S. K. S. Frederiksen, and J. Zhang. Social welfare in one-sided matchings: Random priority and beyond. In *Proc. of 7th International Symposium on Algorithmic Game Theory (SAGT)*, pages 1–12, 2014.
- J. Goldman and A. D. Procaccia. Splidit: Unleashing fair division algorithms. *ACM SIGecom Exchanges*, 13(2):41–46, 2014.
- M. Guo and V. Conitzer. Strategy-proof allocation of multiple items without payments or priors. In *Proc. of 9th AAMAS Conference*, pages 881–888. IFAAMAS, 2010.
- N. Immorlica, B. Lucier, G. Weyl, and J. Mollner. Approximate efficiency in matching markets. In *Web and Internet Economics - 13th International Conference, WINE 2017*, pages 252–265, 2017.
- T. Lu and C. Boutilier. Robust approximation and incremental elicitation in voting protocols. In *Proc. of 22nd IJCAI*, pages 287–293. AAAI Press, 2011.
- A. D. Procaccia and M. Tennenholtz. Approximate mechanism design without money. *ACM Transactions on Economics and Computation*, 1(4), 2013.
- L. Robbins. *An Essay on the Nature and Significance of Economic Science*. Macmillan, 2nd edition, 1935.
- B. Sziklai and E. Segal-Halevi. Resource-monotonicity and population-monotonicity in cake-cutting. *CoRR*, abs/1510.05229, 2015.