Synthesis Lectures on
Operations Research and Applications

SYNTHESIS
COLLECTION OF TECHNOLOGY

Özgün Elçi · John Hooker · Peter Zhang

The Structure of Fair Solutions

Achieving Fairness in an Optimization Model

Springer

Synthesis Lectures on Operations Research and Applications

Series Editor

James Cochran, Ames, USA

This series focuses on the use of advanced analytics in both industry and scientific research to advance the quality of decisions and processes. Written by international experts, modern applications and methodologies are utilized to help researchers and students alike to improve their use of analytics. Classical and cutting-edge topics are presented and explored with a focus on utilization and application across a range in practical situations.

Özgün Elçi · John Hooker · Peter Zhang

The Structure of Fair Solutions

Achieving Fairness in an Optimization Model

Springer

Özgün Elçi
Amazon
Seattle, WA, USA

John Hooker
Tepper School of Business
Carnegie Mellon University
Pittsburgh, PA, USA

Peter Zhang
Carnegie Mellon University
Pittsburgh, PA, USA

ISSN 2770-6303 ISSN 2770-6311 (electronic)
Synthesis Lectures on Operations Research and Applications
ISBN 978-3-031-82189-9 ISBN 978-3-031-82190-5 (eBook)
https://doi.org/10.1007/978-3-031-82190-5

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Preface

Fairness is often an important consideration in the optimization models that lie behind business and public policy decisions. Yet how can fairness be defined mathematically? A number of fairness models have been proposed, but one must somehow select the right model for a given application. In this book, we address this challenge by deducing general characteristics of optimal solutions that result from various fairness models. One can then judge which type of solution is acceptable in a given context. We focus on fairness criteria that combine efficiency with fairness, since typically both are important in practice. Most of the results are new, and some are quite unobvious and even surprising.

Our aim has been to provide a concise but mathematically rigorous presentation that is directed toward several audiences. The practitioner can begin with the self-contained overview of our findings in the final chapter, and then refer to previous chapters as needed for details. Optimization modelers can make use of the linear and integer programming formulations we provide. Our proof techniques may be useful to fairness researchers who wish to obtain further results. Welfare economists may find of interest our marginal analyses of stakeholder incentives and cross-subsidies. Most of all, we hope to encourage and facilitate the current trend toward taking fairness seriously in decision models.

Özgün Elçi's contributions to this work were completed prior to joining Amazon.

Seattle, WA, USA — Özgün Elçi
Pittsburgh, PA, USA — John Hooker
Pittsburgh, PA, USA — Peter Zhang
November 2024

Contents

1 Introduction

1.1 The Fairness Modeling Challenge

There is growing interest in incorporating equity and fairness into optimization models. This concern arises in a number of application areas, including healthcare, disaster relief, facility location, task assignment, telecommunications, and machine learning. The aim is to maximize total benefit or minimize total cost, but only while ensuring that benefits and costs are equitably distributed among the relevant stakeholders.

For example, if donated organs are allocated in the most economically efficient fashion, patients with certain medical conditions may wait far longer for a transplant than other patients [1]. If earthquake shelters are located so as to minimize average distance from residents, persons living in less densely populated areas may have much further to travel [2]. If a machine learning algorithm awards mortgage loans so as to maximize expected earnings, members of a minority group may find themselves unable to obtain loans even when they are financially responsible [3]. If traffic signals at intersections are timed to maximize traffic throughput, motorists on side streets may have to wait forever for a green light [4].

The incorporation of fairness into an optimization model raises the issue of how it should be formulated mathematically. One attractive mechanism for doing so is to introduce a *social welfare function* (SWF) that reflects the desirability of a given distribution of benefits or costs. The SWF is a function of a utility vector, where each component of the vector represents the utility enjoyed by an individual stakeholder. Utility can be any relevant type of benefit, such as profit, health, or negative cost. The welfare economics literature [5] distinguishes *ordinal* SWFs, which rank utility vectors by desirability, from *cardinal* SWFs, which assign numerical values to utility vectors, with more desirable distributions receiving larger values. We study cardinal SWFs because they can serve as the objective function of an optimization model. The modeler can select a SWF that strikes an appropriate balance between total utility and fairness, and then maximize the SWF.

Ö. Elçi et al., *The Structure of Fair Solutions*, Synthesis Lectures on Operations Research and Applications, https://doi.org/10.1007/978-3-031-82190-5_1