

Evaluation report on Grzegorz Pierczynski's PhD Thesis

This report provides an evaluation of the PhD thesis entitled “Proportional Participatory Budgeting”, which has been prepared by Grzegorz Pierczynski in partial fulfilment of the requirements for the degree of Doctor of Philosophy at the University of Warsaw. I have read the PhD thesis carefully, and let me emphasize upfront that, based on the manuscript, I believe that the thesis is sufficient to grant a PhD.

Summary: This is a PhD Thesis in the general area of Computational Social Choice, an important interdisciplinary area at the interface of Theoretical Computer Science (in particular, Algorithms and Complexity Theory) and Economics (in particular, Social Choice Theory).

The thesis consists of nine chapters. Chapter 1 contains an overview of the thesis content, with an overview of the results, including a motivating discussion of the theoretical notions, a high-level presentation of experimental findings, and a discussion on potential applications of the results. Chapter 2 presents the main modelling assumptions, together with basic notation and preliminary definitions. Additional notation is introduced later in the technical chapters.

The rest of the thesis is split into two parts. Part I includes Chapters 3-6 and focuses on the Method of Equal Shares and the proportionality axioms it satisfies. The material presented there appeared originally at the conferences NeurIPS 2021 and IJCAI 2023. Part II includes Chapters 7-9 and considers proportionality axioms which are even stronger than those satisfied by the Method of Equal Shares. The new technical material presented in these chapters appeared originally at WINE 2022 and AAI 2021, while several notions originally defined in the author's NeurIPS 2021 paper are also relevant here.

In particular, Chapter 3 defines the Method of Equal Shares (MES) by refining a method known as Rule X in previous work by Peter Skowron (the author's PhD supervisor) and Dominik Peters. The focus of Rule X has been on approval-based multi-winner elections, while MES is suitable for use in more general settings, such as participatory budgeting with different ballot formats. The refinement of the original definition of MES constitutes an important conceptual contribution of the thesis. The original paper has attracted a lot of attention from the research community as its 90+ citations (according to Google Scholar) indicate.

Chapter 4 presents theoretical results for MES. In particular, it presents several definitions of the extended justified representation (EJR) axiom and adapts it to settings with different agent utilities (binary or cardinal). Then, the axiom EJR-1, extended justified representation up to one alternative/project, is defined and proven to be satisfied by MES. EJR-1 appears to be a reasonable relaxation of EJR, due to the inherent difficulty in finding EJR outcomes. Also, Chapter 4 presents a result about the (approximate) core stability of MES. Let me remark that the “Definition attempt” environment used here is

Dept. of Computer Science

Ioannis Caragiannis
Professor

Date: 16 February 2024

Voice: +45 4189 3304
E-mail: iannis@cs.au.dk
Web: cs.au.dk/~iannis

Page 1/3



very useful so that the reader can follow how the technical definitions are connected to the intuition behind the corresponding axioms.

Chapter 5 presents an empirical evaluation of MES. MES is compared to the utilitarian greedy method on data from 800 participatory budgeting elections that have been held (mainly) in Poland in recent years. Many interesting conclusions are drawn regarding several fairness and efficiency criteria. Interestingly, the definition of MES allows for different variants depending on the completion strategy used to guarantee that the available budget is exhausted. A large number of plots, including election maps, support the findings graphically. Chapter 5 concludes with a discussion on how the ballot format can affect the outcome of participatory budgeting elections when utilitarian greedy or MES is used. Chapter 6 extends MES to ordinal preferences.

Chapter 7 first introduces the axiom of full justified representation (FJR). FJR is proven to be more demanding than EJR in participatory budgeting settings and is not satisfied by MES. The author presents the Greedy Cohesive Rule (GCR), which is proven to produce priceable outcomes, provided that an appropriate completion of the outcome is used. Some drawbacks of GCR compared to MES are also discussed.

Chapter 8 studies the axiom of core stability in multi-winner elections. This is an important proportionality axiom that, unfortunately, is not satisfied in general. Here, the focus is on sufficiently restricted domains, where positive statements about the existence of outcomes in the core can be proven. A long list of restricted elections is considered both for ranking-based (e.g., single-peaked, single-crossing, 1D-Euclidean preferences) and approval-based (e.g., candidate/voter interval and linearly-consistent preferences) elections. The main contribution is the novel Quantile Rule, which is proven to compute outcomes in the core for several restricted types of elections. The chapter concludes with a rather extensive discussion on open problems and directions for future research on the topic.

Chapter 9 presents an interesting perspective that views participatory budgeting with approval preferences as a market-based system. Stable priceability is the important concept introduced here. The main results include implication statements that connect this notion, budget exhaustiveness, core stability, and EJR. The extended version of balanced stable priceability is also discussed. The chapter also contains experimental results on synthetic data.

The main part of the thesis concludes with a summary section, with an overview of the material presented in Chapters 1-9 and a short discussion on open problems and future work. I believe that this part could be extended considerably (see my suggestion below). Some details about the format of inputs of the Pabulib library that is used for the experiments in Chapters 5 and 9 are given in appendix A.

Evaluation: Overall, I found the thesis to be of very good quality, with important conceptual and technical contributions. The refinement of the method of equal shares and the new proportionality axioms defined are very important and will definitely motivate future work by the computational social choice community (they already do, actually). The technical questions considered are challenging, the results are novel and important,



and the proofs are correct (as far as I have checked) and, in general, of sufficient technical depth. The thesis reports work originally published at NeurIPS, AAAI, and IJCAI, the flagship conferences on the foundations of AI, and at the competitive WINE conference. The quality of the write-up is very good, in general.

Let me add two suggestions for possible improvements that the author may consider in the next revision of the thesis text. First, I believe that the thesis would benefit by including a more detailed list of open problems in the final summary chapter. In particular, open problems and future work are explicitly mentioned only in Chapter 8. In the remaining chapters, there are a few references to questions that are not covered by the technical statements that would be interesting to study, but there is no reference to their importance or how they fit in the general picture of proportionality in participatory budgeting. Second, even though the thesis includes material that has previously appeared in the literature in the form of papers published in conference proceedings or in the arXiv repository, the structure followed here is very different (probably with the exception of Chapters 8 and 9). This is an excellent (and rather non-standard) choice by the author and can give the reader a general view of proportional participatory budgeting. However, it is sometimes difficult to trace the origin of the several definitions or statements in the literature. I believe that repeating the pointers to the literature in each chapter and explaining the differences (where applicable), would improve readability.

In conclusion, I deem the thesis as sufficient to grant a PhD.

Sincerely,

Ioannis Caragiannis
Professor

