# Fairness Methods in Optimization and Artificial Intelligence

# Violet (Xinying) Chen

Tepper School of Business Carnegie Mellon University

Committee: John Hooker (Chair) Hoda Heidari Fatma Kılınç-Karzan Alec Morton

This dissertation is submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Operations Research



May 2022

### Acknowledgements

I am extremely fortunate to have received the support, guidance and friendship from many people throughout my journey to the completion of this dissertation and the PhD program.

First of all, I would like to express my deepest gratitude to my advisor, John Hooker. I am indebted to him for opening up my horizon on ethics, an area I never expected to cross into when I started the program. I have been, and will always be, inspired by his passion and insights for interdisciplinary research and diligence to exploring new questions.

I am also deeply grateful to all my committee members. I thank Fatma Kılınç-Karzan for her extensive guidance and advice on research, teaching and service throughout my years in the PhD program. I thank Hoda Heidari for all the interesting discussions related to fairness in artificial intelligence, which greatly expanded my knowledge about the area, and her genuine support throughout my job search. I thank Alec Morton for the insightful questions raised during both my proposal and defense. I enjoyed thinking through these questions and discovering new perspectives.

I would like to thank my wonderful collaborators, Joshua Williams and Derek Leben, in completing Chapter six. I was fortunate to learn new techniques and knowledge from both of them during the collaboration. I also wish to acknowledge the helpful guidance from Thiago Serra, who shared service opportunities with me to get involved with the broader community and offered valuable suggestions for my job search. I also thank all the professors that I have had the opportunity to interact with and learn from throughout these years.

I thank all the precious friendships which have made even the challenging times during the past five years enjoyable. Many thanks to my cohorts: Sagnik Das, Özgün Elçi, Melda Korkut, Kyra Gan, Su Jia, Thomas Lavastida, Yuyan Wang. Special thanks to Yuyan (and her dog Mochi) for sharing a roof with me for three years and being there for all the ups-and-downs, and Sagnik Das for always selflessly helping out whenever needed. I also thank Musa Çeldir, Serim Hwang, Anthony Karahalios, Neda Mirzaeian, Yanhan (Savannah) Tang, Ziye Tang and Ryan Shi for all the happy memories. Additionally, I thank the support from Daniel de Roux, Tom Krumpolc and Mik Zlatin while working together as the board of CMU Informs Chapter.

Lastly, I am deeply indebted to my family in China. To my parents, Jian Chen and Yufeng Gao, I cannot begin to express my gratitude. They are the most supportive and loving parents. They have always believed in me and provided the best opportunities for me. To my dog, Model, I thank her for all the happiness and love.

### Abstract

With increasing deployment of optimization and Artificial Intelligence (AI) to assist high-stake real life decisions, fairness has become an essential factor of consideration for both the designers and users of these tools. This dissertation studies new approaches for formulating, attaining and eliciting fairness. Chapter one begins with a brief introduction of the background on fairness and a selection of common fairness measures.

Chapter two studies balancing fairness and efficiency through optimization models. We propose new social welfare functions (SWFs) as combined measures of two well-known criteria, Rawlsian leximax fairness and utilitarianism. We then design a procedure to sequentially maximize these SWFs with mixed integer/linear programming models to find socially optimal solutions. This approach has practical potentials on a wide range of resource allocation applications, and is demonstrated on realistic size applications in healthcare provision and shelter assignment for disaster preparation.

Chapter three considers an optimization task motivated by fair machine learning (ML). When developing fair ML algorithms, it is useful to understand the computational costs of fairness in comparison to the standard non-fair setting. For fair ML methods that utilize optimization models for training, specialized optimization algorithms have potentials to offer better computational performances than generic solvers. In this chapter, I explore this question for support vector machines (SVMs), and design block coordinate descent type algorithms to train SVMs containing linear fairness constraints. Numerical experiments highlight that the new specialized algorithms are more efficient than an off-the-shelf solver for training fair SVMs.

Chapter four examines social welfare optimization as a general paradigm for formalizing welfarebased fairness in AI systems. Contrary to commonly used statistical bias metrics in fair AI, optimizing a social welfare objective supports broader perspective on fairness motivated by distributive justice considerations. We propose in-processing and post-processing integration schemes between social welfare optimization and AI, in particular, ML and rule-based AI. We implement and evaluate the integration schemes on a simulated loan processing instance. The empirical results demonstrate the advantages of the proposed integration strategies. We conclude this chapter by highlighting research directions to pursue for a holistic view of welfare-based fairness.

The next two chapters explore the human-centric perspective to elicit people's moral values through preference learning. Chapter five studies a general preference learning framework based on online learning (OL) from revealed preferences: a learner learns an agent's private utility function through interactions in a changing environment. Through designing a new convex loss function, we

design a flexible OL framework that enables a unified treatment of usual loss functions from literature and supports a variety of online convex optimization algorithms. This framework has advantages in regret performance and solution time over other OL algorithms from the literature.

Lastly, chapter six explores a moral decision-making inspired task. This chapter considers the modelling and elicitation of people's dynamic ethical judgments in the sequential allocation of resources. We utilize a Markov Decision Process model to represent a sequential allocation task, where the state rewards capture people's moral preferences, thus people's ethical judgments are reflected via policy rewards. We design a preference inference model which relies on active preference-based reward learning to infer the unknown reward function. The learning framework is applied in human-subject experiments on Amazon Mechanical Turk to understand people's moral reasoning in a hypothetical scenario of allocating scarce medical resources.

# **Table of contents**

1	Intro	roduction				
	1.1	Background on Fairness	2			
		1.1.1 Utility-based Fairness	3			
		1.1.2 Parity-based Fairness	7			
	1.2	Outline and Contribution	9			
2	Com	Combining Leximax Fairness and Efficiency in a Mathematical Programming Model				
	2.1	Introduction	. 1			
		2.1.1 Our Approach and Results	.3			
		2.1.2 Basic Definitions	.3			
		2.1.3 Related Literature	.4			
	2.2	Defining the Social Welfare Functions	20			
<ul><li>2.3 Properties of the Social Welfare Functions</li></ul>		Properties of the Social Welfare Functions	22			
		The Sequential Optimization Procedure	26			
	2.5	A Small Example	28			
	2.6	Mixed Integer Programming Model	29			
	2.7	Valid Inequalities	31			
	2.8	Modeling Groups of Individuals	33			
	2.9	Applications	;9			
		2.9.1 Healthcare Resource Allocation	39			
		2.9.2 Shelter Location and Assignment	1			
	2.10	Conclusion	6			
3	Spec	ialized Algorithms for Training Fairness Constrained Support Vector Machines 4	19			
	3.1 Introduction		9			
		3.1.1 Our Approach and Results	50			
		3.1.2 Related Literature	51			
	3.2	Problem Formulation	53			
		3.2.1 Supported Fairness Constraints	55			
	3.3	Fair SVM Training Algorithms    5	56			

		3.3.1	Training Standard SVM	56	
	2.4	3.3.2	Iraining Fair SVM	. 60	
	3.4			05	
		3.4.1	Asymptotic Convergence of Fair SMO for Fair General SVMs		
		3.4.2	Linear Convergence of Fair Coordinate Descent for Fair Linear SVMs.	69	
	<u>а</u> г	3.4.3	Finite Termination of Fair Coordinate Descent for Fair Linear SVMS.	70	
	3.5			73	
		3.5.1		74	
		3.5.2			
		3.5.3			
	2.0	3.5.4 Canali		80	
	3.6	Concil			
4 Fairness through Social Welfare Optimization		ough Social Welfare Optimization	83		
	4.1	Introdu		83	
		4.1.1	Our Approach and Results		
		4.1.2			
	4.2	The Ba		86	
		4.2.1	Advantages of Optimization	87	
	4.3	Examp	ble: Mortgage Loan Processing	88	
	4.4	A General Framework			
	4.5	A Sam	pling of Social Welfare Functions	92	
		4.5.1	Pure fairness measures	92	
		4.5.2	Combining fairness and ef ciency	93	
		4.5.3	Statistical bias metrics	94	
	4.6	Case S	Study: Mortgage Loan Processing	95	
		4.6.1	Decision models	96	
		4.6.2	Dataset and implementation details	97	
		4.6.3	Results and ndings	98	
	4.7	Conclu	usion	100	
5	Onli	ne Con	vex Optimization Perspective for Learning from Dynamically Revealed	b	
	Pref	references 103			
	5.1	Introdu		103	
		5.1.1	Our Approach and Results	104	
		5.1.2	Related Literature	105	
	5.2	Proble	m Setting	107	
		5.2.1	Forward Problem	108	
		5.2.2	Inverse Problem	111	

		5.2.3	Online Inverse Optimization	112	
	5.3	.3 Loss Functions		113	
	5.4	Regret	Performance for <sup>im</sup> -based Online Inverse Optimization	115	
		5.4.1	Perfect Information	115	
		5.4.2	Imperfect Information	118	
	5.5	Online	Learning Algorithms	119	
		5.5.1	Online Convex Optimization with First-Order Oracle	119	
		5.5.2	Implicit Online Learning with a Solution Oracle	120	
		5.5.3	Comparison with the Existing Approaches	121	
	5.6	Compu	utational Study	123	
		5.6.1	Problem Instances	123	
		5.6.2	Implementation Details	125	
		5.6.3	Perfect Information Experiments	127	
		5.6.4	Imperfect Information Experiments	130	
6	Mod	deling and Inferring Dynamic Ethical Judgments Around High-stakes Allocations 133			
	6.1	Introdu		133	
		6.1.1	Related Literature	137	
	6.2	A Mathematical Framework for Dynamic Moral Preferences			
		6.2.1	MDP Model	139	
		6.2.2	Moral Preference Model	142	
		6.2.3	An Active Learning Scheme to Learn Moral Preferences	142	
		6.2.4	Mutual Information-based Query Selection	143	
	6.3	Experimental Design			
		6.3.1	Scenario and Setup	145	
		6.3.2	Data Generation	147	
		6.3.3	Heuristics for Query Generation and Selection	149	
		6.3.4	Establishing Viability through Simulations	150	
	6.4	Experi	mental Findings and Discussion	151	
	6.5			157	
Ro	foron	000		150	
I C	leren	003		100	
Appendix A Appendix to Chapter 3					
	A.1	Pseudo		169	
		A.1.1	Fair Sequential Minimal Optimization for Fair General SVM	169	
		A.1.2	Fair Dual Coordinate Descent for Fair Linear SVM	172	
	A.2	Additio	nal Experiment Results: Accuracy and Fairness Measures	173	

Append	ix B A	ppendix to Chapter 5	179
B.1	Conve	exity Status of <sup>pre</sup> and <sup>est</sup>	179
B.2 ` <sup>sub</sup> -based Online Learning Algorithms			180
	B.2.1	OCO with First-Order Oracle	180
	B.2.2	OCO with Solution Oracle	183
B.3 Formulations for the Solution Oracles Used in the Implicit OL Algorithms			185
	B.3.1	Solution Oracle for <sup>sim</sup> -based Implicit OL Algorithm	186
	B.3.2	Solution Oracle for pre-based Implicit OL Algorithm	186
B.4 ` <sup>sim</sup> -based Online Inverse Optimization under Imperfect Information			188
Append	ix C A	ppendix to Chapter 6	193
C.1	Additi	onal Experiment Results	193

# Chapter 1

# Introduction

<sup>1</sup> Fairness is important to people's daily lives and society. As a key component of ethics, fairness seeks equitable distribution of resources and opportunities to promote equality, justice and social well-beings. In real-life decisions and policies, it is morally desirable and socially sustainable to pursue fairness. With optimization and Arti cial intelligence (AI) tools increasingly applied to assist high-stake decisions, such as, optimizing humanitarian operations with resource allocation models, using machine learning models for criminal risk assessment, fairness is no longer a nice-to-have advantage, but rather a necessary property of practical and trustworthy decision support models and algorithms. This dissertation speaks to the rising demand of better understanding of fairness, and develops new approaches and insights for fairness in optimization and AI.

In contrast to the widely recognized signi cance of fairness, there is no uniformly superior interpretation and operationalization of fairness. Besides the large variety of fairness related theories proposed across academic elds, people hold diverse fairness perceptions as in uenced by their personal backgrounds and speci c decision contexts. This dissertation explores two perspectives for the fundamental questions that is fairness and how to be fair

One perspective is the formulation of fairness optimization models in alignment with ethics principles. Optimization has long been used to support decision making. Most practical optimization problems involves the allocation decisions of some resources. In recent years, optimization also has been extensively studied in arti cial intelligence, especially in machine learning (ML) where optimization models are often core components in a ML framework. Conventionally, optimization models have a ef ciency-driven objective function. For resource allocation decisions, the objective describes seeking the most effective use of resources. For example, a government agency aims to maximize the total population bene ts from healthcare provision, a company tries to minimize costs and maximize pro ts during facility expansion. For optimization-based ML methods, the standard objective focuses on minimizing the total prediction losses to seek high accuracy. By pursuing ef ciency goals, these conventional optimization models may lead to unfair results, so a crucial task is to incorporate fairness and equity into existing optimization models. While it is normally

<sup>&</sup>lt;sup>1</sup>This chapter uses excerpts from a joint work with John Hooker

straightforward to formulate an ef ciency objective, fairness can be understood in multiple ways, with no generally accepted method for representing any of them in a mathematical idiom. Moreover, the growth of integrated methods, such as, optimization based on ML predictions, in modern decisionmaking adds to the challenge of formulating fairness models. As discussed above, this perspective opens up interesting modeling and computational questions in the broad eld of AI. Chapter 2, 3, 4 studies such questions.

The other perspective is to elicit fairness beliefs, and more generally ethical values, from human stakeholders. Conventional approaches to de ning and justifying fairness are driven by principles, namely, the central planner determines what are the suitable ethical values and what should be the resulting fairness notion. Due to the large variety of fairness concepts and the potentially different positions held by the planner and stakeholders, the chosen fairness notion could be incompatible with stakeholders' moral beliefs and ethics principles. The possible incompatibility has led to the recent research thread of human-centric fairness, a bottom-up strategy aiming to learn what people believe to be fair under different decision contexts then bring their judgments into the formulation of fairness. While there are situations where people's perspectives should not be incorporated due to irrationality or bias, in general, decision makers would seek to align fairness interventions with what the people desire, so that they would be more welcoming to these interventions. A core component of this perspective is the modeling, elicitation and learning of preferences: chapter 5 studies a general preference learning question in an online setup, and chapter 6 explores a concrete moral preference learning task.

## 1.1 Background on Fairness

Fairness has a long history of being studied across elds including philosophy, sociology, psychology and economics. These early literature provide ethical and conceptual foundations for the more recent discussion of fairness in operations research and AI, where the goal is to operationalize fairness via optimization and machine learning for concrete applications. A well known framework to distinguish fairness concepts is to perceive fairness via distributive justice versus procedural justice. The highlevel intuition is that distributive justice concerns fairness in the outcomes from a decision, whereas procedural justice emphasizes fairness in the process of decision making. This dissertation focuses on outcome fairness via the distributive justice perspective.

Mathematical formalization of outcome fairness can be further distinguished into utility-based and parity-based de nitions. Utility-based fairness is broadly used in resource allocation to attain a fair distribution of utilities, which can be pro ts, negative costs, or some other bene ts appropriate to the application. Utilities were initially proposed as part of the theory of utilitarianism, which seeks to maximize the overall well-beings of the population. In particular, the standard ef ciency formulation follows utilitarianism and uses the total or average utility as a measure of the population's overall welfare. Parity-based fairness is typically considered in machine learning to seek unbiased treatments towards the involved groups or individuals. There is a large number of fairness de nitions in both types, and samples of popular choices in literature and practice are summarized below.

### 1.1.1 Utility-based Fairness

Suppose  $\mathbf{u} = (\mathbf{u}_1; \ldots; \mathbf{u}_n)$  is a vector of utilities distributed across parties::; n and  $W(\mathbf{u})$  is a utilitybased fairness measure ( $\mathbf{u}$ ) is a function aggregating the utility values to evaluate ith respect to the selected fairness notion. We categorize utility-based fairness criteria into three types.

#### **Inequality Measures**

The rst type measures fairness via the degree of equality in the distribution of utilities, for which several statistical metrics have been proposed (Cowell 2000, Jenkins and Van Kerm 2011). There is a wide variety of philosophical opinion on the ethical signi cance of equality, ranging from the view that we have an irreducible obligation to strive for equality, to the view that inequality is unfair only when it reduces total utility (Frankfurt 2015, Par t 1997, Scanlon 2003). In any event, it is generally acknowledged that equality is not the same concept (or cluster of concepts) as fairness, even when the two are closely related. An equality metric can be appropriate in a context where a speci cally egalitarian distribution is the primary goal, without regard for efficiency or other forms of equity. Inequality measures have been used for inequity averse optimization in a broad range of applications. More recently, inequality measures are also considered in the growing area of algorithmic fairness. We next introduce several commonly studied de nitions.

Measures of relative dispersion<sup>2</sup>. The relative rangeof utilities is an inequality metric, that, when negated, yields the fairness measure,

$$W(\mathbf{u}) = (1 = \bar{u}) u_{max} u_{min}$$

where  $u_{max} = max_i f u_i g$ ,  $u_{min} = min_i f u_i g$ , and  $u = (1=n) a_i u_i$ .

Another dispersion metric is the lative mean deviation which measures inequality more comprehensively by considering all utilities rather than only the minimum and maximum. The corresponding fairness measure is,

$$W(\mathbf{u}) = (1=\bar{u}) \mathop{a}_{i}^{*} ju_{i} \bar{u}_{j}$$

The coef cient of variation is the normalized standard deviation, leading to the fairness measure below. It may be appropriate when large deviations from the mean are disproportionately signi cant, but it has the possible drawback of computational dif culty due to the quadratic component.

$$W(\mathbf{u}) = \frac{1}{\bar{u}} \frac{h_{1}}{n} \overset{*}{a}_{i} (u_{i} \ \bar{u})^{2} \frac{i_{1}}{2};$$

<sup>&</sup>lt;sup>2</sup>All of the following dispersion measures are normalized by the mean utility so as to be invariant under rescaling of utilities.

Gini coef cient and Hoover index. The Gini coef cient is by far the best known measure of inequality, as it is routinely used to measure income and wealth inequality (Gini 1912). It is proportional to the area between the Lorenz curve and a diagonal line representing perfect equality. Lorenz curve plots the proportion of the total wealth or bene ts accumulated to the bottooth the population, thus indicating the degree of inequality via its deviation from the perfect equality line. By de nition, Gini coef cient vanishes under perfect equality. We can use the negative Gini coef cient as a fairness measure,

$$W(\mathbf{u}) = \frac{1}{2\bar{u}n^2} \mathbf{a}_{i;j}^2 \mathbf{j} u_i \quad u_j \mathbf{j}:$$

The Hoover indexis also related to the Lorenz curve, as it is proportional to the maximum vertical distance between the Lorenz curve and a diagonal line representing perfect equality (Hoover 1936). It can be interpreted as the fraction of total utility that would have to be redistributed to achieve perfect equality. The corresponding fairness measure is

$$W(\mathbf{u}) = \frac{1}{2n\bar{u}} \mathop{a}_{i}^{a} ju_{i} \quad \bar{u}j:$$

### Fairness for the Disadvantaged

Rather than focus solely on inequality, fairness measures can prioritize the disadvantaged. Far and away the most famous of such measures is the difference principle of John Rawls (1999), a maximin criterion that is based on careful philosophical argument and debated in a vast literature (surveyed in Freeman 2003, Richardson and Weithman 1999). The difference principle can be plausibly extended to a lexicographic maximum principle. There is also the McLoone index, which is a statistical measure that emphasizes the lot of the less advantaged.

The Rawlsian maximin criterion has been a popular fairness measure for decades. Early works on fair resource allocation, such as bandwidth allocation, often choose the maximin criterion to seek the best possible performance for the worst-off service among services competing for bandwidth. Recent research has applied the criterion to more diverse problem contexts, including peer review paper assignment, ridesharing, etc.

Rawlsian criteria. The Rawlsiandifference principlestates that inequality should exist only to the extent that it is necessary to improve the lot of the worst-off. It is defended with a social contract argument that, in its simplest form, maintains that the structure of society must be negotiated in an "original position" in which people do not yet know their station in society. Rawls argues that one can rationally assent to the possibility of ending up on the bottom only if that person would have been even worse off in any other social structure, whence an imperative to maximize the lot of the worst-off. The principle is intended to apply only to the design of social institutions, and only to the distribution of "primary goods," which are goods that any rational person would want. Yet it can be adopted as a general criterion for distributing utility, namely aximincriterion that maximizes the smallest utility as the fairness measure,

$$W(\mathbf{u}) = \min_{i} u_i g:$$

The maximin criterion can be plausibly extended excoorgraphic maximizatio (leximax). Leximax is achieved by rst maximizing the smallest utility subject to resource constraints, then the second smallest, and so forth.

McLoone index. The McLoone indexcompares the total utility of individuals at or below the median utility to the utility they would enjoy if all were brought up to the median utility. The index is 1 if nobody's utility is strictly below the median, and it approaches 0 if the utility distribution has a very long lower tail (on the assumption that all utilities are positive). The McLoone index bene ts the disadvantaged by rewarding equality in the lower half of the distribution, but it is unconcerned by the existence of very rich individuals in the upper half. The de nition is,

$$W(\mathbf{u}) = \frac{1}{jI(\mathbf{u})j\tilde{u}} \overset{a}{\underset{i \ge I(\mathbf{u})}{a}} u_{i}$$

where  $\tilde{u}$  is the median of utilities in u and (u) is the set of indices of utilities at or below the median, so that  $(u) = f i j u_i$   $\tilde{u}_j$ .

#### Combined Fairness and Ef ciency

The previous examples are all pure fairness measures, which are appropriate when there is no need to balance fairness against the overall well-being of the population. However, practical situations frequently call for both fairness and efficiency to be explicitly considered. We next review three common schemes for combining the two criteria and give examples of the combined measures following each scheme. Chapter 2 provides further discussion of these measures.

Convex combination. The most obvious approach is to maximize a convex combination of fairness and ef ciency. Suppose  $(\mathbf{u})$  is a pure fairness measure, then a combined measure is

$$W(\mathbf{u}) = (1 | \mathbf{i}) \underset{i}{a} u_i + \mathbf{i} F(\mathbf{u}):$$

The combination strategy is applicable to both inequality indices and fairness for the disadvantaged. For instance, Eisenhandler and Tzur (2019) and Mostajabdaveh et al. (2019) propose combined measures with variations of the Gini coef cient  $\mathbf{a}$  (u). Yager (1997) and Ogryczak a (u) = min<sub>i</sub>f u) = min<sub>i</sub>f u<sub>i</sub>g.

Alpha Fairness. Alpha fairness provides an alternative and perhaps more satisfactory means of combing fairness and ef ciency than convex combinations. Alpha fairness regulates the combination with a continuous parameter, where larger values of signify a greater emphasis on fairness. It

(Mo and Walrand 2000, Verloop et al. 2010) is represented by a family of functions having the form

Alpha fairness characterizes a continuum that stretches from a utilitarian criterien0 to a maximin criterion as ! ¥. Lan et al. (2010) provide an axiomatic treatmenta fairness in the context of network resource allocation, and Bertsimas et al. (2012) study worst-case fairness/ef ciency trade-offs implied by this criterion. The parametecan be interpreted as quantifying the trade-off as follows: utility  $u_j$  must be reduced b( $u_j = u_i$ )<sup>a</sup> units to compensate for a unit increase parties, as we desire, with a indicating how much priority. Yet it is not obvious what kind of trade-off, and therefore what value of a, is appropriate for a given application. There is no apparent interpretation independent of its role in thea-fairness function.

Threshold Criteria. Williams and Cookson (2000) suggest two ways to combine utilitarian and maximin objectives using threshold criteria for two persons. One, based on an ef ciency threshold, begins with a maximin criterion but switches to a utilitarian criterion when the overall utility cost of fairness becomes too great. More explicitly, Williams and Cookson uses a maximin criterion when the two utilities are suf ciently close to each other wjth  $u_2j$  D, otherwise it uses a utilitarian criterion. Hooker and Williams (2013) provide appersonextension for this criterion. Parities with  $u_i$   $u_{min}$  D are treated in solidarity with the worst-off through the maximin criterion, and the remaining greater utilities are treated as themselves. The combined fairness measure is the following,

$$W(\mathbf{u}) = (n \quad 1)D + \overset{n}{\underset{i=1}{a}} \max u_i \quad D; u_{\min}g:$$

It is evident tha D = 0 corresponds to a purely utilitarian objective aDel = 0 corresponds to a purely utilitarian objective.

Williams and Cookson also propose the reverse perspective based on a fairness threshold: a utilitarian criterion is used when  $u_1 u_2 j$  D, then a maximin criterion is used since the inequity has become too severe. Following the same argument from Hooker and Williams (2012), this de nition can be extended to the personcase and leads to the combined measure below. The main difference with the previous de nition is that, parities with  $u_{min}$  D are now counted as equal  $u_{\text{Plin}}$  and the other utilities are counted as themselves.

$$W(\mathbf{u}) = nD + \overset{n}{\underset{i=1}{a}} minf u_i \quad D; u_{min}g:$$

Here, D = 0 corresponds to a purely maximin objective **and**  $\neq$  to a purely utilitarian objective.

In Chapter 2, we illustrate that threshold-based combinations that rely on the maximin notion for fairness are sensitive to the utility level of only the very worst-off party. The equity situation of other disadvantaged parties become irrelevant, so long as their utilities are **Dittfithe** lowest. As a result, fairness among the disadvantaged plays almost no role in the solution. This situation can be addressed to a great degree by replacing maximin fairness with leximax fairness. We defer details on these measures combining leximax fairness to Chapter 2.

### 1.1.2 Parity-based Fairness

In the eld of machine learning, there have been rising demands for fairness considerations to eliminate prejudice favoritism toward an individual or a group based on their inherent or acquired characteristics. One well-known example that motivates extensive interest in fair ML is the series of research efforts on whether the COMPAS software, supported by a recidivism risk prediction algorithm, is biased against African-Americans (Angwin et al. 2016, Dieterich et al. 2016). The focus of fair ML is primarily on mitigating this kind of bias and ensuring that certain minority groups, often de ned by law, receive fair treatment. Among ML frameworks, fairness in supervised learning has been most widely studied. The community has seized upon traditional statistical measures of classi cation error to detect bias, so that it can be avoided when possible.

In a typical scenario, an ML model is trained to make yes–no decisions as to who receives a certain bene t, such as a mortgage loan, a job interview, parole, and so forth, based on various features they possess. A fairness test compares decisions for a minority or protected group with those for the remainder of the population. Four outcomes are possible for each individual: a true positive (the ML model correctly selects the individual for a bene t), a false positive (it incorrectly selects), a true negative (it correctly rejects), and a false negative (it incorrectly rejects). We will refer to the number of individuals in these four groups, respectively, as TP, FP, TN, and FN. Various metrics involving these four statistics are compared between the minority group and the rest of the population, each yielding a measure of parity between the groups.

We setd<sub>i</sub> = 1 when individuali should be selected, and = 0 otherwise, and = 1 when i is selected and = 0 otherwise. We leN<sub>1</sub> be an index set for individuals in the protected group, and N<sub>0</sub> for those in the remainder of the population. For simplicity of exposition, we consider the minority group as the protected group, and the majority as the remainder. Group fairness measures W(**d**) are de ned directly in terms of the decision vector Four of the widely studied de nitions are summarized below. We note that fair ML literature also considers other types of de nitions, including individual fairness, counterfactual fairness, and more recently, utility-based fairness (Chapter 4 reviews relevant literature).

Demographic parity. The simplest bias metric is based **de**mographic parityalso known as proportional/statistical parity. This metric was introduced by Dwork et al. (2012). It seeks to equalize the fraction of minority individuals selected for bene ts and the fraction of majority individuals selected. It is de ned by comparing the ra( $i\overline{D}P$ + FP)=(TP+ FP+ TN+ FN) across the two groups.

The demographic parity is de ned as:

$$W(\mathbf{d}) = \frac{1}{jN_{1}j} \overset{a}{\underset{i2N_{1}}{a}} d_{i} \quad \frac{1}{jN_{0}j} \overset{a}{\underset{i2N_{0}}{a}} d_{i}:$$

A main critic of demographic parity is its strict equality of outcomes. By seeking demographic parity and requiring individuals in both groups receive bene ts at equal probabilities, we may discriminate against a minority group that happens to be to be more quali ed for bene ts than the majority on the average.

Equalized odds. The equalized odds metric is based on two related but distinct criteria. One is that the fraction of quali ed minority persons selected is the same as the fraction of quali ed majority persons selected (Hardt et al. 2016). The other is that the fraction quali ed minority persons selected is the same as the fraction of unquali ed majority persons selected (Zafar et al. 2017). The former is also known asquality of opportunity and is de ned by comparing the ratio (TP+ FN):

$$W(\mathbf{d}) = \frac{\dot{a}_{i2N_{1}} d_{i} d_{i}}{\dot{a}_{i2N_{1}} d_{i}} - \frac{\dot{a}_{i2N_{0}} d_{i} d_{i}}{\dot{a}_{i2N_{0}} d_{i}};$$

The latter criterion is based on the ratio={PP+ TN):

$$W(\mathbf{d}) = \frac{\dot{a}_{i2N_1}(1 \quad d_i)d_i}{\dot{a}_{i2N_1}(1 \quad d_i)} - \frac{\dot{a}_{i2N_0}(1 \quad d_i)d_i}{\dot{a}_{i2N_0}(1 \quad d_i)}:$$

Accuracy parity. The two-sided evaluation in equalized odds can be obviated simply by measuring the fraction of predictions that are accurate, which is the (ate + TN)=(TP + TN + FP + FN), namely,

$$W(\mathbf{d}) = \frac{1}{jN_1j} \mathop{a}\limits_{i2N_1}^{a} d_i d_i + (1 d_i)(1 d_i) \frac{1}{jN_0j} \mathop{a}\limits_{i2N_0}^{a} d_i d_i + (1 d_i)(1 d_i) =$$

Predictive rate parity. When one wishes to compare what fraction individuals selected from each group should have been selected, the relevant mea**puzedist**ive rate parity This measure aims to equalize T=(TP+FP), namely,

$$W(\mathbf{d}) = \frac{\dot{a}_{i2N_1} d_i d_i}{\dot{a}_{i2N_1} d_i} - \frac{\dot{a}_{i2N_0} d_i d_i}{\dot{a}_{i2N_0} d_i}.$$

Predictive parity is primarily considered in risk assessment contexts, such as, recidivism prediction (Dieterich et al. 2016, Chouldechova 2017), child maltreatment screening (Chouldechova et al. 2018).

### 1.2 Outline and Contribution

Chapter 2 studies the trade-off between fairness and ef ciency, an important task in many practical decisions. We propose a principled and practical method for balancing these two criteria in an optimization model. Following an assessment of existing schemes, we de ne a set of social welfare functions (SWFs) that combine Rawlsian leximax fairness and utilitarianism and overcome some of the weaknesses of previous approaches. In particular, we regulate the equity/ef ciency trade-off with a single parameter that has a meaningful interpretation in practical contexts. We formulate the SWFs using mixed integer constraints and sequentially maximize them subject to constraints that de ne the problem at hand. We demonstrate the method on problems of realistic size involving healthcare resource allocation and disaster preparation, with solution times of several seconds at most.

Chapter 3 examines an optimization problem rising from training fair support vector machine (SVM). SVM is a popular supervised learning model that constructs hyperplane(s) for tasks including classi cation and regression. Motivated by the computational advantages of specialized algorithms over general-purpose optimization solvers in training SVM, we propose specialized algorithms to train SVMs with fairness constraints. We focus on fairness constraints that are linear functions of the weight vector of the separating hyperplane in a SVM. Utilizing the structure of the dual formulation of this particular form of fair SVM, we design a coordinate descent type subroutine to extend two dual based standard SVM training algorithms, the dual coordinate descent (DCD) algorithm for SVM with a linear kernel and the sequential minimal optimization (SMO) algorithm for SVM with an arbitrary kernel, to train fair SVMs. We establish convergence properties of these new algorithms. In numerical experiments, our training algorithms train fair SVMs much more efficiently than an off-the-shelf quadratic program solver. Moreover, compared to training standard SVMs, our algorithms allow fairness constraints to be included with minor runtime increase.

Chapter 4 proposes optimization as a general paradigm for formalizing welfare-based fairness in AI systems. We argue that optimization models allow formulation of a wide range of fairness criteria as social welfare functions, while enabling AI to take advantage of highly advanced solution technology. We highlight that social welfare optimization supports a broad perspective on fairness motivated by distributive justice considerations, as literature has designed social welfare functions capturing various concepts of equity. We formalize in-processing and post-processing integration schemes between social welfare optimization and AI, in particular machine learning and rule-based AI. We then present empirical results from a simulated loan processing example to demonstrate the viability and potentials of the integration strategies. Our optimization-centric integrated decision framework broadens the research scope of welfare-based fairness, and opens up interesting directions to pursue a more comprehensive view of fair decision-making.

Chapter 5 studies the problem of online learning (OL) from revealed preferences: a learner wishes to learn a non-strategic agent's private utility function through observing the agent's utility-maximizing actions in a changing environment. We adopt an online inverse optimization setup, where the learner observes a stream of agent's actions in an online fashion and the learning performance

is measured by regret associated with a loss function. We design a new loss function that is convex under relatively mild assumptions. Through establishing that the regret with respect to the new loss bounds the regret with respect to the three classical loss functions used in literature, we provide a exible OL framework that enables a uni ed treatment of loss functions and supports a variety of online convex optimization algorithms. We demonstrate with theoretical and empirical evidence that our framework based on the new loss function (coupled in particular with the online Mirror Descent) has signi cant advantages in terms of regret performance and solution time over other OL algorithms from the literature and bypasses the previous technical assumptions as well.

Chapter 6 investigates the modeling and inferring of people's moral preferences. We consider a setting in which a social planner or policymaker has to make a sequence of decisions regarding the allocation of scarce resources in high-stakes social domains. Our goal is to understand stakeholders' moral judgments regarding such allocation policies. In particular, we evaluate the sensitivity of these judgments to the context/history of allocations and their perceived future impact on various socially salient groups. We propose a mathematical model to capture and infer stakeholders' potentially-dynamic moral preferences. We illustrate our model through small-scale human-subject experiments, which elicit crowd workers' moral judgments regarding scarce medical resource distributions during a hypothetical viral epidemic. We observe that participants' preferences are indeed history- and impact-dependent. Additionally, our preliminary experimental results reveal intriguing patterns speci c to medical resources—a topic that is especially salient in the backdrop of the global covid-19 pandemic.

# Chapter 2

# Combining Leximax Fairness and Ef ciency in a Mathematical Programming Model

# 2.1 Introduction

<sup>1</sup> Fairness is an important consideration across a wide range of optimization models. It can be a central issue in health care provision, disaster planning, workload allocation, public facility location, telecommunication network management, traf c signal timing, and many other contexts. While it is normally straightforward to formulate an objective function that re ects efficiency or cost, it is not obvious how to express fairness in mathematical form. When both fairness and efficiency are desired, as is typical in practice, there is the additional challenge of mathematically integrating them in a tractable model.

For example, when a natural disaster brings down the electric power grid, crisis managers may dispatch crews to urban areas rst in order to restore power to more households quickly, thus maximizing ef ciency. Yet this may cause rural areas to experience very long blackouts, which could be seen as unfair. A more satisfactory solution might give some amount of priority to rural customers, but without imposing too much harm on the population as a whole. Similarly, traf c signal timing that minimizes total delay may result in impracticably long wait times for traf c on minor streets that cross a main thoroughfare. Again a balance between equity and ef ciency may be desirable. The issue can be especially acute in health care. Expensive treatments or research programs that prolong the life of a relatively few gravely ill patients may divert funds from preventive health measures that would spare thousands the suffering brought by less serious diseases.

In this chapter, we develop a practical and yet principled approach to balancing ef ciency and fairness that can be implemented with mixiet/get/linear programming (MILP) models. While

<sup>&</sup>lt;sup>1</sup>This chapter is based on joint work with John Hooker, and has been published in Chen and Hooker (2020, 2022a).

there are many possible measures of fairness, we choose a criterion based ultimately on John Rawls' concept of justice-as-fairness (Rawls 1999). One consequence of the Rawlsian analysis is his famous difference principle, which states roughly that a fair distribution of resources is one that maximizes the welfare of the worst-off. Rawls defends the principle with a social contract argument that can be plausibly extended to lexicographic maximization. That is, the welfare of the worst-off is rst maximized subject to resource constraints, then the second worst-off, and so forth. The Rawlsian perspective has been defended by closely reasoned philosophical arguments in a vast literature (Richardson and Weithman 1999, Freeman 2003).

The Rawlsian argument goes roughly as follows. Let's suppose that all concerned parties adopt an agreed-upon social policy in an original position behind a "veil of ignorance" as to their identity. It must be a policy that all parties can rationally accept upon learning who they are. Rawls argues that no rational decision maker will accept a policy in which she is the least advantaged, unless she would have been even worse off under any other policy. A fair outcome should therefore maximize the welfare of the worst-off. The argument can be employed recursively to defend a leximax criterion. Rawls intended his principle to apply only to the design of social institutions, and to pertain only to the distribution of "primary goods," which are goods that any rational person would want. Yet it can be plausibly extended to distributive justice in general, particularly if it is appropriately combined with an ef ciency criterion.

A fundamental question that arises in the integration of equity and ef ciency is how to regulate the trade-off between the two. We nd in a survey of existing models that it is rarely clear how trade-off parameters can be selected and interpreted in a practical context. However, the modeling scheme of Hooker and Williams (2012) offers a potentially appealing approach to this problem. It governs the trade-off between a Rawlsian maximin and a utilitarian criterion with a single para **Drietaer** has the same units as utility and can be related naturally to the problem at hand. The **Vaise** of chosen so that parties whose utility is with **Dro**f the lowest are seen as suf ciently disadvantaged to deserve priority. Larger values **D** fresult in greater equity. The model also has a practical mixed integer/linear programming (MILP) formulation.

The Hooker–Williams (H–W) scheme has a serious limitation, however. Because its fairness component is the maximin criterion, the actual utility levels of disadvantaged parties other than the very worst-off have no bearing on social welfare. As a result, the solution can be insensitive to most equity considerations. This outcome is particularly unsatisfactory when resource limitations tightly constrain the bene ts available to a few parties, a situation we have found to be common in practice. The H–W model awards what utility it can to the most highly-constrained party, whereupon the welfare of other disadvantaged parties becomes irrelevant, and all solutions become virtually indistinguishable with respect to equity. The fairness criterion plays essentially no role in the determination of an optimum among a potentially large number of outcomes considered equally desirable by the H–W scheme, even for arbitrarily large values Dof

### 2.1.1 Our Approach and Results

A natural way to address this problem is to combine ef ciency with a lexicographic criterion rather than a maximin criterion. This allows the utility levels of all disadvantaged parties to factor into social welfare. However, it poses a dif cult modeling challenge at both a theoretical and a computational level. We meet the challenge by maximizing a sequence of social welfare functions that, except for the rst, are quite different from the single function used in the H–W model. Nonetheless the parameter D has a similar interpretation, with increasibg corresponding to greater prioritization of the fairness criterion.

We show how to formulate these optimization problems as a set of practical MILP models. These MILP problems have substantially different constraint sets and polyhedral properties than the H–W formulation. We also describe a family of valid inequalities that can be added to tighten the models. We extend these results to the common situation in which utility is distributed to groups rather than individuals, such as organizations, regions, or demographic groups. We conclude this chapter by demonstrating the practical applicability of our approach on a healthcare resource allocation problem and an emergency preparedness problems. The former allows us to compare results with those reported by Hooker and Williams (2012) on the same problem. The latter is a shelter location and assignment problem of realistic size. We nd that our approach yields reasonable and nuanced socially optimal solutions for both problems, with computation times ranging from a fraction of a second to 18 seconds for a give**D**.

### 2.1.2 Basic De nitions

An optimization model for integrating fairness and ef ciency can be viewed as maximizingial welfare function(SWF)F( $\mathbf{u}$ ). The value of the function is interpreted as measuring the desirability of a given distribution $\mathbf{u} = (u_1; \ldots; u_n)$  of utilities, where  $u_i$  is the amount of utility allocated to party. We assume the optimization model has the general form

$$\max_{\mathbf{u},\mathbf{0}} \mathbf{F}(\mathbf{u}) \quad \mathbf{u} \ge \mathbf{U} \tag{2.1}$$

where U is the set of feasible utility vectors. In practice, is de ned by a constraint set that generally contains additional variables. We illustrate realistic models of this kind in Section 6.3.

We de ne aleximax solution of (2.1) with respect to a nondecreasing ordering of utilities rather than a predetermined ordering, since the former is relevant to the Rawlsian criterion. Thus given two feasible solution  $\mathbf{s}_i$ ,  $\mathbf{u}^0 \ge U$  of (2.1), we let  $i_1$ ;...; $i_n$  be any permutation of;...;n for which  $u_{i_1}$ ,  $u_{i_n}$ , and  $k_1$ ;...; $k_n$  any permutation for which  $\mathbf{u}_{k_1}^0$ ,  $\mathbf{u}_{k_n}^0$ . Then  $\mathbf{u} \ge U$  is a leximax solution of (2.1) if for any  $\mathbf{u}^0 \ge U$  and any 2 f 1;...;ng such that  $\mathbf{u}_{i_j} = \mathbf{u}_{k_j}^0$  for j = 1;...; 1, we have  $\mathbf{u}_{k_n}^{c}$ .

Two properties that SWFs can possess will be helpful for assessing the equity measures discussed in the next section. The igou–Dalton(P–G) condition is frequently used to assess social welfare

functions, particularly those that measure equality (Dalton 1920, Moulin 2004). It is satis ed when any utility transfer from a better-off party to a worse-off party increases (or does not reduce) social welfare. Given a utility vector with  $u_i < u_j$ , we say that  $\mathcal{P}$ -G transferis a transfere > 0 of utility from  $u_j$  to  $u_i$  such that  $u_i + e = u_j = e$ . Then  $F(\mathbf{u})$  satis es the P-G condition if for any, and any  $\mathbf{u}^0$ that results from a P-G transfer, we have  $\mathbf{u}_i = F(\mathbf{u})$ .

The Chateauneuf-Moye( $\mathfrak{C}$ -M) condition is a slightly weaker condition that considers certain utility transfers from a better-off class to a worse-off class rather than from one individual to another. We will refer to these a $\mathfrak{C}$ -M transfers Chateauneuf and Moyes (2005) suggest that their condition is preferable to th $\mathfrak{P}$ -G condition for assessing equity measures. Their simplest argument is that while a pairwise P-G transfer reduces inequality between two individuals, it may increase inequality between those individuals and others. A C-M transfer does not incur this problem, because the donor and recipient classes respectively lie completely above and below the rest of the population.

To de ne the C–M condition formally, we say that a C–M transfer derive from **u** when  $u_1 u_n$  as well as  $u_1^0$   $u_n^0$ , and for some pair of integers h with 1 ` < h n, we have  $u_1 < u_h$  and

$$\mathbf{u}^0 = \mathbf{u} + \frac{\mathbf{e}}{\mathbf{e}} \overset{\mathbf{e}}{\underset{i=1}{\mathbf{a}}} \mathbf{e}_i \quad \frac{\mathbf{e}}{\mathbf{n} + 1} \overset{\mathbf{n}}{\underset{i=h}{\mathbf{a}}} \mathbf{e}_i$$

for some e > 0, where  $e_i$  is the *i*th unit vector. Then  $F(\mathbf{u})$  satisfies the C-M condition i  $F(\mathbf{u}^0) = F(\mathbf{u})$  for any C-M transfer that derives from  $\mathbf{u}$ . Any function that satisfies the P-G condition also satisfies the C-M condition.

### 2.1.3 Related Literature

We now survey some of the primary schemes that have been proposed for combining fairness and ef ciency. Each can be formulated as a social welfare function that can be maximized in an optimization model of the form (2.1).

### **Convex Combinations**

The most obvious device for combining fairness and ef ciency is a convex combination of the two. This corresponds to a SWF of the form

$$F(\mathbf{u}) = (1 \mid i) \underset{i}{a} u_i + i F(\mathbf{u})$$
(2.2)

where F (**u**) is a fairness measure. A number of function (**u**) have been proposed, such as inequality metrics, the Rawlsian maximin principle, and leximax fairness (Cowell 2000, Jenkins and Van Kerm 2011, Karsu and Morton 2015).

A perennial problem with convex combinations is that it is dif cult to interpreparticularly since  $F(\mathbf{u})$  is typically measured in units other than utility. For example, if we select the widely-used Gini coef cient  $G(\mathbf{u})$  as a measure of equity, then we must combine utility with a dimensionless

quantity  $F(\mathbf{u}) = 1$  G( $\mathbf{u}$ ), where

$$G(\mathbf{u}) = \frac{\overset{\bullet}{a}_{i < j} u_{i} u_{j}}{\underset{i}{n \overset{\bullet}{a}_{i} u_{i}}}$$
(2.3)

Another dif culty is that fairness measures are almost always nonlinear, which can pose tractability problems.

Eisenhandler and Tzur (2019) use a product rather than a convex combination of utility and  $G(\mathbf{u})$ , which reduces to an SWF that is easily linearized:

$$F(\mathbf{u}) = \mathop{\text{a}}_{i} u_{i} \quad \frac{1}{n} \mathop{\text{a}}_{i < j}^{i} u_{j} \quad u_{i}j$$

Yet we now have a convex combination of total utility and another equality metric (negative mean absolute difference) in which =  $\frac{1}{2}$ . One may ask why this particular valuelofis suitable.

Mostajabdaveh et al. (2019) use a linear combination that is equival  $a_n u_i = m(1 - G(u)) a_i u_i$ , where m2 [0; 1]. This at least combines quantities measured in the same units. Yet we again have the problem of justifying a weight. In fact, this combination is equivalent to the convex combination implied by the Eisenhandler and Tzur criterion, except that m=(1 + 2m) rather than  $\frac{1}{2}$ .

Since equality is often unsuitable as a fairness measure (Frankfurt 2015, Scanlon 2003), one may wish to use the Rawlsian criterion  $(\mathbf{u}) = \min_i f u_i g$ . It results in a convex combination of quantities that are measured in the same units, but it is again unclear how to select a suitable value tef that if we index utilities so that  $\mathbf{u}_1$  un, the convex combination becomes simply a weighted sum  $\mathbf{u}_1 + (1 \quad \mathbf{I})_{a_{i>1}}^{a} \mathbf{u}_i$  that gives somewhat more weight to the lowest utility. It is unclear how much more weight is appropriate.

One might also attempt to formulate a convex combination of efficiency with a leximax rather than a maximin criterion. Yet it is unclear how to capture leximax in a function when the utilities cannot be ordered by size in advance. Ogryczak and Sliwinski (2006) show how to formulate leximax in an optimization model without pre-ordering, but this requires coefficients that vary enormously in size and can introduce numerical instability. There is also no evident means for incorporating an efficiency criterion into the model.

### Alpha Fairness and Kalai-Smorodinsky Fairness

Alpha fairness is a parameterized combination of equity and efficiency that does not rely on a convex combination. It is based on an SWF of the form

$$F_{a}(\mathbf{u}) = \overset{8}{\underset{i}{\gtrless}} \frac{1}{1 a} \overset{1}{\underset{i}{a}} u_{i}^{1 a} \text{ for } a \quad 0; a \in 1$$
$$\overset{8}{\underset{i}{\gtrless}} \underset{i}{\underset{i}{a}} \log(u_{i}) \text{ for } a = 1$$

This SWF satis es the P–G and therefore the C–M condition for all 0. Larger values of imply a greater emphasis on equity, with= 0 corresponding to a pure utilitarian criterion and ¥ to a maximin criterion. Lan et al. (2010) provide an axiomatic treatment, and Bertsimas et al. (2012) study worst-case equity/ef ciency trade-offs. An interpretation of that utility  $u_j$  must be reduced by  $(u_j=u_j)^a$  units to compensate for a unit increase  $u_i n(< u_j)$  while maintaining constant social welfare. Yet it is again unclear what kind of trade-off, and therefore what value is fappropriate for a given application. There is also the computational issue  $P_i n(a_i)$  is nonlinear.

Another issue with alpha fairness is that it can assign equality the same social welfare as arbitrarily extreme inequality. In a 2-player situation, for example, the distribution (s; s) has the same social welfare value a\$t;T), where

 $t = \begin{pmatrix} s^2 = T & \text{if } a = 1 \\ 2s^{1 a} & T^{1 a} & T^{1 a} & \text{if } a > 1 \text{ and } 2s^{1 a} > T^{1 a} \end{pmatrix}$ 

Thus if we hold social welfare xed, we have 0 asT ! ! for a = 1, and  $! 2^{1=(1 a)} \text{sasT}$  ! ! for a > 1. This means that when 1, alpha fairness can judge an egalitarian solution to be no better than a solution in which one party has arbitrarily more wealth than the other. The same is true of the maximin criterion, of course, since it assigns) and (s; T) the same social value for arbitrarily largeT.

A well-known special case of fairness arises when = 1. This results in proportional fairness, which is equivalent to the Nash bargaining solution (Nash 1950). Nash (1950) showed that his bargaining solution for two persons is implied by a set of axioms for utility theory, including a strong and perhaps questionable axiom of cardinal noncomparability across parties (Hooker 2013). Harsanyi (1977), Rubinstein (1982), and Binmore et al. (1986) showed that the Nash solution is the asymptotic outcome of certain rational bargaining procedures, again based on strong assumptions.

Kalai and Smorodinsky (1975) proposed an alternative to the Nash bargaining solution that minimizes each player's relative concession. The approach is defended by Thompson (1994) and is consistent with the contractarian ethical philosophy of Gauthier (1983). Mathematically, the objective is to nd the largest scalar such that  $\mathbf{u} = (1 \ b)\mathbf{d} + b\mathbf{u}^{max}$  is a feasible utility vector, where each  $u_i^{max}$  is the maximum of  $u_i$  over all feasible utility vectors and each  $d_i$  is the starting utility of i. The bargaining solution is the vector that maximizes. This can be interpreted geometrically as the furthest feasible point from the origin on the line segment connecting the origin  $\mathbf{M}^{max}$  is equivalent to maximizing the SWF

 $F(\mathbf{u}) = \begin{cases} a_i u_i; & \text{if } \mathbf{u} = (1 \ b)\mathbf{d} + b \mathbf{u}^{\text{max}} \text{ for some b with } 0 \ b \ 1 \\ 0; & \text{otherwise} \end{cases}$ 

#### 2.1 Introduction

The SWF is not only discontinuous build build build build be condition. For example, if we have a 2-person utility distribution  $(u_1; u_2) = (b u_1^{max}; b u_2^{max})$  for some bwith 0 < b 1, then a utility transfer that tends to equalize the distribution reduces social welfare.

The K–S scheme provides no parameter for adjusting the equity-ef ciency trade-off. While it can be suitable for such applications as wage or price negotiation, it can yield solutions in other contexts that many would consider unjust. For example, it can divert treatment resources from cancer patients to persons suffering from the common cold to provide them the same fraction of their maximum health potential.

### **Threshold Models**

Williams and Cookson (2000) proposed a pai2offersorSWFs based on a utility or equity threshold. A utility-threshold model uses the maximin criterion unless the sacri ce in total utility exceeds a threshold, in which case it switches to a utilitarian criterion. An equity-threshold model uses a utilitarian criterion unless inequality becomes excessive, when it switches to maximin. Hooker and Williams (2012) extended the utility-threshold model to **thperson** criterion described earlier, and McElfresh and Dickerson (2018) proposed a similar scheme based on a leximax rather than maximin criterion. Our aim in the present paper is likewise to combine leximax and ef ciency in a threshold model, but we will argue that it offers several advantages relative to the McElfresh and Dickerson approach.

The 2-person SWF implied by Williams and Cookson's utility-threshold model can be formulated

1

$$F(u_1; u_2) = \begin{pmatrix} u_1 + u_2; & \text{if } ju_1 & u_2 j & D \\ 2 \min f u_1; u_2 g + D; & \text{otherwise} \end{pmatrix}$$
(2.4)

The function is utilitarian whej $u_1$   $u_2j$  D and represents a maximin criterion otherwise. Indifference curves (contours) of the SWF are illustrated in Fig. 2.1. The maximin cri**terio** $u_1$ ;  $u_2g$  is modi ed in (2.4) to obtain continuous contours. We will see that maintaining continuity is a major factor in the design of threshold-based SWFs.

The feasible set in Fig. 2.1 is the portion of the nonnegative quadrant under the curve. It represents all feasible utility outcomes that are permitted by the resource budget and other constraints. The shape of the curve indicates that when party 1's utility reaches a certain point, further improvement requires extraordinary sacri ce by party 2 due to the transfer of resources. The utilitarian solution (black dot in the gure) might therefore be viewed as preferable to the maximin solution (small open circle) and in fact yields slightly more social welfare as indicated by the contours.

Hooker and Williams (2012) extend this social welfare function persons as follows:

$$F_{1}(\mathbf{u}) = (n \ 1)D + nu_{h1i} + a_{i=1}^{n} (u_{i} \ u_{h1i} \ D)^{+}$$
(2.5)

Fig. 2.1 Piecewise linear social welfare contours for 2 persons.

where  $(a)^+ = \max 0$ ; a g. Here we adopt the convention  $th(at_{h1i}; ...; u_{mi})$  is the tuple  $(u_1; ...; u_n)$  arranged in non-decreasing order. We refer to the function  $b_1$  as cause it will be the rst in a series of functions  $F_1; ...; F_n$  we de ne later. It may be more intuitive to rewrite (2.5) as

$$F_1(\mathbf{u}) = t(\mathbf{u}) \quad 1 \quad D + \overset{t(u)}{\underset{i=1}{\overset{n}{a}}} u_{h1i} + \overset{n}{\underset{i=t(u)+1}{\overset{n}{a}}} u_{hii}$$

where  $t(\mathbf{u})$  is defined so that  $u_{h1i}$ ;:::; $u_{ht}(\mathbf{u})_i$  are within D of  $u_{h1i}$ ; that is,  $u_{h1i}$  D if and only if i  $t(\mathbf{u})$ . We will refer to utilities  $u_{h1i}$ ;:::; $u_{ht}(\mathbf{u})_i$  as being in the fair region and utilities  $u_{h1i}(\mathbf{u})_{+1i}$ ;:::; $u_{hni}$  as being in the utilitarian region. The function  $F_1(\mathbf{u})$  therefore has the effect of summing all the utilities, but with the proviso that utilities in the fair region are counted as equal to  $u_{h1i}$ . The term( $t(\mathbf{u})$  1) D is added to ensure continuity of the function.

The paramete<sup>D</sup> therefore has an interpretation that can be described independently of its role in the SWF. Namely, any party with utility withi<sup>D</sup> of the lowest is viewed as disadvantaged and deserving of special consideration. The SWF then de nes the special consideration to be an identi cation of the disadvantaged party with the worst-off party, which is given disproportionate weight in the summation of utilities—namely, weight equal to the number of utilities with the lowest.

A problem with (2.5), however, is that the actual utility levels of the disadvantaged parties, other than that of the very worst-off, have no effect on the value of the SWF. This is illustrated in the 3-person example of Fig. 2.2, which shows the contou **F**s( $\alpha$ f; u<sub>2</sub>; u<sub>3</sub>) with D = 3 and u<sub>1</sub> xed to zero. The SWF is constant in the shaded region, meaning that the utilities allocated to persons 2 and 3 have no effect on social welfare as measure **G**( $\alpha$ **y**), so long as they remain in the fair region. As

Fig. 2.2 Contours  $oF_1(0; u_2; u_3)$ . The function is constant in the shaded region.

a result, there are in nitely many utility vectors that maximize social welfare, some of which differ greatly with respect to utilities in the fair region. One can add a tie-breaking  $(u_2n_1, u_3)$  to the social welfare function, where > 0 is small, so as to maximize utility as a secondary objective. Yet this still does not account for equity considerations within the fair region.

To obtain a threshold model that is sensitive to the actual utility levels of all the disadvantaged parties, one might combine utility with a leximax criterion rather than a maximin criterion. McElfresh and Dickerson (2018) propose one method of doing so in the context of kidney exchange. Their method is related to the H–W approach, but it relies on the assumption that the parties can be given a preference ordering in advance. It rst maximizes a SWF that combines utilitarian and maximin criteria in a way that treats the most-preferred party as the worst-off. If all optimal solutions of this problem lie in the utilitarian region, a utilitarian criterion is used to select one of the optimal solutions. (Here, a utility vecto**u** is said to be in the fair region ifhaxif u<sub>i</sub>g min<sub>i</sub>f u<sub>i</sub>g D, and otherwise in the utilitarian region.) Otherwise a leximax criterion is used for all of the optimal solutions, subject to the preference ordering. If we index the parties in order of decreasing preference, the SWF is

$$F(\mathbf{u}) = \begin{cases} \sigma \\ < nu_1; & \text{if } ju_i & u_j j & \text{D for all } i; j \\ \vdots & a_i u_i + sgn(u_1 & u_i)D; & \text{otherwise} \end{cases}$$
(2.6)

McElfresh and Dickerson state that (u) has continuous contours, but this is true only for 2. For a counterexample with = 3, we note that (0; 0; D + e) = e and F(0; e; D + e) = 2e D for arbitrarily smalle > 0. The discontinuity of the SWF raises questions regarding its suitability for application, since a slight change in the utility distribution could bring about a large and unexpected change in the measurement of social welfare. We also not **E** (he) tviolates the C–M condition and therefore the P–D condition when 3. For example, a C–M transfer that  $conve(uts; u_2; u_3)$  from (e; 0; D+ e) to (e; e; D) reduces social welfare from 2-D to e.

This approach has two additional limitations. It is often not possible to pre-specify a preference ranking of parties, as was done in the kidney exchange problem. Another is that the leximax criterion is not used until optimal solutions of the SWF are already obtained, and then applied only to the optimal solutions. We wish to allow the leximax criterion to play a role in evaluating all the possible solutions. These limitations are overcome by our proposal, described in the next section.

An earlier version of our scheme appears in a brief conference paper (Chen and Hooker 2020), which uses somewhat different SWFs.

## 2.2 De ning the Social Welfare Functions

To combine leximax and utilitarian criteria in a threshold model, we propose to maximized ancef social welfare function  $\mathbf{F}_1(\mathbf{u})$ ;:::; $\mathbf{F}_n(\mathbf{u})$ , each of which combines maximin and utilitarian measures. The rst function  $F_1(\mathbf{u})$  is the H–W function (2.5) de ned earlier and is maximized over  $(u_1;:::;u_n)$  to obtain a value for  $\mathbf{u}_{h1i}$ . Each subsequent function  $(\mathbf{u})$  is maximized over  $\mathbf{u}_{hki}$ ;:::; $\mathbf{u}_{hni}$ , while xing utilities  $\mathbf{u}_{h1i}$ ;:::; $\mathbf{u}_{hk}$  is to the values already obtained, and while giving a certain amount of priority. The solution of this maximization problem determines the values  $\mathbf{u}_{hi}$  of

The process terminates when maximiz  $\mathbf{F}_{k}(\mathbf{u})$  yields a value of  $\mathbf{u}_{hki}$  that lies outside the fair region. At this point  $\mathbf{F}_{k}(\mathbf{u})$  is utilitarian, and utilities  $\mathbf{u}_{hki}$ ;...;  $\mathbf{u}_{mi}$  are determined simultaneously by maximizing  $\mathbf{F}_{k}(\mathbf{u})$  while xing  $\mathbf{u}_{h1i}$ ;...;  $\mathbf{u}_{hk-1i}$  to the values already obtained. We refer to a utility vector ( $\mathbf{u}_{h1i}$ ;...;  $\mathbf{u}_{mi}$ ) that results from this process **as**cially optimal

We describe this sequential optimization procedure more precisely in Section 2.4, but we must rst de ne and explain the function  $\mathbf{E}_k(\mathbf{u})$  for k 2. Three main considerations govern the design of these functions and give them a signi cantly different character  $\mathbf{F}_k(\mathbf{u})$ .

- The fair region must be viewed as already de ned, because was xed by maximizing  $F_1(\mathbf{u})$ .
- The utility u<sub>tki</sub> should receive less priority asincreases, since the it becomes less disadvantaged relative to the xed lowest utilityu<sub>thi</sub>.
- It turns out that the priority given<sub>tki</sub> cannot depend on the number of utilities in the fair region, as it does fok = 1, because this results in an irreducibly discontinuous SWF. We therefore designF<sub>k</sub>(u) so that the priority depends only an

To develop SWFs that are somewhat analogous to the H–W funct( $\mathbf{u}$ ) while re ecting these considerations, it is helpful to write ( $\mathbf{u}$ ) as

$$F_1(\mathbf{u}) = t(\mathbf{u})u_{h1i} + t(\mathbf{u}) \quad 1 D + a^n a^n u_{h1i}$$
  
 $i = t(\mathbf{u}) + 1$ 

The function assigns weigh(u) to utility  $u_{h1i}$  and weight 1 to utilities in the utilitarian region. This is algebraically equivalent to the formulation

$$F_1(\mathbf{u}) = nu_{h1i} + (n \ 1)D + \overset{n}{\overset{n}{a}}(u_{hii} \ u_{h1i} \ D)$$

We modify this pattern as follows:

$$F_{k}(\mathbf{u}) = \begin{cases} 8 & k \\ \overset{k}{\underset{i=1}{a}} (n & i+1)u_{hii} + \overset{n}{\underset{i=t(\mathbf{u})+1}{a}} (u_{hii} & u_{h1i} & D); & \text{if } t(\mathbf{u}) & k \\ & & \\$$

This SWF assigns weight k + 1 to utility  $u_{hki}$  and weight 1 to utilities in the utilitarian region. As desired, the priority given  $ta_{hki}$  depends only on and decreases as increases. While k = 1 assigns weight  $ta_{h1i}$  equal to the number of utilities in the fair region, this weight is reduced by one each time increases by one, to re ect the fact that one of these utilities is xed and removed from the optimization problem. We also note that additional multiples afe not required to ensure continuity of  $F_k(\mathbf{u})$  when k = 2.

Thus as the function  $\mathbf{g}_{k}(\mathbf{u})$  are sequentially maximized for increasing value  $\mathbf{k}$ ,  $\mathbf{e}\mathbf{f}$  ach utility  $\mathbf{u}_{tki}$ in the fair region receives priority at some point in the process. This scheme incorporates lexicographic optimization in the sense that the smaller utilities are determined earlier in the sequence, although rather than maximizin  $\mathbf{g}_{tki}$  in stepk, we maximize a SWF that gives priority  $\mathbf{t}_{\mathbf{k}i}$ . Utilitarianism in incorporated because each maximization problem considers total utility as well as fairness.

This process yields a purely utilitarian solution where 0. For in this case we have  $(\mathbf{u}) = 1$  for all  $\mathbf{u}$ , and  $F_1(\mathbf{u})$  reduces to a utilitarian criterion. The fair region is the single point and we solve the social welfare problem simply by maximizing  $\mathbf{F}_{g}(\mathbf{u})$ , which yields a utilitarian solution. At the opposite extreme, where  $\mathbf{v}$  is equivalent to maximizing  $\mathbf{g}_{hki}$ . The reason is that for sufficiently largeD,  $t(\mathbf{u}) = n$  for all feasible  $\mathbf{u}$ , and  $\mathbf{F}_k(\mathbf{u})$  is  $(n + 1)\mathbf{u}_{hki}$  plus a constant for  $k = 1; \dots; n$ . (Recall that  $\mathbf{u}_{h1i}; \dots; \mathbf{u}_{hk-1i}$  are xed when  $\mathbf{F}_k(\mathbf{u})$  is maximized.) Thus by sequentially maximizing  $\mathbf{F}_1(\mathbf{u}); \dots; \mathbf{F}_n(\mathbf{u})$ , we maximize the smallest utility, then maximize the second smallest while holding  $\mathbf{u}_i$  xed, and so forth. When there is a tie for the smallest utility, the resulting solution may not be leximax, depending on how the tie is broken, but one of the solutions generated in this manner will be leximax.

Figure 2.3 illustrates how maximizing  $(\mathbf{u})$ ;:::; $F_n(\mathbf{u})$  sequentially is more sensitive to equity than maximizing  $F_1(\mathbf{u})$ , which has the at region shown in Fig. 2.2, as noted earlier. Suppose we determine a value for 1 by maximizing  $F_1(\mathbf{u})$ , say  $u_1 = 0$ . Then the function  $F_2(\mathbf{u})$  has no at regions, as is evident in Fig. 2.3, and the solutions in the at region of Fig. 2.2 are now distinguished.

Fig. 2.3 Contours o $F_2(0; u_2; u_3)$  with D = 3 and contour interval 1.

# 2.3 Properties of the Social Welfare Functions

We now investigate some mathematical properties of the social welfare fun  $\mathbb{E}_{k}(\mathbf{u})$ s First, we note that the contours in Fig. 2.2 are continuous, and the continui  $\mathbb{E}_{k}(\mathbf{u})$  can be shown in general.

Theorem 1. The functions  $\mathbf{k}(\mathbf{u})$  are continuous for  $\mathbf{k} = 1; \dots; n$ .

Proof. To prove continuity of  $F_1(\mathbf{u})$ , it suffices to show that each term of (2.5) is continuous, because a sum of continuous functions is continuous. The rst term of (2.5) is a constant, and the second term is continuous because order statistics are continuous functions. Each term of of the summation is continuous because it is the maximum of two continuous functions. To show  $\overline{W}_k(\mathbf{h})$  tis continuous for k 2, it is convenient to write (2.7) as

$$F_{k}(\mathbf{u}) = \overset{k}{\overset{n}{a}}_{i=1}^{n} (n \quad i+1)u_{hii} + (n \quad k+1)u_{hki} (n \quad k)(u_{hki} \quad u_{h1i} \quad D)^{+} + \overset{n}{\overset{n}{a}}_{i=k+1}^{n} (u_{hii} \quad u_{h1i} \quad D)^{+}$$

which simpli es to

$$F_{k}(\mathbf{u}) = \mathop{a}\limits_{i=1}^{k} (n \quad i+1)u_{hi} + (n \quad k+1) \min u_{hi} + D; u_{hki}g + \mathop{a}\limits_{i=k}^{n} (u_{hi} \quad u_{h1i} \quad D)^{+}$$
(2.8)

Because order statistics are continuous, and  $u_{h1i}$  are continuous functions of Also minf  $u_{h1i}$  + D;  $u_{h1i}$  g and  $(u_{h1i} \quad u_{h1i} \quad D)^+$  are continuous because they are the minimum or maximum of continuous functions.

The SWFs have a monotonicity property as well. BecaFu (se) can be written

$$F_{1}(\mathbf{u}) = (n \ 1)D + a^{n} \underset{i=1}{a} \max u_{i} \ D; u_{h1i}g$$
(2.9)

the following is evident on inspection of (2.8) and (2.9).

Theorem 2.  $F_1(\mathbf{u})$  is monotone nondecreasing. When is xed,  $F_k(\mathbf{u})$  is monotone nondecreasing for k 2.

We also note that while  $\mathbf{E}_{k}(\mathbf{u})$  can violate the P–G condition for 3, it satis es the C–M condition. A counterexample to the P–G condition is illustrated  $\mathbf{F}_{\mathbf{v}}(\mathbf{u})$  in Fig. 2.2, where the utility-preserving transfer from A to B reduces social welfare as measure  $\mathbf{E}_{\mathbf{v}}(\mathbf{u})$ . However, we have the following result.

Theorem 3. The social welfare function  $\mathbf{S}_{k}(\mathbf{u})$  satisfy the Chateauneuf–Moyes condition for 1;:::;n.

Proof. We wish to show that  $\mathbf{k}_k(\mathbf{u})$  satis es the C–M condition for  $\mathbf{k} = 1; \dots; n$ . Recall that a C–M transfer is a transfer of utility from to  $\mathbf{u}^0$  such that  $\mathbf{u}_1$   $\mathbf{u}_n$  as well as  $\mathbf{u}_1^0$   $\mathbf{u}_n^0$ , and for some pair of integers, h with 1 ` < h n, we have  $\mathbf{u}_1 < \mathbf{u}_h$  and

$$\mathbf{u}^0 = \mathbf{u} + \frac{\mathbf{e}}{\mathbf{e}} \overset{\mathbf{e}}{\underset{i=1}{\overset{\mathbf{e}}{\mathbf{e}}}} \mathbf{e}_i \quad \frac{\mathbf{e}}{\mathbf{n} + 1} \overset{\mathbf{n}}{\underset{i=h}{\overset{\mathbf{n}}{\mathbf{e}}}} \mathbf{e}_i$$

for some > 0. ThenF(u) satis es the C–M condition if

$$\mathsf{F}(\mathbf{u}^0) \quad \mathsf{F}(\mathbf{u}) \tag{2.10}$$

for any C–M transfer from  $\mathbf{u}$  to  $\mathbf{u}^0$ .

We rst prove the theorem  $fd\mathbf{k} = 1$ . We wish to show that (2.10) holds for any C–M transfer from  $\mathbf{u}$  to  $\mathbf{u}^0$ . There are three types of C–M transfer, illustrated in Fig. 2.4: (a)  $\mathbf{t}(\mathbf{u})$ , (b)  $\mathbf{t}(\mathbf{u}) < \mathbf{h}$ , and (c)t( $\mathbf{u}$ ) < ` < h. The resulting utility gain by individuals;:::`, and loss by individuals;:::; n, are indicated in Table 2.1. It is clear on inspection of Fig. 2.4 that the gain is at at a start case, and the loss never more the nThe C-M condition is therefore satis ed.

Fig. 2.4 Illustration of C–M transfers relevant  $\mathbb{F}\varphi(\mathbf{u})$ .

Case	Gain	Loss
(a)	$\frac{t(\mathbf{u})}{\mathbf{v}} \mathbf{e} > \mathbf{e}$	$\frac{n t(\mathbf{u})}{n h+1} e < e$
(b)	$\frac{t(\mathbf{u})}{\mathbf{v}} \mathbf{e} > \mathbf{e}$	е
(C)	е	е

Table 2.1 Verifying the Chateauneuf-Moyes condition Fo(ru)

We now prove the theorem for 2. It is clear that a C–M transfer satis  $\notin$  2.10) when k > t(**u**), because in this case (**u**) is simply utilitarian. We therefore need only consider the six types of C–M transfer illustrated in Fig. 2.5, in which t(**u**).

It is convenient to write  $F_k(\mathbf{u})$  in the following form:

$$F_k(\mathbf{u}) = t(\mathbf{u})u_{h1i} + \overset{k}{\underset{i=2}{a}}(n \quad i+1)u_{hii} + \overset{n}{\underset{i=t(\mathbf{u})+1}{a}}(u_{hii} \quad D)$$

The resulting gain by individuals; ::: `, and loss by individuals; :::; n, are indicated in Table 2.2. In cases (b)–(f), it is clear on inspection of Fig. 2.5 that the gain is more dimage ach case, and the loss never more that In case (a), we note rst that the gain can be written

$$n \frac{1}{2} \frac{n t(\mathbf{u})}{2}$$

To show that the loss is no greater than the gain, it sufces to show this lwhen+1, since +1 and the loss is nonincreasing with respectit. To hus it sufces to show

n 
$$\frac{1}{2}$$
  $\frac{n t(\mathbf{u})}{2}$   $\frac{1}{n}$   $\frac{k}{a}$   $(n i+1)+n t(\mathbf{u})$ 

Fig. 2.5 Illustration of C–M transfers relevant  $\mathbf{F}_{Q}(\mathbf{u})$ , k 2.

Case	Gain	Loss
(a)	$\frac{1}{t}$ t( <b>u</b> ) + $a_{i=2}^{n}$ (n i + 1) e	$\frac{1}{n + 1} \stackrel{k}{a}_{i=h}^{k} (n + 1) + n + t(\mathbf{u}) e$
(b)	$\frac{1}{1} t(\mathbf{u}) + a_{i=2}(n i+1) e \frac{t(\mathbf{u})}{1} e > e$	$\frac{n t(\mathbf{u})}{n h+1} \mathbf{e} < \mathbf{e}$
(c)	$\frac{1}{1} t(\mathbf{u}) + \overset{k}{\underset{i=2}{a}} (n  i+1) e  \frac{t(\mathbf{u})}{1} e > e$	$\frac{n t(\mathbf{u})}{n h+1} \mathbf{e} < \mathbf{e}$
(d)	$\frac{1}{2}$ t( <b>u</b> ) + $a_{i=2}^{\circ}$ (n i + 1) e $\frac{t(u)}{2}$ e > e	$\frac{n + 1}{n + 1}e = e$
(e)	$\frac{1}{2} t(\mathbf{u}) + \overset{k}{\underset{i=2}{a}} (n  i+1) e  \frac{t(\mathbf{u})}{2} e > e$	$\frac{n + 1}{n + 1}e = e$
(f)	$\frac{1}{\cdot} t(\mathbf{u}) + \overset{k}{\overset{a}{a}} (n  i+1) + \cdot t(\mathbf{u})  e  e$	$\frac{n + 1}{n + 1}e = e$

Table 2.2 Verifying the Chateauneuf-Moyes condition **F**<sub>Q</sub>(**u**), k 2

Sincek t(u) and each term of the summation is at most`, it sufces to show

$$n \quad \frac{1}{2} \quad \frac{n \quad t(\mathbf{u})}{2} \quad \frac{t(\mathbf{u}) \quad (n \quad ) + n \quad t(\mathbf{u})}{n}$$

Rearranging, we obtain

n t(**u**) 
$$\frac{1}{2} + \frac{1}{n}$$
 1  $\frac{1}{2}$  (2.11)

This inequality is clearly satis ed when the following is false:

$$\frac{1}{n} + \frac{1}{n}$$
 (2.12)

We therefore assume (2.12) is true. Since (2.11) is clearly satis ed  $\dot{w}$  we suppose 2, in which case (2.12) implies <  $^2=(1, 1)$ . Since < h n, we can state

$$+1 n < \frac{2}{1}$$

or  $^{2}$  1 n( 1) <  $^{2}$ . Sincen and are positive integers, this implies + 1, in which case (2.11) reduces to

$$\frac{1}{2} + 1 \quad t(\mathbf{u}) \qquad \frac{1}{2}$$

This holds because  $(\mathbf{u}) \rightarrow +1$ , and the theorem follows.

Finally, we note that unlike maximin and alpha fairness, the functions for the region of the region of the same social welfare value. For a given egalitarian distribution  $\hat{\mathbf{u}} = (s, \dots, s)$ , we have  $F_1(\hat{\mathbf{u}}) = ns+(n-1)D$  and  $F_k(\hat{\mathbf{u}}) = (n-k+1)s$  for k-2. We obtain the same social welfare value, for alk, by setting one utility in  $\mathbf{u}$  to (n-k+1)s+D and the othen 1 utilities to zero. It follows that the maximum gap between utilities k+1 and k+1, rather than arbitrarily large.

## 2.4 The Sequential Optimization Procedure

The sequential optimization procedure combines leximax and utilitarian criteria by rst giving priority to the worst-off party, then somewhat less priority to the second worst-off party, and so forth (the leximax element), but while never demanding excessive sacri ce from the well-off (the utilitarian element). This is accomplished by sequentially maximizing the SFV(Fu);  $F_2(u)$ ;...: to determine the utility level of the worst-off, second worst-off, etc., because the SWFs are designed to adjust the priorities in this manner. The process continues until the utilities of all disadvantaged parties are xed, whereupon the remaining utilities are determined in a purely utilitarian fashion. The disadvantaged parties are those with utilities in the fair region, which the user de nes by setting the parameter

We rst simplify notation by removing the initial constants  $\operatorname{frd}_{FR}(\mathbf{u})$  for k 2, resulting in the SWF

$$\bar{F}_{k}(\mathbf{u}) = (n \ k+1)u_{hki} + a_{i=k}^{n}(u_{hii} \ u_{h1i} \ D)^{+}$$
 (2.13)

This obviously has no effect on the optimal solution that results from maximizing the SWF. For convenience, we de  $n\overline{\mathbf{E}}_{1}(\mathbf{u}) = F_{1}(\mathbf{u})$ .

We next maximize the social welfare function  $\bar{h}_{\vec{n}}(\mathbf{u})$ ;:::; $\bar{F}_n(\mathbf{u})$  sequentially, subject to resource constraints and the condition that the un xed utilities should be no smaller than the largest utility already xed. The un xed<sub>i</sub> with the smallest value in the solution becomes the utility determined by maximizing  $\bar{F}_k(\mathbf{u})$ .

We indicate resource limits by writing 2 U. In practice, they would be formulated in a MILP model by introducing variables and constraints that specify resource limitations and how resource allocations to individual parties translate to utilities. This will be illustrated in our experiments in Section 6.3.

To state the optimization procedure more precisely, we recursively de ne a sequence of maximization problem  $P_1$ ;...;  $P_n$ , where  $P_1$  maximizes  $F_1(\mathbf{u})$  subject to  $\mathbf{u} \ge U$ , and  $P_k$  for k = 2;...; n
is

$$\max F_{k}(\mathbf{u})$$
s.t.  $u_{i} \quad \bar{u}_{i_{k-1}}; i \ge I_{k}$ 
 $u_{i_{j}} = \bar{u}_{i_{j}}; j = 1; \dots; k \quad 1$ 
 $\mathbf{u} \ge U$ 

$$(2.14)$$

The indices i are de ned so that i is the utility determined by solving. Thus

$$i_j = \underset{i \ge I_j}{\operatorname{argminf}} u_i^{[j]} g$$

where  $\mathbf{u}^{[j]}$  is an optimal solution o $\mathbf{P}_j$  and  $\mathbf{I}_j = f 1; \dots; ng n f i_1; \dots; i_{j-1}g$ . We denote by  $\bar{\mathbf{u}}_{i_j} = \mathbf{u}^{[j]}_{i_j}$  the solution value obtained for  $\mathbf{i}_{i_j}$  in  $\mathbf{P}_j$ . We need only solv  $\mathbf{e}_k$  for  $k = 1; \dots; K + 1$ , where K is the largest k for which  $\bar{\mathbf{u}}_{i_k} = \bar{\mathbf{u}}_{i_1} + D$ . The solution of the social welfare problem is then

$$u_{i} = \begin{matrix} ( & \bar{u}_{i} & \text{for } i = i_{1}; \dots; i_{K-1} \\ & u_{i}^{[K]} & \text{for } i \ge I_{K} \end{matrix}$$

As is frequently the case with optimization models, alternate optimal solutions may exist. In particular, there may be multiple utilities with the same minimum value in the solution  $\mathbf{u}^{[k]}$  of a given problem  $\mathbf{P}_k$ . Any of these utilities may be xed when solvin  $\mathbf{g}_{+1}$ , possibly (but not necessarily) giving rise to multiple socially optimal solutions. This is illustrated in the example of the next section.

At least one optimal solution of each problem is Pareto optimal over the set, due to the monotonicity of each  $F_k(\mathbf{u})$  (Theorem 2). Yet the nal socially optimal solution is not guaranteed to be Pareto optimal for an arbitrary feasible set. For example **3** if **a** rsorproblem has only two feasible solution  $\mathbf{su} = (1;2;2); (1;3;2)$ , then both solutions are socially optimal where 3, but only one is Pareto optimal. However, we will nd that the great majority of socially optimal solutions are Pareto optimal in our experiments. Furthermore, a simple postprocessing step checks whether a socially optimal solution  $\mathbf{u}$  is Pareto optimal and, if not, converts it to a Pareto optimal solution. We need only solve the optimization problem

$$\max a_{i=1}^{n} u_{i}$$
s.t.  $u_{i} \quad u_{i}$ ;  $i = 1; ...; n$ 
 $u \ge U$ 
(2.15)

If the optimal solution  $\hat{\mathbf{u}}$  has  $\hat{\mathbf{u}}_i > \mathbf{u}_i$  for some i, then  $\mathbf{u}_i$  is not Pareto optimal, but is. We nd in the experiments that the adjustment to achieve Pareto optimality is slight even when it is necessary.

# 2.5 A Small Example

We illustrate the sequential optimization procedure on an example3withities,  $\mathbf{u} = (u_1; u_2; u_3)$ . Suppose there are exactive feasible solutions, shown in the rst column of Table 2.3. The remainder of the table displays SWF values  $\mathbf{F}_{2}(\mathbf{u}); \mathbf{F}_{2}(\mathbf{u}); \mathbf{F}_{3}(\mathbf{u})$  for variousD settings. Only  $\mathbf{F}_{1}(\mathbf{u})$  is shown for D = 0, since only P<sub>1</sub> is solved in this case. Recall that for 2,  $\mathbf{F}_{k}(\mathbf{u})$  is modified from  $\mathbf{F}_{k}(\mathbf{u})$  by removing constant terms.

Table 2.3 Values  $o\overline{F}_1(\mathbf{u}); \overline{F}_2(\mathbf{u}); \overline{F}_3(\mathbf{u})$  for a small example. Only  $\overline{F}_1(\mathbf{u})$  is shown forD = 0.

u	D= 0	D= 2	D= 5
<b>u</b> <sup>1</sup> = ( 4; 6; 6)	16	16,12,6	22,12,6
<b>u</b> <sup>2</sup> = ( 2;6;9)	17	17,19,14	18,14,11
<b>u</b> <sup>3</sup> = (1;1;14)	16	18,13,25	21,10,22
<b>u</b> <sup>4</sup> = (1;2;13)	16	17,14,23	20,11,20
<b>u</b> <sup>5</sup> = ( 2; 1; 13)	16	17,14,23	20,11,20

Socially optimal solutions are found as follows.

- D = 0: ProblemP<sub>1</sub> is purely utilitarian, and u<sup>2</sup> maximizes F<sub>1</sub>(u). The process stops with socially optimal solution u = u<sup>[1]</sup> = u<sup>2</sup> without solving P<sub>2</sub> and P<sub>3</sub>, becaus e<sup>[1]</sup><sub>i1</sub> = u<sup>[1]</sup><sub>1</sub> = 2 is in the utilitarian region.
- D= 2: ProblemP<sub>1</sub> has optimal solution  $\mathbf{u}^{[1]} = \mathbf{u}^3$  becaus  $\mathbf{u}^3$  maximizes  $\mathbf{F}_1(\mathbf{u})$ . We can select  $\mathbf{u}_1$  or  $\mathbf{u}_2$  as the utility  $\mathbf{u}_{i_2}$  xed by this solution, since both are minimum  $\mathbf{u}^3$ . If we select  $\mathbf{u}_1$ , then the feasible set for  $\mathbf{f}_2$  is f  $\mathbf{u}^3$ ;  $\mathbf{u}^4$ g, and  $\mathbf{u}^{[2]} = \mathbf{u}^4$  maximizes  $\mathbf{F}_2(\mathbf{u})$  over this feasible set. This leaves  $\mathbf{u}^4$  as the only feasible solution  $\mathbf{e}_3$ , and the socially optimal solution  $\mathbf{u}_3 = \mathbf{u}^4$ . If we select  $\mathbf{u}_2$  as the xed utility, P<sub>2</sub> has the feasible set  $\mathbf{u}^3$ ;  $\mathbf{u}^5$ g and optimal solution  $\mathbf{u}^{[2]} = \mathbf{u}^5$ , resulting in socially optimal solution  $\mathbf{u} = \mathbf{u}^5$ . Thus  $\mathbf{u}^4$  and  $\mathbf{u}^5$  are both socially optimal.
- D = 5: ProblemP<sub>1</sub> has optimal solution<sup>[3]</sup> = u<sup>1</sup>. Hereu<sup>[1]</sup><sub>i1</sub> = u<sup>[1]</sup><sub>1</sub> = 4, and the only feasible solution of P<sub>2</sub> and P<sub>3</sub> is u<sup>1</sup>, which is socially optimal.

We highlight some key observations from this example. The procedure sa utilitarian solution when D = 0 and a leximax solution when D = 5, while the intermediate value = 2 yields a solution that is neither purely utilitarian nor purely leximax. Unlike the H–W criterion, it is sensitive to the utility of a disadvantaged player other than the very worst-off. The H–W criterion prefers solution  $\mathbf{u}^3$  to solutions  $\mathbf{u}^4$  and  $\mathbf{u}^5$  when D = 2, while our method has the opposite preference because it considers the lot of the second-worst-off player. We also prefer solutificates du<sup>5</sup> to solution  $\mathbf{u}^1$ when D = 2 because of the large sacri  $\mathbf{u}^4$  imposes on player 3. When  $\mathbf{b}$  is increased to 5, however, the improved lot of players 1 and 2  $\mathbf{u}^1$  outweighs this sacri ce due to the greater importance of equity.

# 2.6 Mixed Integer Programming Model

For practical solution of the optimization problems we wish to formulate them as MILP models. We drop the resource constraints U from problems  $P_1; \ldots; P_n$  to obtain  $P_1^0, \ldots; P_n^0$ , because we wish to analyze the MILP formulations of the SWFs without the complicating factor of resource constraints. These constraints can later be added to the optimization models before they are solved. In addition, problems  $P_1^0, \ldots; P_n^0$  contain innocuous auxiliary constraints that make the problems MILP representable.

The MILP model for  $P_1^0$  follows a different pattern than the models  $R_2^0$ ; ...;  $P_n^0$ , and we therefore treat the two cases separately. Prob

Constraints (c) ensure MILP representability, as explained in Hooker and Williams (2012), because they imply that the hypograph of (2.16) is a nite union of polyhedra having the same recession cone (Jeroslow 1987). The constraints have no practical import for sufficiently Margeethough for theoretical purposes we assume old by D.

The MILP model for  $P_1^0$  can be written as follows:

$$\begin{array}{ll} \max z_{1} \\ \text{s.t.} z_{1} & (n \ 1)D + \mathop{a}\limits^{n}_{i=1} v_{i} \\ u_{i} & D \ v_{i} \ u_{i} \ Dd_{i}; \ i = 1; \dots; n \\ w \ v_{i} \ w + (M \ D)d_{i}; \ i = 1; \dots; n \\ u_{i} \ 0; \ d_{i} \ 2 \ f \ 0; \ 1g; \ i = 1; \dots; n \end{array}$$
(2.17)

The following is proved in Hooker and Williams (2012).

Theorem 4. Model (2.17) is a correct formulation of P

Whenk 2, the expression (2.13)  $f\bar{\Phi_k}(\mathbf{u})$  implies that problem  $\mathbf{P}_k^0$  can be written

The constraints (2.18c) are included to ensure that the problem is MILP representable  $\bar{u}_{i_1}$ Sine constant, the hypograph is a union of bounded polyhedra whose recession cones consist of the origin only and are therefore identical.

The MILP model for  $P_k^0$  when k = 2; ...; n is

maxz <sub>k</sub>		
s.t. z <sub>k</sub> (n k+ 1)s + <mark>å</mark> v <sub>i</sub> i²I <sub>k</sub>	(a)	
$0 v_i Md_i; i 2 I_k$	(b)	
$v_i  u_i  \bar{u_{i1}}  D+M(1  d_i); \ i \ 2 \ I_k$	(c)	
s ū <sub>i1</sub> + D	(d)	
S W	(e)	(2 10)
w $u_i$ ; i 2 $I_k$	(f)	(2.19)
u <sub>i</sub> w+M(1 e <sub>i</sub> );i2I <sub>k</sub>	(g)	
$a_{i2l_k} e_i = 1$	(h)	
w ū <sub>ik 1</sub>	(i)	
u <sub>i</sub> ū <sub>i1</sub> M;i2I <sub>k</sub>	(j)	
d <sub>i</sub> ;e <sub>i</sub> 2 f 0;1g; i 2 I <sub>k</sub>		

Theorem 5. Model (2.19) is a correct formulation  $o_k^{0}$  for k = 2; ...; n.

Proof. We rst show that given an  $(\mathbf{y}\mathbf{u}; \mathbf{z}_k)$  that is feasible for (2.18), where  $\mathbf{u}_{i_j} = \bar{\mathbf{u}}_{i_j}$  for  $j = 1; \dots; k = 1$ , there exists  $\mathbf{v}; \mathbf{d}; \mathbf{e}; w, s$  for which  $(\mathbf{u}; \mathbf{z}_k; \mathbf{v}; \mathbf{d}; \mathbf{e}; w, s)$  is feasible for (2.19). Constraint (2.1) of ollows

directly from (2.1&). To satisfy the remaining constraints in (2.19), we set

$$(d_{i}; e_{i}; v_{i}) = \begin{cases} x \\ y \\ (0; 0; 0); \\ (0; 1; 0); \\ (0; 1; 0); \\ (0; 1; 0); \\ (1; 0; u_{i} & \bar{u}_{i_{1}} & D); \\ (1; 0; u_{i} & \bar{u}_{i_{1}} & D); \\ (1; 1; u_{i} & \bar{u}_{i_{1}} & D); \\ (2.20) \end{cases}$$

$$(2.20)$$

$$w = u_{k}$$

$$s = \min f \bar{u}_{i_{1}} + D; u_{k} g$$

where k is an arbitrarily chosen index ih such that  $u_k = u_{hki}$ . It is easily checked that these assignments satisfy constraint ((h). They satisfy (i) because (2.12) implies that  $u_k = \bar{u}_{i_{k-1}}$ . To show they satisfy (2.12), we note that (2.12) is implied by (2.12) because  $\bar{u}_{i_1} + D$ ;  $u_k g = s$  and  $(u_i = \bar{u}_{i_1} = D)^+ = v_i$  for i 2 I<sub>k</sub>. Since (2.12) is satisfied by (u; z), it follows that (2.12) is satisfied by (2.20).

For the converse, we show that for  $a(\mathbf{u}, \mathbf{z}_k; \mathbf{v}; \mathbf{d}; \mathbf{e}; \mathbf{w}; \mathbf{s})$  that satis es (2.19)( $\mathbf{u}; \mathbf{z}_k$ ) satis es (2.18). Constraint (2.16) follows from (2.19) and (2.19), and (2.16) is identical to (2.19). To verify that (2.16a) is satis ed, we left be the index for which  $\mathbf{e}_k = 1$ , which is unique due to (2.19). It suf ces to show that (2.16a) implies (2.16a) when the remaining constraints of (2.19) are satis ed. For this it suf ces to show that

s minf 
$$\bar{u}_{i_1}$$
 + D;  $u_k g$  (2.21)

$$v_i (u_i \bar{u}_{i_1} D)^+; i 2 I_k$$
 (2.22)

(2.21) follows from (d), (e), and (f) of (2.19). (2.22) follows from (d) and (c) of (2.19). This proves the theorem.

# 2.7 Valid Inequalities

In this section, we identify some valid inequalities that can strengthen the MILP model for k 2. The MILP model (2.17) foP<sub>1</sub><sup>0</sup> is already sharp, meaning that the inequality constraints of the model describe the convex hull of the feasible set, and there is therefore no bene t in adding valid inequalities. The sharpness property may be lost when budget constraints are added, but the resulting model may remain a relatively tight formulation. When 3, the model P<sub>k</sub><sup>0</sup> for k 2 become sharp when the valid inequalities described below are added. This is not true nwheth but the valid inequalities nonetheless tighten the formulation.

The sharpness of the MILP model (2.17)  $\mathbf{R}$  is proved in Hooker and Williams (2012). We present a simpler proof in Appendix<sup>2</sup>2.

Theorem 6. The MILP model (2.17) is a sharp representation  $\beta$ (2.16).

<sup>&</sup>lt;sup>2</sup>The proof in Hooker and Williams (2012) can be simplied by using only the multiplæers  $\frac{M}{nD}$   $a_i$  1+  $\frac{D}{M}$  for i = 1; ...; n, because each 1 D=M. The multipliers  $b_{ij}$  in their proof are unnecessary.

We now describe a class of valid inequalities that can be added to the MILP model ( $2P_k^{09}$ ) of for k 2 to tighten the formulation.

Theorem 7. The following inequalities are valid for k = 2:

$$z_k \stackrel{a}{}_{i2l_k} u_i$$
 (2.23)

$$z_k$$
 (n k+1) $u_i$  + b  $a_{i_{k-1}}(u_j - \bar{u}_{i_{k-1}})$ ; i 2 I<sub>k</sub> (2.24)

where

$$b = \frac{M}{M} \frac{D}{(\bar{u}_{i_{k-1}} - \bar{u}_{i_{1}})} = 1 \frac{D}{M} - 1 \frac{\bar{u}_{i_{k-1}} - \bar{u}_{i_{1}}}{M}$$
(2.25)

Proof. It sufces to show that for any  $(\mathbf{u}; \mathbf{z}_k; \mathbf{v}; \mathbf{d}; \mathbf{e}; \mathbf{w})$  that satis es (2.19), where  $\mathbf{e}_{i_j} = \bar{\mathbf{u}}_{i_j}$  for  $j = 1; \dots; k$  1, the vector  $\mathbf{u}$  satis es (2.23) and (2.24). Since we know from Theorem 5 **thist** feasible in (2.18), it sufces to show that (2.18) implies (2.23) and (2.24). To derive (2.23), we write (2.18a) as

$$z_k \stackrel{a}{}_{i^2 l_k} \min \bar{u}_{i_1} + D; u_{hki} g + (u_i \bar{u}_{i_1} D)^+$$
 (2.26)

For any term in the summation, we consider two cases  $u_i$  If  $\bar{u}_{i_1} + D$ , then  $u_{hki}$   $\bar{u}_{i_1} + D$  (because  $u_{hki}$   $u_i$ ), and the term reduces  $tq_{ki}$ . If  $u_i > \bar{u}_{i_1} + D$ , term becomes

minf 
$$\overline{u}_{i_1}$$
 + D;  $u_{hki}$  g + ( $u_i$   $\overline{u}_{i_1}$  D) = minf 0;  $u_{hki}$   $\overline{u}_{i_1}$  Dg +  $u_i$   $u_i$ 

In either case, termis less than or equal to, and (2.23) follows.

To establish (2.24), it is enough to show that (2.24) is implied by (2.18) for ia2dh. We consider the same two cases as before.

Case 1: $u_i = \bar{u}_{i_1}$  D, which implies  $u_{hki} = \bar{u}_{i_1}$  D. Since u satistics es (2.18a), we have

$$z_{k} (n \ k+1)u_{hki} + \mathop{a}\limits_{\substack{j \\ j \\ u_{j} \\ \bar{u}_{i}, > D}} (u_{j} \ \bar{u}_{i_{1}} D)$$
(2.27)

It suf ces to show that this implies

$$z_{k} (n \ k+1)u_{i} + b \ \overset{o}{a}(u_{j} \ \overline{u}_{i_{k-1}}) + \overset{o}{a}(u_{j} \ \overline{u}_{i_{k-1}}); \qquad (2.28)$$

because (2.28) is equivalent to the desired inequality (2.24). But (2.27) implies (2.28) because by de nition of  $u_{rki}$ ,  $u_j = \bar{u}_{i_{k-1}} = 0$  for all j 2 I<sub>k</sub> due to (2.18), and it can be shown that

$$b(u_j \ \bar{u}_{i_{k-1}}) \ u_j \ \bar{u}_{i_1} \ D$$
 (2.29)

for any j 2  $I_k$ . To show (2.29), we note that the de nition **b**fimplies the following identity:

$$\bar{u}_{i_1}$$
  $\bar{u}_{i_{k-1}} + D = (1 \ b)(M + \bar{u}_{i_1} \ \bar{u}_{i_{k-1}})$ :

Adding  $\begin{pmatrix} 1 & b \end{pmatrix} \bar{u}_{i_{k-1}}$  to both sides, we obtain

$$\bar{u}_{i_{k-1}} \quad b \bar{u}_{i_{k-1}} + D = (1 \ b)(M + \bar{u}_{i_1}) \quad (1 \ b)u_j;$$
 (2.30)

where the inequality holds becauld  $\bar{u}_{i_1}$   $u_j$  due to (2.18). We obtain (2.29) by rearranging (2.30).

Case  $2u_i$   $\bar{u}_{i_1} > D$ . It again suf ces to show that (2.18) implies (2.28). Due to the case hypothesis, we have from (2.18) that

$$z_k$$
 (n k+1) minf  $\bar{u}_1$  + D;  $u_{hki}$  g+( $u_i \ \bar{u}_{i_1}$  D) +  $a_k^*(u_j \ \bar{u}_{i_1}$  D)<sup>+</sup>  
 $j 2 I_k n f i g u_i \ \bar{u}_{i_1} > D$ 

This can be written

$$\begin{array}{ll} z_k & (n \ k+1)u_i & (n \ k+1) \ u_i & minf \, \bar{u}_1 + D; u_{hki} \, g \\ & + (u_i \ \bar{u}_{i_1} \ D) + \mathop{a}\limits^{*}_{a} (u_j \ \bar{u}_{i_1} \ D)^{-1} \\ & & j \, 2 I_k n f i g \\ & & u_j \ \bar{u}_{i_1} > D \end{array}$$

which can be written

$$z_{k} (n k+1)u_{i} (n k) u_{i} \min \overline{u}_{1} + D; u_{hki}g$$

$$\overline{u}_{1} + D \min \overline{u}_{1} + D; u_{hki}g + \overset{\circ}{a}(u_{j} \overline{u}_{i_{1}} D)^{+} (2.31)$$

$$j 2 I_{k} n f ig u_{i} \overline{u}_{i_{1}} > D$$

The second term is nonpositive because  $\bar{u}_1 + D$  by the case hypothesis, and  $u_{hki}$ . The third term is clearly nonpositive. Thus (2.31) implies (2.28) because $\bar{u}_{i_{k-1}} = 0$  and (2.29) holds for j 2  $I_k$  as before.

# 2.8 Modeling Groups of Individuals

In many applications, utility is naturally allocated to groups rather than individuals, where individuals within each group receive an equal allocation. This occurs in the examples of Section 6.3, in which groups correspond to classes of patients with the same disease/prognosis or to neighborhood populations. In other applications, the number of individuals may be too large for practical solution, since problemP<sub>i</sub> must be solved for each individual In such cases, individuals can typically be grouped into a few classes within which the individual differences are small or irrelevant, thus making

the problem tractable and the results easier to digest. We therefore modify the above SWFs to accommodate groups rather than individuals.

We suppose there aregroups of possibly different sizes. We letdenote the utility of each individual in groupi ands the number of individuals in the group.

We begin by deriving  $G_1(\mathbf{u})$ . Let  $u_{i^0}^0$  be the utility of individual  $i^0$ , and let  $u_i$  be the utility of each individual in group i. There are  $n^0$  individuals and groups. Let be the size of group so that

$$n^{0} = \mathop{a}_{i=1}^{n} s_{i}$$
 (2.32)

Then

Since  $u_{h1i} = u_{h1i}^0$  and group has size, we have

$$G_{1}(\mathbf{u}) = \overset{n}{\underset{i=1}{a}} s_{i} 1 D + \overset{n}{\underset{i=1}{a}} s_{i} u_{Mi} + \overset{n}{\underset{i=1}{a}} s_{i} (u_{i} u_{Mi} D)^{+}$$
(2.33)

This is the formula used in Hooker and Williams (2012). Hooker and Williams prove the following. Theorem 8. The problem  $\stackrel{\text{P}}{\text{P}}$  modi ed for groups, is equivalent to the MILP model

We now derive  $\overline{G}_k(\mathbf{u})$  for k 2. Recall that the SWF for individuals is

$$\bar{F}_{k^{0}}(\mathbf{u}^{0}) = (n^{0} k^{0} + 1) \min \{u_{M1i}^{0} + D; u_{Hk^{0}}^{0}g + a_{i^{0}=k^{0}}^{n^{0}}(u_{Hi^{0}}^{0} \bar{u}_{M1i}^{0} D)^{+}$$
(2.35)

To obtain  $\bar{G}_k(\mathbf{u})$ , we again assume the individuals in each grobup ve the same utility. The rst individual in (2.35) that belongs to group individual

$$k^{0} = 1 + \mathop{a}_{j=1}^{k} \mathbf{s}_{j}^{1}$$
 (2.36)

Due to (2.32) and (2.36), the rst term on the RHS of (2.35) is

$$n^{0} = 1 = a^{n}_{j=1} s_{i_{j}} + 1 = minf u_{Hki} + D; u_{Hki} g = a^{n}_{i=k} s_{Hii} = minf u_{Hki} + D; u_{Hki} g$$

since all the utilities in a group are the same. Thus we have

$$\bar{G}_{k}(\mathbf{u}) = \overset{n}{\underset{i=k}{a}} s_{\text{tii}} \quad \text{minf } u_{\text{h1i}} + D; u_{\text{hki}} g + \overset{n}{\underset{i=k}{a}} s_{\text{tii}} (u_{\text{hii}} u_{\text{h1i}} D)^{+}$$
(2.37)

Theorem 9. The functions  $\overline{G}_k(\mathbf{u})$  are continuous in  $\mathbf{W}_i$ ; ...;  $\mathbf{u}_{mi}$  for k = 1; ...; n.

Proof. It suf ces to show each term of (2.37) is a continuous function  $\mathbf{h}_{R}f$ ;...; $\mathbf{u}_{hni}$ , with  $\mathbf{u}_{h1i}$ ;...; $\mathbf{u}_{hk-1i}$  and the corresponding group size  $\mathbf{x}_{R}\mathbf{x}_{f}$ ;...; $\mathbf{x}_{hk-1i}$  xed. The rst term is continuous because it is equal to a constant time the minimum of order statisting and  $\mathbf{u}_{hki}$ , which are continuous functions **u** f Similarly, each term of the summation is a constant times the maximum of a continuous expression and zero.

The MILP model is very similar to the one we developed for (2.18):

$$\max z_{k}$$
s.t.  $z_{k}$   $\stackrel{a}{a}$  s s +  $\stackrel{a}{a}$  s v<sub>i</sub> (a)  
(2:19b)–(2:19j) (b)–(j)  
d<sub>i</sub>; e 2 f 0; 1g; i 2 l<sub>k</sub>
(2.38)

Theorem 10. The problem **P** reformulated for groups, is equivalent to (2.38) for  $k_2; ...; n$ .

Proof. We rst show that given any  $(\mathbf{u}; \mathbf{z}_k)$  that is feasible for (2.18), where  $\mathbf{u}_{i_j}$  for  $j = 1; \dots; k = 1$ , there exist,  $\mathbf{d}; \mathbf{e}; w, \mathbf{s}$  for which  $(\mathbf{u}; \mathbf{z}_k; \mathbf{v}; \mathbf{d}; \mathbf{e}; w, \mathbf{s})$  is feasible for (2.38). Constraint (2.38) follows directly from (2.1&). To satisfy the remaining constraints in (2.38), we assign value as  $(\mathbf{z}, \mathbf{z}, \mathbf{v}; \mathbf{d}; \mathbf{e}; w, \mathbf{s})$  is feasible for (2.38). Constraint (2.38) we assign value as a in (2.20), where is an arbitrarily chosen index in such that  $\mathbf{u}_k = \mathbf{u}_{ki}$ . It is easily checked that these assignments satisfy constraints (2)-3(2.38). They satisfy (2.38) because (2.16) implies that  $\mathbf{u}_k = \bar{\mathbf{u}}_{i_k - 1}$ . To show they satisfy (2.38), we note that (2.38) is implied by (2.18) because minf  $\bar{\mathbf{u}}_{i_1} + D; \mathbf{u}_k g$  s and  $(\mathbf{u}_i = \bar{\mathbf{u}}_{i_1} = D)^+ = \mathbf{v}_i$  for i 2 I<sub>k</sub>. Since (2.18) is satisfied by ( $\mathbf{u}; \mathbf{z}$ ), it follows that (2.38) is satisfied by (2.20).

For the converse, we show that for  $a(\mathbf{u}, \mathbf{z}_k; \mathbf{v}; \mathbf{d}; \mathbf{e}; \mathbf{w}; \mathbf{s})$  that satis es (2.38)( $\mathbf{u}; \mathbf{z}_k$ ) satis es (2.18). Constraint (2.16) follows from (2.19) and (2.19), and (2.16) is identical to (2.38). To verify that (2.16a) is satis ed, we let be the index for which  $\mathbf{e}_k = 1$ , which is unique due to (2.36). It suf ces to show that (2.36a) implies (2.16a) when the remaining constraints of (2.38) are satis ed. For this it suf ces to show that

s minf 
$$\bar{u}_{i_1}$$
 + D;  $u_k g$  (2.39)

$$v_i (u_i \bar{u}_{i_1} D)^+; i 2 I_k$$
 (2.40)

(2.39) follows from (1), (e), and (f) of (2.38). (2.40) follows from (a) and (c) of (2.38). This proves the theorem.

Hooker and Williams prove that (2.17) is a sharp representation  $p_1^0$ , and (2.34) a sharp representation of  $P_1^0$  reformulated for groups. We present a simpler proof of both theorems. It is necessary only to prove the latter, because the former is a special case of it.

Theorem 11. The MILP model (2.34) is a sharp representation of feformulated for groups.

of Theorems 6 and 11We prove Theorem 11, of which Theorem 6 is a special case in vehichfor eachi. It suffices to show that any inequality  $\mathbf{a}^T \mathbf{u} + \mathbf{b}$  that is valid for  $P_1^0$  is a surrogate (nonnegative linear combination) of inequalities in (2.34). Let

We rst show that the following is a surrogate of (2.34) for any

$$z_1$$
 (N 1)D+  $s_i$  + (N  $s_i$ ) $\frac{D}{M}$   $u_i$  + 1  $\frac{D}{M}$   $\mathring{a}_{j \in i} s_j u_j$ : (2.41)

We then show that  $\mathbf{a}^{\mathsf{T}}\mathbf{u} + \mathbf{b}$  is a surrogate of the inequalities (2.41). The theorem follows.

To show that (2.41) is a surrogate of (2.34), we rst note that the following is a linear combination of the upper bounds on in (2.34b) and (2.34c), using multipliers D and H(M = D), respectively:

$$v_j = \frac{D}{M}w + 1 = \frac{D}{M}u_j$$
: (2.42)

We also have the following from (2.34b) and (2.34c):

$$v_i \quad u_i$$
: (2.43)

$$w v_i$$
: (2.44)

We now obtain the following, for any giverand j, as a linear combination of (2.42) and (2.44), using multipliers 1 and D=M, respectively:

$$v_j = \frac{D}{M}v_i + 1 = \frac{D}{M}u_j$$
: (2.45)

Finally, we obtain (2.41) for any given by summing (2.34a) with multiplier 1, (2.43) with multiplier

and (2.45) over all 6 i with multiplier si.

It remains to show that  $\mathbf{a}^T \mathbf{u} + \mathbf{b}$  is a surrogate of (2.41) for  $\mathbf{i}_1, \dots, \mathbf{w}$  rst observe that  $(\mathbf{u}; z) = (0; (N \ 1)D)$  is feasible in  $P_1^0$  and must therefore satisfy  $\mathbf{a}^T \mathbf{u} + \mathbf{b}$ , which implies b (N 1)D. We can assume without loss of generality that  $(N \ 1)D$ , since otherwise we can add an appropriate multiple of the valid inequality b to obtain the desired inequality  $\mathbf{a}^T \mathbf{u} + \mathbf{b}$ . We also note that  $(\mathbf{u}; z) = (M; \dots; M; NM + (N \ 1)D)$  is feasible and must satisfy  $\mathbf{a}^T \mathbf{u} + (N \ 1)D$ , which means

$$\overset{\text{m}}{\overset{\text{a}}{\underset{j=1}{a_j}}} a_j \quad N \tag{2.46}$$

Finally, we note that  $(\mathbf{u}; \mathbf{z}) = (M\mathbf{e}_1; (N \ 1)D + \mathbf{s}(M \ D))$  is feasible for  $\mathbf{P}_1^0$ , where  $\mathbf{e}_1$  is the ith unit vector. Substituting this int  $\mathbf{z}_1 \ \mathbf{a}^T \mathbf{u} + (N \ 1)D$ , we obtain

$$a_i = 1 - \frac{D}{M} s_i$$
 (2.47)

Due to (2.46), we can suppose without loss of generality  $a_{j} = N$ , since otherwise we can add appropriate multiples of the valid inequalities  $a_{j}$  to obtain  $z_{1} = \mathbf{a}^{T}\mathbf{u} + \mathbf{b}$ .

To obtain  $\mathbf{a}^{\mathsf{T}}\mathbf{u}$  + b as a surrogate of (2.41), we sum (2.41) over jalsing the multipliers

$$a_i = \frac{M}{ND} a_i \qquad 1 \quad \frac{D}{M} s_i \qquad (2.48)$$

for eachi. It is easily checked that  $a_{i=1}^{m} a_{i} = 1$ , so that the linear combination has the form

$$z d^{T}u + (N 1)D$$
(2.49)

We wish to show that  $\mathbf{a} = \mathbf{a}$ . Note that

d<sub>i</sub> = s<sub>i</sub> + N
$$rac{D}{M}$$
 a<sub>i</sub> + 1  $rac{D}{M}$  s<sub>jei</sub>aj

Using the fact that  $a_{i=1}^{m} a_{j} = 1$ , this becomes

$$d_i = N \frac{D}{M} a_i + 1 \frac{D}{M} s_i$$

which immediately reduces  $t\mathbf{d}_i = a_i$ . We conclude that (2.49) is a linear combination of the inequalities (2.41) using multipliers. It remains to show that each is nonnegative, but this follows from (2.47) and (2.48).

Finally, we describe a set of valid inequalities for the MILP model (2.38) kfor2.

Theorem 12. The following inequalities are valid for the group problem  $f \theta r k = 2$ :

$$z_k \overset{a}{\underset{i2l_k}{a}} su_i$$
 (2.50)

$$z_k = \mathop{a}_{j2l_k} s_i u_j + b \mathop{a}_{j2l_k} s_j(u_j = \bar{u}_{i_{k-1}}); i 2 l_k$$
 (2.51)

whereb is given by (2.25).

Proof. Recall that the group version  $\Theta_k^p$  for k 2:

$$\begin{array}{cccc} \max z_{k} & \\ \text{s.t. } z_{k} & \underset{i2l_{k}}{a} & \min f \, \bar{u}_{i_{1}} + D; u_{hki} \, g + \underset{i2l_{k}}{a} & (u_{i} & \bar{u}_{i_{1}} & D)^{+} & (a) \\ & & & \\ u_{i} & \bar{u}_{i_{k-1}}; \ i \ 2 \ l_{k} & (b) \\ & & & \\ u_{i} & \bar{u}_{i_{1}} & M; \ i \ 2 \ l_{k} & (c) \end{array}$$

It suf ces to show that for an yu;  $z_k$ ; v; d; e; w) that satis es (2.38), when  $a_j = \bar{u}_{ij}$  for j = 1; ...; k = 1, the vector **u** satis es (2.23) and (2.24). Since we know from Theorem 10 **u** fiatfeasible in (2.52), it suf ces to show that (2.52) implies (2.50) and (2.51). To derive (2.50), we write  $a_2^2 a_3^2$ 

$$z_k = \mathop{a}_{i21_k} s_i \min f \bar{u}_{i_1} + D; u_{hki} g + (u_i = \bar{u}_{i_1} = D)^+$$
 (2.53)

For any term in the summation, we consider two cases  $u_i$  If  $\bar{u}_{i_1} + D$ , then  $u_{hki}$   $\bar{u}_{i_1} + D$  (because  $u_{hki}$   $u_i$ ), and the term reduces  $\mathbf{s} u_{hki}$ . If  $u_i > \bar{u}_{i_1} + D$ , term becomes

 $s_i \min \overline{u}_{i_1} + D; u_{hki}g + (u_i \overline{u}_{i_1} D) = s_i \min 0; u_{hki} \overline{u}_{i_1} Dg + u_i s_i u_i$ 

In either case, termis less than or equal top, and (2.50) follows.

To establish (2.51), it is enough to show that (2.51) is implied by (2.52) for ia2dh. We consider the same two cases as before.

Case 1: $u_i$   $\bar{u}_{i_1}$  D, which implies  $u_{tki}$   $\bar{u}_{i_1}$  D. Since **u** satistics (2.52a), we have

$$z_k \overset{a}{\underset{j^2 I_k}{a}} s_j u_{hki} + \overset{a}{\underset{j^2 I_k nfig}{a_{i_1} > D}} s_j(u_j \quad \overline{u_{i_1}} \quad D)$$
(2.54)

It suf ces to show that this implies

$$Z_{k} = \mathop{a}_{j \ge I_{k}} s_{j} u_{i} + b = \mathop{a}_{s} s_{j}(u_{j} = \bar{u}_{i_{k-1}}) + \mathop{a}_{s} s_{j}(u_{j} = \bar{u}_{i_{k-1}}); \quad (2.55)$$

because (2.55) is equivalent to the desired inequality (2.51). But (2.54) implies (2.55) because by de nition of  $u_{rki}$ ,  $u_j = \bar{u}_{i_{k-1}} = 0$  for all j 2  $I_k$  due to (2.52), and (2.12) for all j 2  $I_k$ .

Case  $2u_i$   $\bar{u}_{i_1} > D$ . It again suf ces to show that (2.52) implies (2.55). Due to the case hypothesis, we have from (2.5*a*) that

$$\begin{array}{cccc} z_k & \mathop{a}\limits^{a} s_j & minf \, \bar{u}_1 + D; u_{hki} g + s_i (u_i & \bar{u}_{i_1} & D) + & \mathop{a}\limits^{a} s_j (u_j & \bar{u}_{i_1} & D) \\ & & & & j^{2l_k nfig} \\ & & & u_j & \bar{u}_{i_1} > D \end{array}$$

This can be written

which can be written

$$\begin{array}{ll} z_{k} & \overset{*}{a} s_{j} u_{i} & \overset{*}{a} s_{j} u_{i} & \min f \bar{u}_{1} + D; u_{hki} g \\ s_{j2l_{k}} & & j_{2l_{k}nfig} \\ s_{i} \bar{u}_{1} + D & \min f \bar{u}_{1} + D; u_{hki} g + \overset{*}{a} s_{j} (u_{j} \ \bar{u}_{i_{1}} D) \end{array}$$

$$\begin{array}{l} (2.56) \\ & & j_{2l_{k}nfig} \\ u_{j} \ \bar{u}_{i_{1}} > D \end{array}$$

The second term is nonpositive because  $\bar{u}_1 + D$  by the case hypothesis, and  $u_{hki}$ . The third term is clearly nonpositive. Thus (2.56) implies (2.55) becaujes  $\bar{u}_{i_{k-1}} = 0$  and (2.29) holds for j 2  $I_k$  as before.

# 2.9 Applications

We now implement our approach on a healthcare resource allocation problem and a disaster management problem. We solve all MILP instances using Gurobi 8.1.1 on a desktop PC running Windows 10.

### 2.9.1 Healthcare Resource Allocation

A proper balance between fairness and ef ciency is crucial in the allocation of healthcare resources. Hooker and Williams (2012) study a problem in which treatments are made available to patients on the basis of their disease and prognosis. In discussing this case, we caution that the results we report should not be taken as general recommendations for the allocation of medical resources. They are based on cost and clinical data speci c to a particular set of circumstances. We use this example because it allows comparison with the published H–W results on the same problem instance. Patients are divided into groups based on their disease and prognosis. There is one treatment potentially available to each patient group, and for policy consistency, it is provided to either all or none of the group members. Binary variates 1 if groupi receives the recommended treatment and 0 otherwise. The average utility experienced by members of groups measured in terms of quality adjusted life years (QALYs)q<sub>i</sub> is the net gain in QALYs for a member of group when receiving the recommended treatment, ands the expected QALYs experienced with medical management without the treatment. Thus

$$u_i = a_i + q_i y_i; i = 1;...;n$$
 (2.57)

The budget constraint is

$$a_{i}^{n}$$
sıçıy<sub>i</sub> B (2.58)

wheres is the group sizec, the cost of treating one patient in groupandB the total available budget. The budget is set so as to force some hard decisions. The constraints (2.57)–(2.58), along with



Fig. 2.6 Runtime in the shelter allocation example



Fig. 2.7 Utility distributions in shelter allocation instance cp@=(50, m = 25)

Figs. 2.7 and 2.8 show the evolution of per capita utility in individual neighborhood strats eases. The shaded region indicates which utilities are in the fair region (w the worst). We see immediately that the problem is highly constrained, because the lowest utilities quickly reach a plateau and remain at a low level even for largevalues. These neighborhoods are located at a considerable distance from candidate shelter locations, and so they remain disadvantaged even when given high priority. In this type of situation, it is particularly important to use a leximax rather than a maximin criterion of fairness, so as to take into account the situation of disadvantaged neighborhoods other than the very worst-off. If a maximin criterion is used, only the most distant neighborhood is given special priority, and the other disadvantaged neighborhoods bene t only from tie-breaking and the fact that maximizing the worst-off imposes a oor on their utility level. This can be seen in Fig. 2.9,

SVM models, Zafar et al. (2017a, 2019) directly apply solvers for convex and non-convex programs, and Olfat and Aswani (2018) devise an iterative algorithm where each iteration relies on a non-convex program solver. As off-the-shelf solvers are designed to handle general models with only necessary structural assumptions, they may not be ef cient for all problem instances, hence these fair SVM training methods may encounter computational issues in complex, large-scale real applications.

## 3.1.1 Our Approach and Results

An overlooked opportunity in algorithmic development of fair SVM is to build connections with standard SVM training, then make use of the extensive SVM training techniques in literature. In this paper, we examine this opportunity and study whether algorithmic and computational techniques used to train standard SVMs can be extended to train fair SVMs. Our main contribution is that we extend two dual-based SVM training algorithms, the DCD algorithm for linear SVM and the SMO algorithm for general SVM, to handle fair SVM formulations with fairness constraints that are linear in the weights of SVM. The extension utilizes a coordinate descent type subroutine for updating the dual multipliers corresponding to fairness constraints. Moreover, the numerical techniques of shrinking and caching that are effective for improving standard SVM training are valid for the modi ed algorithms as well. We evaluate the empirical performance of these specialized training methods for fair SVM and demonstrate that their computational costs are comparable with the corresponding standard counterparts. Moreover, they can be more efficient than a state-of-the-art solver on large instances.

SVM Formulations and Notations. We study the standard soft-margin C-SVM formulation for binary classi cation. Suppose a training set containing ample points  $i\mathfrak{D} = f(x_i; y_i; z_i)g_{i=1}^n$ , where  $x_i 2 X$  are feature vectors  $y_i 2 f 1$ ; 1g is i's true label, and  $z_i 2 f 0$ ; 1g indicates i's group membership in a protected class (e.g. gender, race). Note that the group features will be used to de ne fairness constraints. Let( $x_i$ ) denote a map of feature and C be the penalty parameter, the  $(2n_1)$ and (3.2) are respectively the primal and dual SVM formulations.

$$\min_{q;b;x} \frac{1}{2} kqk^2 + C \mathring{a}_{i=1}^n x_i \quad \text{s.t. } x_i \quad 1 \quad y_i(hq; f(x_i)i + b); x_i \quad 0.8i:$$
(3.1)

$$\min_{m} e^{T}m + \frac{1}{2}m^{T}Qm \quad \text{s.t.} \quad \overset{a}{a}_{i=1}^{n}my_{i} = 0; 0 \quad m \quad C \; 8i: \qquad (3.2)$$

In (3.2), m2 R<sup>n</sup> denotes the vector of dual variables **anis** the all1 vector in R<sup>n</sup>. Q is an n matrix with  $Q_{i;j} = y_i y_j K(x_i;x_j)$ , where  $K(x_i;x_j) = hf(x_i)$ ;  $f(x_j)$  is represented the kernel function. For notation ease, we use the abbreviation  $K(x_i;x_j)$ . One notable advantage of SVM is that using different kernel functions in (3.2) leads to a wide range of predictive models. For example, two popular options are: linear function kernel  $(x_i;x_j) = hx_i;x_ji$ , Gaussian radial basis function (RBF) kernel  $K(x_i;x_j) = exp(x_i + x_j)^2 = s^2$ .

Supposen is an optimal dual solution, then the weight and biasb in the optimal classi eq can be computed as follows. We use;  $y_{sv}$  to denote a training sample point that is a support vector, namely its corresponding dual optimal solution satismes 2 (0;C):

$$q = \mathbf{a}_{i=1}^{n} \mathbf{m}_{i} y_{i} f(x_{i}); \mathbf{b} = y_{sv} \quad \mathbf{a}_{i=1}^{n} \mathbf{m}_{i} Y_{i} K(x_{i}; x_{sv}):$$
(3.3)

A special class of SVM that is often studied separately is the linear SVM, which uses a linear kernel function, or equivalently  $(x_i) = x_i$ . In this case(3.1) can be simplied by treating the bias term b as an additional dimension in As shown in Hsieh et al. (2008), the dual formulation for linear SVM does not need the equation constraint and the optimal can be obtained together from the optimal dual solution.

$$\min_{m} e^{T}m + \frac{1}{2}m^{T}Qm \quad s.t. \quad 0 \quad m \quad C \; 8i:$$
 (3.4)

$$q = a_{i=1}^{n} m y_i x_i; b = a_{i=1}^{n} m y_i:$$
 (3.5)

## 3.1.2 Related Literature

There is a large number of SVM training algorithms in literature. As summarized in the survey paper Shawe-Taylor and Sun (2011), majority of proposed methods solve the dual form (Bag) and well-studied methodologies include interior point method, decomposition method, active set method and coordinate descent. We focus on reviewing literature related to decomposition method for general SVM and coordinate descent for linear SVM, as they are more relevant to our study.

Decomposition methods solve the complete dual formulation, a quadratic program (QP), by breaking it down into a series of smaller QP subproblems. This technique is useful for handling the potential computational challenge rising from a large number of training samples and features. Each subproblem optimizes the same objective function but with the restriction that only a subset of the dual variables, typically referred to asworking set can be modi ed. A decomposition method iteratively uses subproblems to search among feasible solutions until convergence to an optimal solution. Two popular software packages for decomposition methods and light (Joachims (1998)) and LIBSVM (Chang and Lin (2011)).

Early works such as Osuna et al. (1997a) and Osuna et al. (1997b) consider constant size working sets in decomposition methods. The seminal paper Joachims (1998) formalize the working set selection as an optimization problem and design **BMe**/l<sup>light</sup> algorithm as an ef cient implementation of decomposition methods to train SVM with arbitrary kernel function. Later, Platt (1998) propose the sequential minimal optimization (SMO) algorithm, a simpler and faster decomposition method where each working set contains exactly two variables. Follow-up works have been improving SMO algorithm with more ef cient working sets selection heuristics, e.g. Fan et al. (2005); Glasmachers et al. (2006); Keerthi et al. (2001); List and Simon (2004); Yang et al. (2019). Some key developments include the rst-order information bas**ed** aximum Violating Paiselection rule from Keerthi et al.

(2001), second-order information based heuristics from Fan et al. (2005); Glasmachers et al. (2006); List and Simon (2004), and more recently Yang et al. (2019) incorporate higher order information to select working sets. The LIBSVM library Chang and Lin (2011) implements SMO as described in Fan et al. (2005) for common SVM formulations and kernel functions.

Parallel with the research on algorithm design, another stream of papers focus on theoretical properties of various decomposition methods for SVM. Keerthi and Gilbert (2002) prove the nite termination of the SMO algorithm stated in Keerthi et al. (2001) which set the stopping criteria as satisfying the KKT conditions associated w(th 1) and(3.2) within up tot violation. Lin (2002b) extend the analysis to general decomposition methods with working sets of size greater than two such aSVM<sup>light</sup>. As noted in papers including Lin (2001b), these nite termination results are not equivalent to convergence to the optimal solution due to the violation toletarkine (2001b) prove the asymptotic convergence of the SMO algorithm in Keerthi et al. (2001). In addition, Lin (2001a) prove the linear convergence of both algorithms Joachims (1998); Keerthi et al. (2001) when the dual objective function is strictly convex, which is possible for some kernel choices such as the Gaussian RBF kernel. A comprehensive convergence analysis of general decomposition methods for SVMs can be found in Chen et al. (2006).

Coordinate descent is another optimization technique that has been applied to train SVMs. Coordinate descent is an iterative algorithm that nds an optimal solution through successive coordinate-wise optimization. Early research on applying coordinate descent methods to the dual SVM formulations can be found in Mangasarian and Musicant (1999, 2001). Later works explore the potentials of this technique in large-scale linear SVMs, for example, Chang et al. (2008) and Hsieh et al. (2008) respectively design coordinate descent methods to solve the primal and the dual linear SVM models. The algorithms from Hsieh et al. (2008) are implemented in the LIBLINEAR (Fan et al. (2008)) library for large-scale linear classi cation. The convergence analysis can take advantage of relevant results for general descent methods from convex programming. For instance, Hsieh et al. (2008) apply the techniques from Luo and Tseng (1993) to prove the global convergence of their dual coordinate descent algorithm for linear SVM, which implies both asymptotic convergence and nite termination to a desirable accuracy level.

Numerical techniques in the implementation of these fore-mentioned specialized algorithms are useful for further speeding up SVM training. Two best-known techniqueshaiteking and caching both initially proposed in Joachims (1998). They are routinely implemented in decomposition methods (e.g. Fan et al. (2005); Joachims (1998)) and coordinate descent methods (e.g. Hsieh et al. (2008)). The shrinking technique takes advantage of the fact that an optimal dual solution contains a relatively small number of variables not at bounds, which correspond to the support vectors, and the remaining variables are bounded, i.e. = 0 or C. To accelerate training, we can identify dual variables that are likely to be bounded at optimality and temporarily remove these variables to reduce the size of the dual problem, namely, we ignore the shrunk elements when choosing working sets. The caching technique refers to storing some recently ulique in the available memory to reduce the needed

number of kernel function evaluations and computation time. This technique is particularly useful for large problems wher Q is too large to be fully stored an  $Q_{ij} = y_i y_j K_{ij}$  has to be calculated as needed.

Another related line of literature is fair machine learning (ML). Fair ML methods are typically designed by modifying standard ML methods before, during or after the training phase, respectively known as pre-, in-, or post-processing fair ML. Besides the variety in fairness seeking strategies, there is also a large number of fairness de nitions, which can be categorized as group versus individual fairness. In the contexts of SVMs, or more generally, classi cation models, group fairness de nitions with respect to the confusion matrix of classi cation are the most common. We refer to Mehrabi et al. (2019a) for a comprehensive survey of fair ML.

Next, we focus on a speci c in-processing approach that is closely related to our study. One can train a fair SVM using the following modi ed SVM formulation, where the spectrum desirable fairness levels.

$$\min_{q;b;x} \frac{1}{2} kqk^2 + C \mathring{a}_{i=1}^n x_i \text{ s.t. } x_i = 1 \quad y_i(hq; f(x_i)i + b); x_i = 0.8i; f_i(q; b; D) \quad e \in 8I: \quad (3.6)$$

Olfat and Aswani (2018); Zafar et al. (2017a, 2019) have studied this general model. Zafar et al. (2017a) de ne a convex proxy for the well-known demographic parity constraint using the covariance associated with the SVM decision boundary. Their constraint en strain en strain a convex program which can be readily solved by commercial solvers in some cases. Donini et al. (2018a) considers SVM as a special case of empirical risk minimization, and applies a similar proxy technique to approximate equalized opportunity constraints. A main motivation for our paper is that even though state-of-the-art solvers are highly ef cient f(36.6) with a linear kernel, they tend to suffer noticeable performance drop with a non-linear kernel such as RBF kernel. Specialized fair SVM algorithms, as we demonstrate, could be more robust to kernel choices. Olfat and Aswani (2018) also focus on demographic parity. They note that the convex proxy used in Zafar et al. (2017a) is a relatively weak relaxation of the exact demographic parity de nition. They propose a non-convex quadratic constraint as a more accurate proxy, then designed an iterative algorithm to (sch)ey solving a sequence of structured convex-concave programs with available solvers. More recently, Zafar et al. (2019) de ne a continuous proxy for the equalized odds constraints, which require equal false positive rate and false negative rate across groups. The new measure le(ads) to the form of a Disciplined Convex-Concave Program, which are still solvable with some existing solvers. These latter two papers are not explicitly comparable with our methods, since they work with a more dif cult non-convex fair SVM formulation whereas we aim to design specialized algorithms to solve convex fair SVM formulations.

## 3.2 Problem Formulation

We use(3.6) as the fair SVM model in this paper. In addition, we focus on a speci c form of fairness constraints.

Assumption 1. The fairness measure functions of q; D; moreover, they are linear in. For notation ease, we denote as:  $f_1(q; D) := hp_1(D); qi + q_1$ .

Note that  $p_l(D)$  is dependent only on the training data and its denition relects the desirable notion of fairness. To further simplify notation, we spet=  $p_l(D)$ . Assumption 1 ensures the convexity of (3.6) and makes it convenient to obtain the following Lagrangian dual formulation of fair SVM.

For easy comparison wit(3.2), we again usen 2 R<sup>n</sup> as the dual multipliers for the constraints on x<sub>i</sub>. Suppose(3.6) containsk fairness constraints, then we tep 2 R<sup>k</sup> denote dual multipliers for the fairness constraints, and 2 R<sup>k</sup> represents the constraint parameters whith eq. c<sub>i</sub> for all I = 1; ...; k. In addition, we use  $\bar{Q}$  to denote a matrix de ned similarly as in the standard SVM case:  $\bar{Q} = \begin{bmatrix} Q & A \\ A^T & P \end{bmatrix}$ , Q is the same as i(8.2) with Q<sub>ij</sub> = hy<sub>i</sub>f(x<sub>i</sub>); y<sub>j</sub>f(x<sub>j</sub>)i, A is ak n matrix with A<sub>li</sub> = h p<sub>i</sub>; y<sub>i</sub>f(x<sub>i</sub>)i, and P is ak k with P<sub>10</sub> = hp<sub>1</sub>; p<sub>10</sub>i.

$$\min_{\substack{m,g \\ m,g}} e^{T}m d^{T}g + \frac{1}{2} \frac{m}{g} \frac{\bar{Q}}{g} s.t \mathring{a}_{i=1}^{n} my_{i} = 0 (a); 0 m C; 8i; g 0; 8l: (3.7)$$

Proposition 1. Under Assumption 1(3.7) is the Lagrangian dual of the fair SVM formulation (3.6). In the linear SVM case, namely  $(x_i) = x_i$  in (3.6), constraint (a) is not needed in the dual formulation.

Proof. We use  $m; l_i; g: i \ge [n]; l \ge [k]g$  to denote the dual multipliers respectively for the three sets of constraints in (3.6). Under Assumption 1, the Lagrangian function for (3.6) is:

$$L(q;b;x;m|;g) = \frac{1}{2}kqk^{2} + C\underset{i=1}{\overset{n}{a}}x_{i} + \underset{i=1}{\overset{n}{a}}m(1 \quad y_{i}(hq;f(x_{i})i + b) \quad x_{i}) + \underset{i=1}{\overset{n}{a}}I_{i}(x_{i}) + \underset{i=1}{\overset{k}{a}}g(hp_{i};qi + c_{i} \quad q_{i})$$

By de nition, the Lagrangian dual is:

$$\max_{\substack{j:m,g2R_{+}^{n} \in R_{+}^{K}}} \min_{\substack{R_{+}^{K} \neq (b;x2P(q;b;x)}} L(q;b;x;m;l;g);$$

where P(q;b;x) represents the primal feasible region. We can simplify the Lagrangian dual from the above max-min form to a maximization problem in terms of only the dual variables by taking

advantage of the following KKT conditions.

$$\frac{\P L}{\P q} = 0 ) q \qquad a_{i=1}^{n} my_{i}f(x_{i}) + a_{l=1}^{k} gp_{l} = 0$$
$$\frac{\P L}{\P b} = 0 ) \qquad a_{i=1}^{n} my_{i} = 0$$
$$\frac{\P L}{\P x_{i}} = 0 ) C \qquad m \quad I_{i} = 08i$$

Given a primal optimal solution; b; x and a dual optimal solution; m; g, they must satisfy these KKT conditions. We use the rst condition to replace by its equivalent formula in terms on f; g, and use the other two conditions to eliminate terms involving dx in the Lagrangian function. In addition, we include the latter two conditions as constraints of f. Thel i multipliers can be eliminated by replacing the equations with inequalities. After applying simple algebra to reformulate the objective function, we obtain (3.7) as the Lagrangian dual.

In the case when a linear kernel is use( $3r_6$ ), same as the standard linear SVM, we can simplify (3.6) by treating b as an additional dimension in, then all the above derivation steps involving re no longer needed. In particular, we do not need to include the KKT condition corresponding to the term,  $a_{i=1}^{n}my_i = 0$ , in the dual formulation.

Supposen; g is an optimal solution t(3.7), the optimal q and b for general kernel and linear kernel respectively satisfy (3.8) and (3.9).

$$q = \mathbf{a}_{i=1}^{n} \mathbf{m} y_{i} f(x_{i}) \quad \mathbf{a}_{i=1}^{k} g p_{i}; b = y_{sv} h q; f(x_{sv})i:$$
(3.8)

$$q = \mathbf{a}_{i=1}^{n} \mathbf{m}_{i} \mathbf{y}_{i} \mathbf{x}_{i} \quad \mathbf{a}_{i=1}^{k} \mathbf{g}_{i} \mathbf{p}_{i}; \mathbf{b} = \mathbf{a}_{i=1}^{n} \mathbf{m}_{i} \mathbf{y}_{i}:$$
(3.9)

We observe that (3.7) is a quadratic program (QP) with a very similar structur (3td). In addition, the variablesmandg are separable in the constraints, which plays a key role in the design of our specialized methods. We also note t (3a7) and the optimal formulas for and simplify to their corresponding forms in standard SVM if gill= 0.

#### 3.2.1 Supported Fairness Constraints

We next state two examples of fairness constraints that satisfy Assumption 1. We specify these constraints for binary classi cation on the training  $d\mathbf{D}a= f(x_i; y_i; z_i)g_{i=1}^n$ . Both examples characterize statistical group fairness notions which seek fairness by eliminating certain disparity among the involved groups. For simplicity of presentation, we consider a single group dadbel0; 1g indicating whether belongs to the protected group of interest, but it is possible to extend the formulations to handle multiple groups by comparing each pair separately.

Among group fairness de nitions proposed for fair ML, three well studied notion **Dare**ographic Parity, Equalied OddsandPredictive Rate ParityIn binary classi cation, exact formulations of these fairness conditions need integer variables to denote the predicte  $\hat{g}_i | \hat{a}_i \hat{b}_i \hat{c}_i \hat{s}_i | \hat{a}_i \hat{b}_i \hat{c}_i \hat{s}_i | \hat{a}_i \hat{b}_i \hat{c}_i \hat{s}_i \hat{s}$ 

$$\frac{1}{n} \overset{n}{a}_{i=1}^{n} (z_{i} \quad \overline{z})(hq; f(x_{i})i + b) \quad e; \\ \frac{1}{n} \overset{n}{a}_{i=1}^{n} (z_{i} \quad \overline{z})(hq; f(x_{i})i + b) \quad e:$$
(3.10)

Both constraints in (3.10) satisfy Assumption 1; more speci cally, we have  $= \frac{1}{n} a_{i=1}^{n} (z_i - \overline{z}) f(x_i)$ and  $p_2 = p_1$ .

The same technique applies equalized odds namely true positive rate parity which requires  $jP(\hat{y} = 1jz = 1; y = 1)$   $P(\hat{y} = 1jz = 0; y = 1)j$  e. For notation ease,  $lgp = \frac{y_i+1}{2}$ ,  $w_i = z_i y_i^0$ , clearly  $y_i^0$ ;  $w_i \ge f 0$ ; 1g. In addition, we de  $ney = \frac{1}{n} a_{i=1}^n y_i^0$ ,  $\bar{w} = \frac{1}{n} a_{i=1}^n w_i$ .

$$\frac{1}{n} \overset{n}{a}_{i=1}^{n} (y_{i}^{0} \bar{w} \quad \bar{y} w_{i})(hq; f(x_{i})i + b) \quad e;$$

$$\frac{1}{n} \overset{n}{a}_{i=1}^{n} (y_{i}^{0} \bar{w} \quad \bar{y} w_{i})(hq; f(x_{i})i + b) \quad e$$
(3.11)

Similar to the covariance parity constraint, we conclude (Batt1) is a pair of constraintsp<sub>1</sub> =  $\frac{1}{n} a_{i=1}^{n} (y_{i}^{0} \bar{w} - \bar{y}w_{i}) f(x_{i})$  and  $p_{2} = p_{1}$ , satisfying Assumption 1.

We conclude this section by noting two limitations with this substitute technique. First, recent papers Bendekgey and Sudderth (2021); Wu et al. (2019) investigate the theoretical guarantees of these relaxed group fairness constraints and observe that they may be insufficient to provide fairness guarantees for certain rangeofBendekgey and Sudderth (2021) propose an alternative substitute strategy: use a logistic surrogate instead of the simple linear surrogate( $x_i$ ) i + b. With a logistic surrogate, we can no longer derive the exact dual fair SVM formulation, so our dual-based algorithms are not directly applicable. Second, the substitute technique does not generalize to predictive rate parity, which is conditional on the predicted labels instead of the true labels as in demographic parity and equalized odds. Formulating tractable continuous surrogates of predictive rate parity is a challenging open question in literature.

# 3.3 Fair SVM Training Algorithms

We develop specialized methods to train fair SVM via solv( $\hat{a}$ ,  $\hat{g}$ ). As our methods extend dualbased standard SVM algorithms, we begin with a review of specialized algorithms for standard SVM with a general kernel and a linear kernel.

#### 3.3.1 Training Standard SVM

### Decomposition Method: General SVM

Decomposition method solv(3.2) by decomposing the complete QP into a series of smaller QP subproblems, and relies on the Karush-Kuhn-Tucker (KKT) conditions to choose working sets, update working set variable values and determine whether an optimal solution is found. We refer to the fight (Joachims (1998)) as a standard implementation of decomposition method.

SVM<sup>light</sup> initiates with a feasible solutiom<sup>(0)</sup> and maintains a feasible throughout the procedure. In iterationt, SVM<sup>light</sup> examines whether the current solution<sup>(†)</sup> is optimal based on the KKT conditions of(3.2). The algorithm terminates and returns the optimal solution if one has been found, otherwise it proceeds to the next subproblem. In each subproblem, the complete variable spelt into a working set B, of free variables to be updated, and image tive set N, of xed variables to be temporarily treated as constants.

Supposemandq; b are respectively feasible solutions  $(\mathfrak{D}2)$  and (3.1), then they are both optimal when they satisfy the following KKT conditions.

$$\begin{split} m &= 0 \qquad ) \quad y_i(hq; f(x_i)i + b) \quad 1 \\ m &2 (0; C) \qquad ) \quad y_i(hq; f(x_i)i + b) = 1 \qquad (3.12) \\ m &= C \qquad ) \quad y_i(hq; f(x_i)i + b) \quad 1 \end{split}$$

We denote the objective function in (3.2) has the state of the stat

$$\tilde{N}_{m}h(m) = a_{i=1}^{n}Q_{ij}m_{j}$$
 1: (3.13)

Using these gradient formulas, an equivalent representation of (3.12) is:

$$m > 0$$
)  $\hat{N}_{m}h(m) + by_{i} \quad 0; m < C$ )  $\hat{N}_{i}h(m) + by_{i} \quad 0$  (3.14)

Therefore, a feasible is optimal if there exists such that the KKT condition (3.14) are satis ed. We can easily check whether a valide xists by keeping track of the following bounds:

$$\begin{split} m(m) &= \max_{i \ge I_{up}(m)} y_i \tilde{N}_m h(m); \text{ where } I_{up}(m) \coloneqq fi: m < C; y_i = 1 \text{ or } m > 0; y_i = 1g \\ M(m) &= \min_{i \ge I_{low}(m)} y_i \tilde{N}_m h(m); \text{ where } I_{low}(m) \coloneqq fi: m < C; y_i = 1 \text{ or } m > 0; y_i = 1g \end{split}$$
(3.15)

Since  $y_i 2 f 1$ ; 1g, (3.14) is satisfied when there exists (m) b M(m). In practice, decomposition algorithms typically do not require a perfect ful llment (3.14). SVM<sup>light</sup> considers the following termination criterion allowing up to violation to the KKT conditions.

If (3.16) does not hold at a givern, then KKT conditions(3.14) are violated at some. These indices are the candidates for the next working set, in other words, updating these indices may improve the solution.

In SVM<sup>light</sup>, working set selection is performed via an optimization problem 7) solves for a feasible direction that leads to the steepest descent in a rst-order approximation of the objective function.

$$\min_{d} \tilde{N}_{m}h(m)^{T}d \text{ s.t. } y^{T}d = 0; \text{ if } d_{i} : d_{i} \in 0 g j \quad q$$

$$d_{i} \quad 0 \text{ if } m_{i} = 0; d_{i} \quad 0 \text{ if } m_{i} = C; \quad 1 \quad d_{i} \quad 1; 8i$$

$$(3.17)$$

Joachims (1998) show th( $\mathfrak{A}$ .17) can be solved with a simple iterative procedure. The obtained directiond contains at most nonzero components. The variables corresponding to these nonzero indices are included in the working  $\mathfrak{B}$  that  $\mathfrak{B} = \mathfrak{f} i : d_i \mathfrak{G}$  og and its size is controlled by the choice of a first selecting a working  $\mathfrak{s} \mathfrak{S} \vee M^{light}$  proceeds to solve a new subproblem obtained from restricting the changing variables  $\mathfrak{m} \mathfrak{g} := \mathfrak{f} \mathfrak{m} : i 2 \operatorname{Bg} \mathfrak{and} \operatorname{xing} \mathfrak{m}_N := \mathfrak{f} \mathfrak{m} : j \mathfrak{G} \mathfrak{B} \mathfrak{g}$ . Note that  $\tilde{N}^2_{\mathfrak{m};\mathfrak{m}} h(\mathfrak{m})$  denotes the second order derivative  $\mathfrak{m}(\mathfrak{m})$  with respect tom;  $\mathfrak{m}$ , and we further derive  $\tilde{N}^2_{\mathfrak{m};\mathfrak{m}} h(\mathfrak{m}) = Q_{ij}$ .

$$\min_{m_B} \mathbf{\mathring{a}}_{i2B} \tilde{N}_m h(m) m + \frac{1}{2} \mathbf{\mathring{a}}_{i;j2B} \tilde{N}_{m;m}^2 h(m) m m$$
s.t.  $m_B^T y_B + m_N^T y_N = 0; 0 m C; 8i 2 B$ 
(3.18)

(3.18) is a convex quadratic program with a small number of variables, so it can be handled by off-the-shelf solvers, for instance, SVM<sup>t</sup> use a solver based on interior-point method.

To speed up the decomposition algorith  $\mathfrak{B}N/M^{light}$  implements caching to store recently used  $Q_{ij}$  in the available memory to reduce the needed number of kernel evaluations primited by reduce the problem size by temporarily neglecting variables that are likely to be bounded or C) at optimality. We note that both techniques are broadly applicable to SVM training methods. In addition, SVM<sup>light</sup> maintains an update  $\tilde{M}_m h(m)$  along with the updates on to eliminate unnecessary computations. In particular, when an iteration updates the working set variable  $\mathfrak{sr}_B$  room  $\mathbb{R}^{new}$ , we can easily update the gradient at an arbitrary index.

$$\tilde{N}_{m}h(m^{\text{pew}}) = \tilde{N}_{m}h(m) + a_{i2B}(m^{\text{pew}}_{j} m)Q_{ij}$$
(3.19)

Sequential Minimal Optimization: General SVM

Sequential minimal optimization (SMO) is a special type of decomposition method where all working sets have a xed size of two. Compared to a generic decomposition method SOVIMAS<sup>ht</sup>, SMO follows the same main steps but use simpler working set selection and variable update subproblems.

For working set selection in SMO, one approach is to a(p)/7 with q = 2. This method is studied in Keerthi et al. (2001) as the aximal Violating Pair (MVP)rule. Recall that a feasible is

not optimal if we cannot ndb that satis es all the KKT conditions i(3.14) at m A pair of indices f i; j : i 2  $I_{up}(m)$ ; j 2  $I_{low}(m)$ g leads to violation if  $\tilde{N}_m h(m) > \tilde{N}_m h(m)$ . The MVP rule picks the pair with the most violation, namely 2  $I_{up}(m)$  and j 2  $I_{low}(m)$  that respectively correspond **fm**(m) and M(m). Namely, at a givern, the MVP rule selects the following working set:

$$B = fi; jg; where i = argma_{\gamma_2} y_{i0} \tilde{N}_{m_0} h(m); j = argmin_{\gamma_2} y_{i0} \tilde{N}_{m_0} h(m): (3.20)$$

After selecting a working set, SMO proceeds to solve the subproblem of the f(3m18), which simpli es to the following QP.

Since this subproblem is a QP of two variables, its closed form optimal solution is easy to obtain. SMO applies the closed form formulas to updatem. The optimal updates

$$m_{i}^{\text{new}} = \begin{array}{ccc} 8 & 8 \\ \gtrless & 0 & \text{if } m_{i} < 0 \\ R_{i}^{\text{new}} = \begin{array}{ccc} C & \text{if } m_{i} > C \\ R_{i}^{\text{new}} = \begin{array}{ccc} C & \text{if } m_{i} > C \\ R_{i}^{\text{new}} = \begin{array}{ccc} C & \text{if } m_{i} > C \\ R_{i}^{\text{new}} = \begin{array}{ccc} C & \text{if } m_{i} > C \\ R_{i}^{\text{new}} = \begin{array}{ccc} C & \text{if } m_{i} > C \\ R_{i}^{\text{new}} = \begin{array}{ccc} C & \text{if } m_{i} > C \\ R_{i}^{\text{new}} = \begin{array}{ccc} C & \text{if } m_{i} > C \\ R_{i}^{\text{new}} = \begin{array}{ccc} C & \text{if } m_{i} > C \\ R_{i}^{\text{new}} = \begin{array}{ccc} C & \text{if } m_{i} > C \end{array} \end{array}$$

are clipped from the following unconstrained optimal solution  $m_{j}$  to t in the [0;C] bounds. Note that d > 0 is a small constant.

$$m = m + y_{i} \frac{b_{ij}}{a_{ij}}; m_{j} = m_{j} y_{j} \frac{b_{ij}}{a_{ij}};$$

$$wherea_{ij} = \begin{cases} K_{ii} + K_{jj} & 2K_{ij} \text{ if } K_{ii} + K_{jj} & 2K_{ij} > 0 \\ \vdots & \text{d if } K_{ii} + K_{jj} & 2K_{ij} = 0 \end{cases}; b_{ij} = y_{j} \tilde{N}_{m} g(m,g) y_{i} \tilde{N}_{m} g(m,g):$$
(3.22)

Dual Coordinate Descent Method: Linear SVM

The previous two methods are designed to handle general kernel functions. Among all kernel choices, the linear kernel has been studied separately in literature as the simple structure of linear SVM provides great potentials for large-scale problems. We next review the dual coordinate descent (DCD) algorithm that Hsieh et al. (2008) propose for linear SVMs. Note that this is one of the dual-based algorithms implemented in the LIBLINEAR (Fan et al. (2008)) package.

Coordinate descent is a well-known optimization technique that nds an optimal solution via coordinate-wise optimization, namely, it successively updates one variable at a time via a single-variable subproblem until reaching certain termination criteria. The DCD method in Hsieh et al. (2008) essentially apply this standard technique to s(3)ve). The algorithm initiates with a feasible

 $m^{(0)}$ . An iteration of updating  $n^{(t-1)}$  to  $m^{(t)}$  consists of inner iterations, each of which solves a single-variable optimization model to nd the optimal update along each coordinate.

Let  $m^{(i;t-1)}$  denote the solution obtained in the inner iteration, the  $m^{(0;t-1)} = m^{(t-1)}$ ,  $m^{(i;t-1)} = [m^{(t)}_1; \dots; m^{(t-1)}_1; \dots; m^{(t-1)}_n]$  for  $i = 1; \dots; n-1$ , and  $m^{(n;t-1)} = m^{(t)}$ . The i-th inner iteration solves the following optimization model to update i;t-1 to  $m^{(i+1;t-1)}$ :

$$\min_{d} h(m^{(i;t-1)} + de_i) \text{ s:t: } 0 \quad m^{(t-1)}_i + d \quad C$$
(3.23)

Here is then-dimensional unit vector with  $(a_i)_i = 1$ . (3.23) thereby nds the optimal feasible update along the direction of  $a_i$ . We further evaluate to simplify (3.23) to a quadratic program:

$$\min_{d} \frac{1}{2} Q_{ii} d^2 + \tilde{N}_{m} h(m^{(i;t-1)}) d st: 0 \quad m^{(t-1)}_{i} + d \quad C$$
(3.24)

These update subproblems also provide a convenient way to check optimal it is optimal if all n subproblems have an optimal solution dat = 0, namely, no further improvement is possible along any coordinate. We note that = 0 is an optimal solution to (3.24) if and only if the projected gradient  $\tilde{N}_m^P h(m) = 0$  for all i, where,

$$\tilde{N}_{m}^{P}h(m) = \min_{\substack{m \neq 0 \\ M_{m}}} \min_{\substack{m \neq 0 \\ M_{m}}} (m) = 0 \qquad (3.25)$$

Similar to decomposition methods, instead of seeking the exact optimality, Hsieh et al. (2008) use a relaxed stopping criteria which requires  $x_m f \tilde{N}_m^P h(m)g = \min_m f \tilde{N}_m^P h(m)g < t$  and  $max_m f \tilde{N}_m^P h(m)jg < t$ . Another implementation detail is the shrinking procedure: Hsieh et al. (2008) demonstrate that the same shrinking techniques used in decomposition methods are valid for the DCD algorithm.

### 3.3.2 Training Fair SVM

We next develop algorithms for training fair SVM with the dual formulat(3n7). The key challenge of solving (3.7) is that we need to search the optimal values of two sets of variablass dg. We observe that mandg are fully separable in the constraints, and they share the same functional format in the objective function. Given a xed, (3.7) apparently simpli es to the dual formulation(3.2) or (3.4), used in standard SVM. Given a xed, (3.7) becomes a convex quadratic program with only non-negativity constraints on the variables:ming  $d^Tg + (Am)g + \frac{1}{2}g^TPg + constant g = 0; 8I$ . Note that a quadratic program of this form is solvable with the coordinate descent method (Luo and Tseng (1993)). In fact, the dual formulation(3.4) of standard linear SVM has the same format, which allows Hsieh et al. (2008) to design the dual coordinate descent method. Drawing motivation from these observations, we propose to **(GIXP)** with separate treatment of mandg updates; more explicitly, we apply the update subroutine from a standard SVM algorithm to updatemand we design a coordinate descent based update subroutinealities. Utilizing the newg update subroutine, we extend SMO and DCD to specialized algorithms, respectively for fair SVMs with a general kernel and a linear kernel. We note that the same extension is also applicable to generic decomposition methods.

### Main Steps and Termination Criteria

Same as standard SVM algorithms, our fair SVM algorithms initiate with a feasible solution  $(m^{(0)}; g^{(0)})$ . In iterationt, we examine whether the current solution satis es the termination criteria. If the termination criteria are not reached yet, we rst apply an update iteration from a standard SVM algorithm (Section 3.3.2) to update<sup>(t 1)</sup> to  $m^{(t)}$  with  $g^{(t 1)}$  treated as x constants, then we update $g^{(t 1)}$  to  $g^{(t)}$  with  $m^{(t)}$  xed using coordinate descent (Section 3.3.2).

We rely on the KKT conditions to select the stopping criteria. (g(ett, g) denote the objective function in (3.7).

Proposition 2. The KKT conditions for fair SVMs de ned (6.6), (3.7) are:

$$\begin{array}{l} m > 0 \ ) \quad \tilde{N}_{m}g(m,g) + by_{i} \quad 0; \ m < C \ ) \quad \tilde{N}_{m}g(m,g) + by_{i} \quad 0; \\ g = 0 \ ) \quad \tilde{N}_{q}g(m,g) \quad 0; \ g > 0 \ ) \quad \tilde{N}_{q}g(m,g) = 0; \end{array}$$

$$(3.26)$$

where the gradient formulas are:

$$\tilde{N}_{m}g(m,g) = \overset{n}{\underset{j=1}{a}} Q_{ij}m_{j} + \overset{k}{\underset{l=1}{a}} A_{li}g \quad 1; \\ \tilde{N}_{g}g(m,g) = \overset{n}{\underset{j=1}{a}} A_{lj}m_{j} + \overset{k}{\underset{l\subseteq 1}{a}} P_{l} g_{0} \quad (c_{l} e_{l}): \quad (3.27)$$

Proof. We rst derive the gradient formulas for the dual objective function f(m, g). Recall that  $g(m, g) = e^{T}m d^{T}g + \frac{1}{2} \frac{m}{g} \bar{Q} \frac{m}{g}$ , where  $\bar{Q} = \frac{Q}{A^{T}} \frac{A}{P}$ : Q is the same n matrix used in the standard dual SVM with  $Q_{ij} = hy_i f(x_i); y_j f(x_j) i$ , A is a k n matrix with  $A_{ii} = h p_i; y_i f(x_i) i$ , and P is a k k with  $P_{I_i} = hp_i; p_{I_i} i$ . It is easy to compute the gradients as follows, which are exactly (3.27).

$$\tilde{N}_m g(m,g) = \underset{j=1}{\overset{n}{\overset{n}{a}}} Q_{ij} m_j + \underset{l=1}{\overset{k}{\overset{n}{a}}} A_{li} g \quad 1; \ \tilde{N}_g g(m,g) = \underset{j=1}{\overset{n}{\overset{n}{a}}} A_{lj} m_j + \underset{l=1}{\overset{k}{\overset{n}{a}}} P_{l} \circ g \circ \quad (c_l \quad e_l):$$

These gradients have the following equivalent representations when we plug in the formulas for the matrix entries.

$$\begin{split} \tilde{N}_{m}g(m,g) &= hy_{i}f(x_{i}); \overset{n}{a}_{j=1}^{n}m_{j}y_{j}f(x_{j})i \ h \ y_{i}f(x_{i}); \overset{k}{a}_{l=1}^{k}gp_{l}i \ 1; \\ \tilde{N}_{g}g(m,g) &= h \overset{n}{a}_{i=1}^{n}m_{j}y_{j}f(x_{j}); p_{l}i + hp_{l}; \overset{k}{a}_{l=1}^{k}g_{0}p_{l}oi \ (c_{l} \ e_{l}): \end{split}$$

Supposeq; b and m, g are respectively the optimal solution (**a**.6) and (3.7), then the KKT conditions of this primal-dual pair are:

$$\begin{split} m &> 0 \ ) \quad y_i(hq; f(x_i)i+b) \quad 1; \ m < C \ ) \quad y_i(hq; f(x_i)i+b) \quad 1; \\ g &= 0 \ ) \ h \ p_i; qi+q_i < q; \ g > 0 \ ) \ h \ p_i; qi+q_i = q; \end{split}$$

We recall from(3.8) that the optimal solutions satisfy =  $a_{i=1}^{n} my_i f(x_i) = a_{i=1}^{k} gp_i$ . When we replaced in the above KKT conditions with this optimal format in terms of dual variables, we obtain the format in (3.26).

$$\begin{split} & \mathfrak{m} > 0 \ ) \ y_i(h \mathring{a}_{i=1}^n \mathfrak{m} y_i f(x_i) \ \ \mathring{a}_{i=1}^k g p_i; f(x_i) i + b) \ 1 \ \ \tilde{N}_{\mathfrak{m}} g(\mathfrak{m}, g) + b y_i \ 0; \\ & \mathfrak{m} < C \ ) \ y_i(h \mathring{a}_{i=1}^n \mathfrak{m} y_i f(x_i) \ \ \mathring{a}_{i=1}^k g p_i; f(x_i) i + b) \ 1 \ \ \tilde{N}_{\mathfrak{m}} g(\mathfrak{m}, g) + b y_i \ 0; \\ & g = 0 \ ) h \ p_i; \mathring{a}_{i=1}^n \mathfrak{m} y_i f(x_i) \ \ \mathring{a}_{i=1}^k g \circ p_i \circ i + c_i < e_i \ \ \tilde{N}_g g(\mathfrak{m}, g) > 0; \\ & g > 0 \ ) h \ p_i; \mathring{a}_{i=1}^n \mathfrak{m} y_i f(x_i) \ \ \mathring{a}_{i=1}^k g \circ p_i \circ i + c_i = e_i \ \ \ \tilde{N}_g g(\mathfrak{m}, g) = 0; \end{split}$$

In fair SVM algorithms, we apply the stopping criteria from standard SVM algorithms with the only change that  $\tilde{N}g(m,g)$  replace  $\tilde{N}h(m)$ . Suppose we allow up to violation to the KKT conditions. For training a fair SVM with a general kernel, we terminate the fair SMO algorithm when the obtained solution (m,g) satis esm(m,g) M(m,g) t where

$$m(m,g) = \max_{i \ge l_{up}(m)} y_i \tilde{N}_m g(m,g); \ M(m,g) = \max_{i \ge l_{low}(m)} y_i \tilde{N}_m g(m,g):$$
(3.28)

For a linear kernel, we can further simpli(§.26) to eliminate the terms including through the augmentation trick, that is, we add an additional element equator is for all i and an additional zero element top for all I. Fair linear SVM algorithm thereby uses a similar stopping condition as the DCD method in Section 3.3.1, that is  $\tilde{N}_m^P g(m, g)g = \min_i f \tilde{N}_m^P g(m, g)g = t$ , maxif  $\tilde{N}_m^P g(m, g)jg = t$ , where the projected gradie  $\tilde{N}_m^P g(m, g)$  has an analogous denition to  $\tilde{N}_i^P h(m)$ .

Update Dual Variables for Standard SVM Constraints

In iterationt, we rst updatem<sup>(t)</sup> to m<sup>(t+1)</sup> with g<sup>(t)</sup> xed. The update requires solving a subproblem containing only the nvariables. Similar to the stopping criteria, compared to the subproblems used in standard SVM algorithms, the only change needed is to substitute the gradient de nitions with  $\tilde{Ng}(m, g^{(t-1)})$ .

For fair SVMs with a general kernel, we modify the update step in the SMO algorithm. In each iteration, we choose the working softrom mvariables by solving 3.17) with  $\tilde{N}_ih(m)$  replaced by  $\tilde{N}_mg(m,g)$ . It sufficient to use the following modified maximal violation pair rule.

$$B = fi; jg; where i = argma_{\chi_2 I_{up}(m)} \quad y_i \circ \tilde{N}_{m_0} g(m; g); j = argmin_{j^0 \ge I_{low}(m)} \quad y_j \circ \tilde{N}_{m_0} g(m; g): (3.29)$$

Then we solve for the optimal update with the following subproblem.

$$\min_{m_{B}} \hat{a}_{i2B} \tilde{N}_{m} g(m,g) m + \frac{1}{2} \hat{a}_{i;j2B} \tilde{N}_{m;m}^{2} g(m,g) m m$$
s.t.  $m_{B}^{T} y_{B} + m_{N}^{T} y_{N} = 0; 0 m C; 8i 2 B$ 
(3.30)

For fair linear SVMs, we utilize the DCD method: the outer iteration updatifig<sup>1)</sup> to m<sup>(t)</sup> consists of coordinate-wise optimization subproblems. **The** subproblem is stated below.

$$\min_{d} \frac{1}{2} Q_{ii} d^2 + \tilde{N}_{m} g(m^{(i;t-1)}; g^{(t-1)}) d st: 0 \quad m^{(t-1)}_{i} + d \quad C:$$
(3.31)

Update Dual Variables for Fairness Constraints

We now introduce the coordinate descent update**g** formables. Recall that there akefairness constraints, thus dual variables. The outer iteration of computing (b) consists of k inner iterations. For  $I = 0; \ldots; k$ , we useg<sup>(I;t 1)</sup> to denote the solution obtained from the inner iteration, then  $g^{(0;t 1)} = g^{(t 1)}, g^{(I;t 1)} = [g_1^{(t)}; \ldots; g_{t+1}^{(t)}; \ldots; g_k^{(t-1)}]$  for  $I = 1; \ldots; k$  1, and  $g^{(k;t 1)} = g^{(t)}$ . The I-th inner iteration optimizes the objective function  $g^{(m(t)}; g)$  along the direction of g with the following optimization problem where is the k-dimensional unit vector with the the element equal to 1.

$$\min_{d} g(m^{(t)}; g^{(l-1;t-1)} + dq) \text{ s.t. } g^{(t-1)} + d = 0$$
(3.32)

The closed form solution to (3.32) is stated in the following proposition.

Proposition 3. Supposen<sup>(t)</sup> is a given feasible solution in iteration for I = 1; ...; k, the optimal solution to(3.32) is  $d_i = \max d_i; g^{(t-1)}g$ , where

$$d_{l} := \frac{a_{l=1}^{n} A_{li} m_{l}^{(t)} + a_{l=1}^{k} P_{l} g_{0}^{(l-1;t-1)} (c_{l} - e_{l})}{P_{l}} = \frac{\tilde{N}_{g} g(m^{(t)}; g^{(l-1;t-1)})}{P_{l}}:$$

 $\label{eq:constraint} Therefore \underline{q}^{(t)} = \mbox{ maxf } q^{(t-1)} \quad \frac{\tilde{N}_{q} \mbox{ g}(m^{(t)}; g^{(l-1;t-1)})}{P_l}; 0g \mbox{ for all } l.$ 

Proof. The objective function of (3.32) simpli es to the following format:

$$g(m^{(t)};g^{(l-1;t-1)} + de_{l}) = \frac{1}{2} a_{i=1}^{n} m_{i}^{(t)} y_{i} f(x_{i}) a_{i=1}^{k} g_{0}^{(l-1;t-1)} p_{10} dp_{1}^{2} (c_{i} e_{i}) dp_{1}^{2}$$

We observe that the objective function is a convex functiod.dfn addition, the gradient has the following formula:

$$\begin{split} \tilde{N}_{d}g(m^{(t)};g^{(l-1;t-1)} + de_{i}) &= h p_{i}; \overset{a}{a}_{i=1}^{n} m^{(t)}_{i}y_{i}f(x_{i}) & \overset{a}{a}_{l^{Q}=1}^{k}g_{0}^{(l-1;t-1)}p_{l^{0}} dp_{l}i \quad (c_{i} e_{i}) \\ &= dP_{i} + \overset{a}{a}_{i=1}^{n}A_{li}m^{(t)}_{i} + \overset{a}{a}_{l^{Q}=1}^{k}P_{l} g_{0}^{(l-1;t-1)} \quad (c_{i} e_{i}) \\ &= dP_{i} + \tilde{N}_{a}g(m^{(t)};g^{(l-1;t-1)}) = dP_{i} d_{i}P_{i} \end{split}$$

The last two equations follow directly from the de nition  $\delta f_g g(m^{(t)}; g^{(l-1;t-1)})$  and d<sub>l</sub>. Since (3.32) contains a single bound constraint, if the stationary point of the objective function satis es the constraint, then the stationary point is the optimal solution. Namely,

$$d_i = d_i \text{ when } g^{(t-1)} + d_i = 0$$
:

Otherwise,  $ifg^{(t-1)} + d_l < 0$ , which further implies that  $\tilde{N}_g g(m^{(t)}; g^{(l-1;t-1)}) > g^{(t-1)} P_l$  0, therefore the objective function is minimized at the smallesteasible. Namely,

$$d_{l} = q^{(t-1)} when q^{(t-1)} + d_{l} < 0$$

The optimal formula fod, implies that  $g^{(t)} = \max f g^{(t-1)} + d_i$ ; 0g.

An alternative interpretation of the subproble  $\mathfrak{Bn32}$  is that in the t-th outer iteration with a given feasiblem<sup>(t)</sup>, in each inner iteration fdr= 1;:::;k, we update to x the violated g-related KKT condition atm<sup>(t)</sup> and the latest update  $\mathfrak{m}^{1,t-1}$ . This is formalized in the next proposition.

Proposition 4. Supposen<sup>(t)</sup> is a given feasible solution iteration and  $g^{(1;t-1)}$ ;...; $g^{(k;t-1)}$  is the sequence of updates generated by coordinate descent method with subp(3632) then for all I = 1;...;k,  $(m^{(t)}; g^{(l;t-1)})$  satis es the following conditions where  $g^{(t)} = g^{(l;t-1)}$  is the update from the l-th inner iteration:

$$q^{(t)} = 0 \ ) \quad \tilde{N}_g \, g(m^{(t)}; g^{(l;t-1)}) \quad 0; \ q^{(t)} > 0 \ ) \quad \tilde{N}_g \, g(m^{(t)}; g^{(l;t-1)}) = \ 0;$$

Proof. We have derived that the gradientop with respect tog is:

$$\tilde{N}_{g}g(m,g) = \mathop{a}\limits^{n}_{j=1} A_{ij}m_{j} + \mathop{a}\limits^{k}_{I_{0}=1} P_{I_{0}}g_{0} \quad (c_{I} \quad e_{I}):$$

In iterationt, given a feasible  $m^{(t)}$ , the gradien  $\tilde{N}_{g}$  g after the the inner iteration satis es:

$$\tilde{N}_{g} g(m^{(t)}; g^{(l;t-1)}) = \tilde{N}_{g} g(m^{(t)}; g^{(l-1;t-1)}) + P_{l}(q^{(t)} - q^{(t-1)}):$$

We next discuss in two separate cases. Fig(s) > 0, by the derived update formula in Proposition 3, this happens when  $\mathbf{q}^{(t)} = \mathbf{q}^{(t-1)} - \frac{\tilde{N}_{g} g(\mathbf{m}^{(t)}; \mathbf{g}^{(l-1;t-1)})}{P_{l}}$ , therefore we can conclude

$$\tilde{N}_{g}g(m^{(t)};g^{(1;t-1)}) = \tilde{N}_{g}g(m^{(t)};g^{(1-1;t-1)}) + P_{H}(g^{(t)} - g^{(t-1)}) = 0:$$

The second case  $\dot{g}^{(t)} = 0$ , which happens when  $g^{(t)} = g^{(t-1)} - \frac{\tilde{N}_{g} g(m^{(t)}; g^{(t-1;t-1)})}{P_{I}}$ . This further implies

$$\tilde{N}_{g}g(m^{(t)};g^{(1;t-1)}) = \tilde{N}_{g}g(m^{(t)};g^{(1-1;t-1)}) + P_{I}(g^{(t)} - g^{(t-1)}) = 0$$

#### Implementation techniques

We use caching to reduce the needed kernel evaluations in both fair SMO and fair DCD. In addition, both algorithms maintain an efficient update of the gradients gradients along with the variable updates. In Fair SMO, after computing in iterationt, we update all gradients once with the following formulas:

$$\tilde{N}_{m}g(m^{(t)};g^{(t-1)}) = \tilde{N}_{m}g(m^{(t-1)};g^{(t-1)}) + \mathring{a}_{j=1}^{n}Q_{ij}(m_{j}^{(t)} - m_{j}^{(t-1)}) 8i$$

$$\tilde{N}_{g}g(m^{(t)};g^{(t-1)}) = \tilde{N}_{g}g(m^{(t-1)};g^{(t-1)}) + \mathring{a}_{j=1}^{n}A_{ij}(m_{j}^{(t)} - m_{j}^{(t-1)}) 8i$$
(3.33)

Then during the coordinate descent steps updag(fnd) to  $g^{(t)}$ , we update all gradients after each inner iteration. Namely, after theth iteration compute $g^{(t)}$ , we update new gradients as below:

$$\tilde{N}_{m}g(m^{(t)};g^{(1;t-1)}) = \tilde{N}_{m}g(m^{(t)};g^{(1-1;t-1)}) + A_{li}(q^{(t)} - q^{(t-1)}) 8i$$

$$\tilde{N}_{g_0}g(m^{(t)};g^{(1;t-1)}) = \tilde{N}_{g_0}g(m^{(t)};g^{(1-1;t-1)}) + P_{l^0}(q^{(t)} - q^{(t-1)}) 8l^{0}.$$

$$(3.34)$$

In fair DCD, an iteration updates all gradients in each inner iteration molpdate with (3.33) and in each inner iteration of update with (3.34).

# 3.4 Theoretical Properties of Fair SVM Algorithms

We prove that both specialized algorithms for fair SVM asymptotically converge to a optimal solution of (3.7). In particular, we establish the stronger linear convergence property for fair DCD, which implies that fair DCD returns a solution with up toviolation to the KKT conditions, namely reaches the termination criteria, in nitely many iterations. Moreover, we show that the shrinking heuristic for standard DCD is applicable to fair DCD and guarantees a nite termination.

## 3.4.1 Asymptotic convergence of fair SMO for fair general SVMs

We observe that fair SMO algorithm is a block coordinate descent method, as each iteration nds a new feasible solution by either updating a block of size two amongneriables or a block of size one among the variables. Tseng and Yun (2010) prove the asymptotic convergence of a block coordinate descent method for linearly constrained smooth optimization when the block is chosen to provide suf cient descent in the objective value. For completeness, we state the convergence result from Tseng and Yun (2010).

Theorem 13(Theorem 4.1 Tseng and Yun (2010) iven a linearly constrained smooth optimization problem

$$\min_{x \in \mathcal{X}} f(x) \text{ s.t. } Ax = b; I \quad x \quad u; x \ge R^{n}:$$

In a block coordinate descent algorithm, iteration t consists of the following steps:

- 1. select a nonempty block<sup>(t)</sup> from all x coordinates and a symmetric mathik<sup>(t)</sup> 2 R<sup>n n</sup> with B  $_{B^{(t)}B^{(t)}}^{T}$  B  $_{B^{(t)}}^{(t)}$  B  $_{B^{(t)}}^{(t)}$  0, where B  $_{B^{(t)}}$  is a matrix whose columns form an orthonormal basis for Null(A<sub>B(t)</sub>).
- 2. updatex<sup>(t+1)</sup> = x<sup>(t)</sup> + a<sup>(t)</sup>d<sup>(t)</sup>, with a<sup>(t)</sup> > 0 and d<sup>(t)</sup> = argmin<sub>d2R<sup>n</sup></sub>f  $\tilde{N} f(x)^T d + \frac{1}{2}d^T H^{(t)}d$ : A(x<sup>(t)</sup> + d) = b; I x<sup>(t)</sup> + d u; d<sub>i</sub> = 0 8i 6 $\mathcal{B}^{(t)}$ g.

Every limit point off  $x^{(t)}$  g is an optimal solution if the following conditions are satis ed:

(a)  $f(x^{(t)} + a^{(t)}d^{(t)}) = f(x^{(t)} + \tilde{a}^{(t)}d^{(t)})$  for all t, where  $f\tilde{a}^{(t)}g$  is chosen by the Armijo rule.

(b) There exist  $\mathfrak{S} < v$  1 such that for all t,

$$\min \tilde{N} f(x)^{\mathsf{T}} d + \frac{1}{2} d^{\mathsf{T}} \operatorname{diag}(\mathsf{H}^{(t)}) d : \mathsf{A}(x^{(t)} + d) = b; \mathsf{I} \quad x^{(t)} + d \quad u; d_{\mathsf{i}} = 0 \, 8\mathsf{i} \, 6\mathcal{B}^{(t)} g \\ \operatorname{vminf} \tilde{N} f(x)^{\mathsf{T}} d + \frac{1}{2} d^{\mathsf{T}} \operatorname{diag}(\mathsf{H}^{(t)}) d : \mathsf{A}(x^{(t)} + d) = b; \mathsf{I} \quad x^{(t)} + d \quad ug:$$

$$(3.35)$$

In addition, there exist  $< \underline{l}$  is such that for all t,  $\overline{l}$  diag(H<sup>(t)</sup>)  $\underline{l}$ .

To prove the asymptotic convergence of fair SMO algorithm, it sufces to show that fair SMO satis es all the needed conditions in (Tseng and Yun, 2010, Theorem 4.1), thus the convergence result applies. We rst prove the following lemma which states the degree of improvement achieved with each update.

Lemma 1. Suppose  $(m^{(t)}; g^{(t)})g$  is the sequence generated by fair SMO. In the outer iteration suppose  $g^{(l;t)}g_{l^2[k]}$  is the sequence generated by the inner iterations updating the sequence generated by the inner iteration set of the sequence generated by the inner iterations updating the sequence generated by the inner iteration set of the sequence generated by the inner iteration set of the sequence generated by the inner iteration set of the sequence generated by the inner iteration set of the sequence generated by the inner iteration set of the sequence generated by the inner iteration set of the sequence generated by the inner iteration set of the sequence generated by the inner iteration set of the sequence generated by the inner iteration set of the sequence generated by the inner iteration set of the sequence generated by the inner iteration set of the sequence generated by the inner iteration set of the sequence generated by the inner iteration set of the sequence generated by the inner iteration set of the sequence generated by the set of the set of the set of the sequence generated by the set of the sequence generated by the set of the set

$$g(m^{(t+1)};g^{(t)}) \quad g(m^{(t)};g^{(t)}) \qquad \frac{1}{4}a_{i_t j_t} m^{(t+1)} m^{(t)^{-2}}; \qquad (3.36)$$
$$g(m^{(t+1)};g^{(l;t)}) \quad g(m^{(t+1)};g^{(l-1;t)}) \qquad \frac{1}{2}\mathsf{P}_{\mathsf{I}}(g^{l;t} \quad g^{l-1;t})^{2}; \tag{3.37}$$

where  $f_{i_t}$ ;  $f_{t_t}$  is the working set used in iteration t and  $f_{j_t} = f(x_{i_t}) - f(x_{j_t})^2$ . Moreover,

$$g(m^{(t+1)};g^{(t+1)}) \quad g(m^{(t)};g^{(t)}) \qquad \frac{1}{4}a_{i_t j_t} \quad m^{(t+1)} \quad m^{(t)} \stackrel{2}{=} \frac{1}{2} \overset{k}{\overset{k}{a}} P_{I}(g^{(t+1)} \quad g^{(t)})^2:$$
(3.38)

Proof. Suppose fair SMO choosés;  $j_tg$  as the working set in iteration For notation ease, we denote(i; j) = ( $i_t$ ;  $j_t$ ), then the optimal updates without the bound constraints are:

$$m_i = m_i^{(t)} + y_i d = m_i^{(t)} + y_i \frac{b_{ij}}{a_{ij}}; m_j = m_j^{(t)} \quad y_j d = m_j^{(t)} \quad y_j \frac{b_{ij}}{a_{ij}}$$

With the additional constraints  $m_i^{(t)} + y_i d$  C and  $m_j^{(t)} + y_j d$  C, we suppose d U is the feasible domain of. Then the optimal feasible updates are:

$$m_i^{(t+1)} = m_i^{(t)} + y_i \bar{d}; m_j^{(t+1)} = m_j^{(t)} \quad y_j \bar{d}; \text{ where} \bar{d} := \min \{\frac{b_{ij}}{a_{ij}}; Ug:$$

Recall from the working set selection rule; 0, so we have  $\bar{d}$  0. Since the other indices are not modilised in iteration, we conclude  $m^{(t+1)}$   $m^{(t)}^2 = 2d^2$ . Using the denition of g(m, g), we have:

$$\begin{split} g(\mathbf{m}^{(t+1)}; \mathbf{g}^{(t)}) &= (\frac{1}{2} \mathsf{K}_{ii} + \frac{1}{2} \mathsf{K}_{jj} - \mathsf{K}_{ij}) \vec{d^2} + (y_i \tilde{\mathsf{N}}_{\mathfrak{m}} g(\mathbf{m}^{(t)}; \mathbf{g}^{(t)}) - y_j \tilde{\mathsf{N}}_{\mathfrak{m}} g(\mathbf{m}^{(t)}; \mathbf{g}^{(t)})) \vec{d} \\ &= \frac{1}{2} a_{ij} \vec{d^2} - b_{ij} \vec{d} = (a_{ij} \vec{d} - b_{ij}) \vec{d} - \frac{1}{2} a_{ij} \vec{d^2} \\ &= \frac{1}{2} a_{ij} \vec{d^2} = -\frac{a_{ij}}{4} - \mathbf{m}^{(t+1)} - \mathbf{m}^{(t)} - \frac{2}{2} : \end{split}$$

The last inequality follows from  $a_{ij}d^{\bar{j}}b_{ij} = 0$  and  $\bar{d} = 0$ .

Now, to prove(3.37), we again begin with the denition of to simplify the function difference formula. Recall from Proposition  $\mathfrak{A}_{l} = \max \left\{ \begin{array}{c} \tilde{N}_{g} \, g(m^{(t+1)}; g^{(l-1;t)}) \\ P_{li} \end{array} \right\}; \quad g^{(t)}g.$ 

$$\begin{split} g(\mathbf{m}^{(t+1)}; g^{(l;t)}) & g(\mathbf{m}^{(t+1)}; g^{(l-1;t)}) = \frac{1}{2} \mathsf{P}_{\mathsf{I}}(\mathsf{d}_{\mathsf{I}})^2 \quad (\mathsf{e}_{\mathsf{I}} \quad \mathsf{e}_{\mathsf{I}} + \overset{\mathsf{n}}{\overset{\mathsf{n}}{\underset{\mathsf{I}=1}{\overset{\mathsf{n}}{\mathsf{m}}}} \mathfrak{m}_{\mathsf{I}_{\mathsf{I}}\mathsf{I}} + \overset{\mathsf{k}}{\overset{\mathsf{n}}{\underset{\mathsf{I}\cong 1}{\overset{\mathsf{n}}{\mathsf{m}}}} \mathfrak{g}_{\mathsf{I}} \mathsf{P}_{\mathsf{I}}) \mathsf{d}_{\mathsf{I}} \\ & = (\mathsf{P}_{\mathsf{I}}\mathsf{d}_{\mathsf{I}} + \tilde{\mathsf{N}}_{\mathsf{g}} \mathfrak{g}(\mathbf{m}^{(t)}; \mathfrak{g}^{(l-1;t)})) \mathsf{d}_{\mathsf{I}} \quad \frac{1}{2} \mathsf{P}_{\mathsf{I}}(\mathsf{d}_{\mathsf{I}})^2 \end{split}$$

In the case when  $\mathbf{d}_{l} = \frac{\tilde{N}_{q_{l}} g(\mathbf{m}^{(t+1)}; g^{(l-1;t)})}{\mathsf{P}_{l}}$ , the RHS in the above formula simpli es to  $\frac{1}{2}\mathsf{P}_{l} (d)^{2}$ . In the other case when  $\mathbf{d}_{l} = \mathbf{q}^{(t)} = \mathbf{q}^{(t)$ 

After all k inner iterations, we obtaig (t+1). As we have shown, for each  $1; \ldots; k$ , updating the l-th coordinate of greduces the value of by at least  $\frac{1}{2}P_{l}(d_{l})^{2}$ . Therefore, considering all the updates,

$$g(m^{(t+1)}; g^{(t+1)}) \quad g(m^{(t)}; g^{(t)}) = g(m^{(t+1)}; g^{(k+1)}) \quad g(m^{(t+1)}; g^{(k-1)}) + g(m^{(t+1)}; g^{(t)}) + g(m^{(t+1)}; g^{(t)}) + g(m^{(t+1)}; g^{(t)}) = g(m^{(t+1)}; g^{(t)}) \quad g(m^{(t+1)}; g^{(t)}) + g(m^{(t+1)}; g^{(t)}) = g(m^{(t)}; g^{(t)}) = \frac{1}{2}P_{kk}(d_{l})^{2} \quad \frac{1}{2}P_{k-1;k-1}(d_{l})^{2} \quad \frac{1}{2}P_{11}(d_{l})^{2} \quad \frac{1}{4}a_{i_{t}j_{t}} \quad m^{(t+1)} \quad m^{(t)} \quad g^{(t)} = \frac{1}{2}a_{i_{t}j_{t}}^{k} \quad m^{(t+1)} \quad m^{(t)} \quad g^{(t)} = \frac{1}{2}a_{i_{t}j_{t}}^{k} \quad g^{(t)} = \frac{1}{2}a_{i_{t}j_{t}}^{k} \quad g^{(t+1)} = \frac{1}{2}a_{i_{t}j_{t}}^{k} \quad g^{(t)} = \frac{1}{2}a_{i_{t}j_{t}}^{k} \quad g^{(t+1)} = \frac{1}{2}a_{i_$$

Theorem 14. Suppose (m<sup>(t)</sup>; g<sup>(t)</sup>)g is the in nite sequence generated by fair SMO with tolerance t = 0, a limit point ( $\bar{m}, \bar{g}$ ) of f ( $m^{(t)}; g^{(t)}$ ) g an optimal solution to (3.7).

Proof. Fair SMO solves the dual fair SVM formulation (6.7), where the objective function (m,g) is smooth, and the constraint matrix is a (n + k) matrix with y<sub>i</sub> in the rst n indices and in the remaining k indices. In iteration, the working block is  $B^{(t)} = fm_{t}; m_{t}g \text{ or } B^{(t)} = fgg$ . Then along the descent directiond<sup>(t)</sup>, fair SMO solves the subproble(8.30) or (3.32) to nd the optimal step sizea<sup>(t)</sup>. In addition, fair SMO usels  $\vec{P} = \vec{Q} = \begin{bmatrix} Q & A \\ A^T & P \end{bmatrix}$ , the Hessian of the objective function for all t. In iterationt, if  $\vec{B}^{(t)} = f \vec{m}_t$ ;  $\vec{m}_t g$ , then  $\vec{H}_{\vec{B}^{(t)}\vec{B}^{(t)}} = \begin{bmatrix} Q & A \\ A^T & P \end{bmatrix}$ , the Hessian of the objective function for all t.  $\vec{Q}_{i_t j_t} = \begin{bmatrix} Q & A \\ P & P \end{bmatrix}$ , the Hessian of the objective function for all t.  $\vec{Q}_{i_t j_t} = \begin{bmatrix} P & P & P \\ Q_{i_t j_t} & Q_{j_t j_t} & Q_{j_t j_t} \end{bmatrix}$  and  $\vec{B}_{\vec{B}^{(t)}} = \begin{bmatrix} P & P & P \\ P & P & P \end{bmatrix}$ . If  $\vec{B}^{(t)} = \vec{D}_t$ 

f g g, then  $H_{B^{(t)}B^{(t)}}^{(t)} = P_{I}$  and  $B_{B^{(t)}} = 1$ . In both cases, we derive that  $T_{B^{(t)}}^{T} H_{B^{(t)}B^{(t)}}^{(t)} B_{B^{(t)}}$  0. Together with the de nitions ofd<sup>(t)</sup>, we have veri ed that fair SMO is a block coordinate descent algorithm as stated in Theorem 13.

We next show that (m<sup>(t)</sup>; g<sup>(t)</sup>); B<sup>(t)</sup>; H<sup>(t)</sup>; d<sup>(t)</sup>; a <sup>(t)</sup>g from fair SMO satis es Theorem 13(a) and (b). Since fair SMO applies a minimization rule to rad<sup>(t)</sup>, which is at least as well as the Armijo rule, we immediately conclude (a) is satis ed. To verify (b), we observe that

$$g(\mathbf{m}^{(t+1)}; \mathbf{g}^{(t)}) \quad g(\mathbf{m}^{(t)}; \mathbf{g}^{(t)}) = \tilde{N}_{\mathbf{m}_{B^{(t)}}} g(\mathbf{m}^{(t+1)}; \mathbf{g}^{(t)})^{\mathsf{T}} \mathbf{d}^{(t)} + \frac{1}{2} \mathbf{d}^{(t)}^{\mathsf{T}} \operatorname{diag}(\mathbf{H}^{(t)}) \mathbf{d}^{(t)};$$
  
$$g(\mathbf{m}^{(t+1)}; \mathbf{g}^{(1,t)}) \quad g(\mathbf{m}^{(t+1)}; \mathbf{g}^{(1-1;t)}) = \tilde{N}_{g} g(\mathbf{m}^{(t+1)}; \mathbf{g}^{(1-1;t)})^{\mathsf{T}} \mathbf{d}^{(t)} + \frac{1}{2} \mathbf{d}^{(t)}^{\mathsf{T}} \operatorname{diag}(\mathbf{H}^{(t)}) \mathbf{d}^{(t)};$$

Therefore, Lemma 1 implies that for all the LHS of (3.35) is negative. Moreover, the RHS of (3.35) is no smaller than the LHS, as the maximal objective descent when allowing all indices to change is at least as large as the maximal descent where the second exists 0< v 1 such that (3.35) holds for all Lastly, by de nition of H<sup>(t)</sup>, we easily conclude that diag( $H^{(t)}$ ) <u>I</u> with  $\overline{I}$  = maxf  $Q_{ii}$ ;  $P_{I}g$  and <u>I</u> = minf  $Q_{ii}$ ;  $P_{I}g$ . Ī

Therefore, Theorem 13 implies that a limit  $po(\bar{mt}, \bar{g})$  of  $f(m^{(t)}; g^{(t)})g$  is an optimal solution of to (3.7).

Another useful property is to show the nite termination of fair SMO in this general case. Extending the nite termination proofs of standard SMO turns out to be insuf cient for this purpose: Chen et al. (2006) prove that standard SMO terminates in nite iterations by using the continuity of  $\tilde{N}_m h(m)$  and a counting argument to show that SMO will run outmindices to update. In fair SMO, the counting technique is invalid because the grad  $\tilde{N}_m g(m,g)$  are updated with both and g value changes. In Tseng and Yun (2010), besides the asymptotic convergence result used above to establish the asymptotic convergence of fair SMO, they also proved linear convergence for a working set selection utilizing second order informat  $\tilde{N}_{R;m} h(m^{(t)})$ , which is computationally more costly than the maximal violating pair rule we use but is more efficient in reaching the optimal solution. One possible future step is to investigate fair SMO with this alternative working set selection rule, and utilize the linear convergence result in Tseng and Yun (2010) to establish stronger convergence guarantees, including nite termination guarantees.

#### 3.4.2 Linear convergence of coordinate descent for fair linear SVMs

Our algorithm for fair linear SVM is a standard coordinate descent method, with each iteration updating one variable in the steepest descent direction. In fact, our algorithm can be viewed as a direct extension of the coordinate descent method that Hsieh et al. (2008) proposed for training standard linear SVMs. We therefore follow the proof techniques used in Hsieh et al. (2008) to establish theoretical properties of our algorithm.

Theorem 15. Suppose  $m^{(t)}$ ;  $g^{(t)}g$  is an in nite sequence generated by the coordinate descent algorithm for fair linear SVMs, then the sequence globally converges to an optimal solution with a convergence rate that is at least linear. Namely, there exists a < 1 and an iteration  $t_0$  such that  $g(m^{(t+1)}; g^{(t+1)}) = g(m; g) = a(g(m^{(t)}; g^{(t)}) = g(m; g))$  for all  $t = t_0$ .

Proof. In Luo and Tseng (1992), they proved the linear convergence of coordinate descent method for the following convex optimization problems with 2 [ $\pm$ ;  $\pm$ ) and U<sub>i</sub> 2 ( $\pm$ ;  $\pm$ ]:

$$\min_{x} f(Ex) + b^{\mathsf{T}}x \text{ s.t. } L_i \quad x_i \quad U_i:$$
(3.39)

The linear convergence holds when the following conditions are satis ed (ma)s no zero columns; (b) The set of optimal solutions (3.39) is nonempty; (c)f is strongly convex and twice differentiable everywhere. We next show that the dual fair linear SVM problem satis es all these conditions. Recall from (3.7), fair linear SVM has the following dual formulation:

h i h  $i_{T}$  h  $i_{2}$ In the notation of (3.39),  $E = y_1 x_1 ::: y_n x_n p_1 ::: p_1 and f(E m g) = \frac{1}{2} E m g$ .

We observe that (a) holds by our assumption that there is no alkzer $\varphi_l$  as both are trivial cases that can be eliminated from consideration. (c) holds by the format **b** addition, (b) can be shown with strong duality. In the primal problemin<sub>q;b;x</sub> f  $\frac{1}{2}$ kqk<sup>2</sup> + Cå  $_{i=1}^{n}$ x<sub>i</sub> : x<sub>i</sub> 1 y<sub>i</sub>(q<sup>T</sup>x<sub>i</sub> + b);x<sub>i</sub> 0 8i;q<sup>T</sup>p<sub>l</sub> e<sub>i</sub> 8lg, we can ndq;b andx > 0 that satis es all constraints as strict inequalities, namely, the Slater's condition is satis ed. Therefore, strong duality holds for fair SVM with a linear kernel, implying that the dual program has a nonempty optimal solution set and the optimal dual value is equal to the optimal primal value.

Therefore, by (Luo and Tseng, 1992, Theorem 2.1), coordinate descent method on fair linear SVMs converges linearly to an optimal solution.

#### 3.4.3 Finite termination of coordinate descent with shrinking for fair linear SVMs

The standard DCD from Hsieh et al. (2008) adopt the same shrinking technique used in standard SMO, that is, remove from the current set of active indices  $\dot{m}^{(t)} = 0$  or  $\dot{m}^{(t)} = C$ . In fair DCD, since we have both and g variables, we can shrink the variables at bounds, namely 0 or  $\dot{m} = C$  and g = 0, from both sets. We will show that fair DCD with this shrinking heuristic reaches the KKT based termination criteria in nite iterations. We begin by observing some properties of the optimal solutions to fair linear SVM, and the objective descent in each iteration.

Theorem 16. Let m; g be the limit point of  $m^{(i;t)}$ ;  $g^{(l;t)}$ g generated by coordinate descent algorithm.

- (a) If  $m_i = 0$  and  $\tilde{N}_m g(m; g) > 0$ , then there exists such that  $t_i$ , 8s,  $m_i^{st} = 0$ .
- (b) If m = C and  $\tilde{N}_m g(m; g) < 0$ , then there exists such that  $t_i$ , 8s,  $m^{s,t} = C$ .
- (c) If g = 0 and  $\tilde{N}_g g(m; g) > 0$ , then there exists such that  $t_I$ ,  $8m, g^{m,t} = 0$ .
- (d)  $\lim_{t! \neq max} \max \tilde{N}_{m}^{P}g(m^{(i;t)};g^{(0;t)}); \max \tilde{N}_{g}^{P}g(m^{(n;t)};g^{(1;t)})g$ =  $\lim_{t! \neq minf min_{i}} \tilde{N}_{m}^{P}g(m^{(i;t)};g^{(0;t)}); \min_{l} \tilde{N}_{q}^{P}g(m^{(n;t)};g^{(l;t)})g = 0.$

Proof. From the global convergence proved in Theorem 15g is an optimal solution to dual fair linear SVM. By strong duality and optimality conditions, =  $a_{i=1}^{n} m y_i \frac{x_i}{1} = a_{i=1}^{k} g \frac{p_i}{0}$  is an optimal primal solution. Moreover, is the unique primal optimal solution since the primal problem has a strictly convex objective function. For the sake of this proof, we re-assign the indices of the in nite sequence  $m^{(i;t)}$ ;  $g^{(1;t)}g$  as  $m^{(t)}$ ;  $g^{(t)}g$ , where updating  $m^{(t)}$ ;  $g^{(t)}$  to  $m^{(t+1)}$ ;  $g^{(t+1)}$  corresponds to solving the sub-problem at one coordinate (eithermicor g). Let  $q^{(t)} = a_{i=1}^{n} m_i^{(t)} y_i \frac{x_i}{1} = a_{i=1}^{k} g^{(t)} \frac{p_i}{0}$ . From the global convergence to optimality, we have

$$\lim_{t^{!}} \sum_{i=1}^{n} m_{i}^{(t)} y_{i} \frac{x_{i}}{1} = \sum_{l=1}^{k} q^{(t)} \frac{p_{l}}{0} = \lim_{t^{!}} q^{(t)} = q :$$

This further implies

$$\lim_{t! \neq i} \tilde{N}_{m}g(m^{(t)};g^{(t)}) = \lim_{t! \neq i} y_{i}hq^{(t)}; \quad \begin{matrix} \\ x_{i} \\ 1 \end{matrix} i \quad 1 = y_{i}hq \ ; \quad \begin{matrix} \\ x_{i} \\ 1 \end{matrix} i \quad 1 = \tilde{N}_{m}g(m;g); \quad (3.40)$$

#

$$\lim_{t_{l}} \tilde{N}_{g} g(m^{(t)}; g^{(t)}) = \lim_{t_{l}} h q^{(t)}; \begin{array}{c} p_{l} \\ p_{l} \\ 0 \end{array} i + d_{l} = h q ; \begin{array}{c} p_{l} \\ p_{l} \\ 0 \end{array} i + d_{l} = \tilde{N}_{g} g(m; g): \quad (3.41)$$

There exists<sub>i</sub> such that for all ti.

$$\tilde{N}_{m}g(m;g) > 0! \quad \tilde{N}_{m}g(m^{(t)};g^{(t)}) > 0; \quad \tilde{N}_{m}g(m;g) < 0! \quad \tilde{N}_{m}g(m^{(t)};g^{(t)}) < 0:$$
(3.42)

Similarly, there exists such that for alt  $t_{\rm I}$ ,

$$\tilde{N}_{g}g(m;g) > 0! \quad \tilde{N}_{g}g(m^{(t)};g^{(t)}) > 0; \quad \tilde{N}_{g}g(m;g) < 0! \quad \tilde{N}_{g}g(m^{(t)};g^{(t)}) < 0:$$
(3.43)

When we update  $f^{(t)}$ ;  $g^{(t)}$  to  $m^{(t+1)}$ ;  $g^{(t+1)}$  by changing m, the KKT condition of the subproblem implies:

$$\tilde{N}_{m}g(m^{(t+1)};g^{(t+1)}) > 0! \quad m^{(t+1)} = 0; \ \tilde{N}_{m}g(m^{(t+1)};g^{(t+1)}) < 0! \quad m^{(t+1)} = C:$$
(3.44)

These relations imply, after  $t_i$ , m is always0 or alwaysC. We also note that wit $\tilde{N}_m g(m^{(t+1)}; g^{(t+1)}) \in$ 0 for all t t<sub>i</sub>, it is impossible that 2 (0;C), because it will contradict the optimality condition. Therefore,(3.42)and(3.44)imply (a) and (b). To obtain the exact statements, we need to switch back to the original indices m<sup>(s;t)</sup>g wheres 2 f 1;...;ng.

Similarly, when updating  $m^{(t)}$ ;  $g^{(t)}$  to  $m^{(t+1)}$ ;  $g^{(t+1)}$  changes g, the KKT condition of the subproblem implies:

$$\tilde{N}_{g}g(m^{(t+1)};g^{(t+1)}) > 0! \quad m^{(t+1)} = 0:$$
 (3.45)

After t  $t_i$ , g is always0. We now conclude (c) from (3.43) and (3.45), and switching back to the index g<sup>(m;t)</sup> wherem 2 f 1; ...; kg.

Moreover, (3.42) and (3.43) imply that fort  $t_i$ , if  $\tilde{N}_m g(m;g) \in 0$ , then  $\tilde{N}_m^P g(m^{(t)};g^{(t)}) = 0$ . The other cas $\tilde{\mathbf{A}}_{m}g(m;g) = 0$  directly implies  $\tilde{N}_{m}^{P}g(m^{(t)};g^{(t)}) = 0$ . Similarly, fort  $t_{l}$ , we conclude that  $\tilde{N}_{\alpha}^{P}g(m^{(t)};g^{(t)}) = 0$ . Changing back to the original indices, we have shown (d).

Lemma 2. Supposé (m<sup>(t)</sup>; g<sup>(t)</sup>) g is the sequence generated by coordinate descent algorithm. In the outer iteration t, f  $m^{(i;t)}$ ;  $g^{(l;t)}$  g is the sequence generated by the inner iterations. Then

$$g(m^{(i;t)};g^{(0;t)}) \quad g(m^{(i-1;t)};g^{(0;t)}) \qquad \frac{1}{2}Q_{ii}(m^{(i;t)},m^{(i-1;t)})^2; \tag{3.46}$$

$$g(m^{(n;t)};g^{(l;t)}) \quad g(m^{(n;t)};g^{(l-1;t)}) \qquad \frac{1}{2}\mathsf{P}_{\mathsf{I}}(g^{(l;t)} \quad g^{(l-1;t)})^2:$$
(3.47)

Proof. When updating  $m^{(i-1;t)}; g^{(0;t)}$  to  $(m^{(i;t)}; g^{(0;t)})$ , we solve the subproblem (3.31). From the KKT conditions of the subproblem, we have

$$\tilde{N}_{m}g(m^{(i;t)};g^{(0;t)}) > 0! \quad m^{(i-1;t)}_{i} \quad m^{(i;t)}_{i} \quad 0; \quad \tilde{N}_{m}g(m^{(i;t)};g^{(0;t)}) < 0! \quad m^{(i-1;t)}_{i} \quad m^{(i;t)}_{i} \quad 0:$$

Therefore,

$$g(\mathbf{m}^{(i;t)}; \mathbf{g}^{(0;t)}) \quad g(\mathbf{m}^{(i-1;t)}; \mathbf{g}^{(0;t)}) = \tilde{N}_{m}g(\mathbf{m}^{(i;t)}; \mathbf{g}^{(0;t)})(\mathbf{m}^{(i;t)} - \mathbf{m}^{(i-1;t)}) - \frac{1}{2}Q_{ii}(\mathbf{m}^{(i;t)} - \mathbf{m}^{(i-1;t)})^{2} \\ - \frac{1}{2}Q_{ii}(\mathbf{m}^{(i;t)} - \mathbf{m}^{(i-1;t)})^{2}:$$

Next, when updating (n;t);  $g^{(l-1;t)}$ ) to  $(m^{(n;t)}; g^{(l;t)})$ , we solve the subproble (3.32). From the KKT conditions of the subproblem, we have

$$\tilde{N}_{g}g(m^{(n;t)};g^{(1;t)}) > 0! \quad q^{(1;t)} = 0; \ \tilde{N}_{g}g(m^{(n;t)};g^{(1;t)}) = 0! \quad q^{(1;t)} \quad 0:$$

Therefore,

$$g(\mathbf{m}^{(n;t)}; \mathbf{g}^{(l;t)}) \quad g(\mathbf{m}^{(n;t)}; \mathbf{g}^{(l-1;t)}) = \tilde{N}_{g} g(\mathbf{m}^{(n;t)}; \mathbf{g}^{(l;t)}) (\mathbf{g}^{(l;t)} \quad \mathbf{g}^{(l-1;t)}) \quad \frac{1}{2} \mathsf{P}_{\mathsf{I}} (\mathbf{g}^{(l;t)} \quad \mathbf{g}^{(l-1;t)})^{2} \\ \frac{1}{2} \mathsf{P}_{\mathsf{I}} (\mathbf{g}^{(l;t)} \quad \mathbf{g}^{(l-1;t)})^{2} :$$

We next prove the nite termination of coordinate descent with the selected shrinking heuristic amongmvariables on fair linear SVM.

Theorem 17. Algorithm 5 terminates in nite iterations.

Proof. For the sake of contradiction, we assume Algorithm 5 does not terminate in nite iterations. Then the algorithm will generate an in nite sequerfor  $g^{(t)}$ ;  $g^{(t)}$ : t = 1; 2; ::: g. Throughout the algorithm, at an iteration, if the active indices in 1 and all theg indices satisfy the stopping criteria M m t; jMj t; jmj t, we will un-shrinkl back to allmindices and go through and in the next outer iteration (with + k inner iterations). We collect all these outer iterations with un-shrinking as a subsequent  $f^{(t)}; g^{(t)}g_{R}$ .

We rst claim that  $fm^{(t)}; g^{(t)}g_R$  must be in nite. Assumé  $m^{(t)}; g^{(t)}g_R$  is nite, then we consider the iterations after the last iteration in this set. Theorem 16(d) implies that these future iterations will reach the stopping condition  $fm^{(t)}$  m t; jMj t; jmj t at some point. Thus,  $m^{(t)}; g^{(t)}g$  will be a nite sequence, contradicting our assumption that Algorithm 5 does not terminate.

Next, consider a subsequer  $\bar{\mathbf{k}}e$  R such that  $m^{(t)}; g^{(t)}: t 2 \bar{R}g$  converges to m; g). Lemma 2 has shown that  $g(m^{(i;t)}; g^{(1;t)})g$  with i 2 f 1;...; ng and 2 f 1;...; kg is decreasing. From the proof

of Theorem 13, we have th $\mathbf{ag}(\mathbf{m}^{(i;t)}; \mathbf{g}^{(l;t)})\mathbf{g}$  is lower bounded by the optimal dual value. Therefore, we conclude

$$\lim_{t \ge \bar{R}; t!} g(m^{(i;t)}; g^{(0;t)}) \quad g(m^{(i-1;t)}; g^{(0;t)}) = 0; \ \lim_{t \ge \bar{R}; t!} g(m^{(n;t)}; g^{(1;t)}) \quad g(m^{(n;t)}; g^{(1-1;t)}) = 0;$$

Moreover, taking the limits of (3.46) and (3.47), we have

$$\lim_{t \ge \bar{R}; t!} (m^{(1;t)}; g^{(0;t)}) = \lim_{t \ge \bar{R}; t!} (m^{(1-1;t)}; g^{(0;t)}) = \lim_{t \ge \bar{R}; t!} (m^{(n;t)}; g^{(1;t)}) = \lim_{t \ge \bar{R}; t!} (m^{(n;t)}; g^{(1-1;t)})$$
$$= \lim_{t \ge \bar{R}; t!} (m^{(t+1)}; g^{(t+1)}) = (m; g):$$

We claim that the limit poin(m;g) must be an optimal dual solution. Assume it is not optimal, then there exists an indexwhere we can solve the corresponding subproblem and obtain an optimal descent such that  $g((m;g) + de_s) < g(m;g)$ . Note that the variable at indexcan be either or g. In each outer iteration, the inner iteration at this indexsolves the subproble ( $\mathfrak{B}.31$ ) or (3.32) to optimality. Suppose the variable at index m, then  $g((m^{(i;t)};g^{(0;t)}) + de_s) = g(m^{(i+1;t)};g^{(0;t)})$ . Taking the limit on both sides gives  $(m;g) + de_s) = g(m;g)$ . We can derive the same inequality wheng is the variable at index. Therefore, we have reached a contradiction. So the desired  $\mathfrak{A}$  an optimal dual solution must be true.

We have shown that any limit point  $\delta fn^{(t)}; g^{(t)}g_R$  is a dual optimal solution. By Lemma 2, the objective values are decreasing throughout all iterations  $m_{i}^{(t)}; g^{(0;t)}g$  for any i and  $fn^{(n;t)}; g^{(1;t)}g$  for any I also converge to a dual optimal solution. We can apply the proof techniques used in Theorem 16 to conclude that  $\tilde{N}^P g(m^{(i;t)}; g^{(0;t)})$  and  $\tilde{N}^P g(m^{(n;t)}; g^{(1;t)})$  globally converges to. Therefore, the stopping conditions M m t; jMj t; jmj t will be reached by Algorithm 5 in nite iterations.

# 3.5 Numerical Experiments

We implement fair SMO and fair DCD to train SVMs constrained with covariance p(arity) and/or true positive rate parity(3.11) constraints. We apply fair SMO to nonlinear SVMs with a Gaussian RBF kernel, and fair DCD to linear SVMs. Our main goal is to demonstrate the practical potentials of these specialized fair SVM algorithms. To this end, we compare the runtime of fair SVM algorithms against their corresponding standard algorithms without fairness constraints. In addition, we also compare the runtime of both fair SVM algorithms with an off-the-shelf quadratic program solver.

Both algorithms are coded in C++ using Visual Studio 16.11.10. In addition, we use the CVXOPT QP solver with the default setting in Python as the off-the-shelf solver. All experiments are conducted in Windows 10 Pro 64-bit on a machine with Intel(R) Core(TM) i7-7600U CPU @ 2.80GHz processors and 24 GB of RAM.

#### 3.5.1 Datasets

We run experiments on two real-world datasets that are suitable for training and testing fair classi cation algorithms: the German credit dataset Hofmann (1994) and the Adult income dataset adu (1996). As our implementations are based on LIBSVM/LIBLINEAR, we use both datasets in LIBSVM format provided in Chang and Lin (2011). The German credit dataset consisted of datasets with 24 features (transformed from 0 attributes), and the laby indicates whether a person has a good (1) or bad  $\langle y = -1 \rangle$  credit risk. For this dataset, we choose the features of housing status as the group attribute: z = 1 represents 'rent', and z = 0 represents 'own' or 'for free'; we respectively refer to them as 'renters' and 'homeowners'. In the Adult income dataset, each point 20 test (transformed from 24 attributes), and the laby irepresents whether a person has income a 50,000 dollars (y = 1) or not (y = -1). We choose gender as the sensitive feature, and assignfor women and z = 0 for men. From this dataset, we use the 'a4a' instance with 4781 data points.

#### 3.5.2 Implementation Details

We implement fair SMO algorithm by modifying the SMO implementation in LIBSVM (Chang and Lin (2011)), and fair DCD by modifying the standard DCD implemented in LIBLINEAR (Fan et al. (2008)). For both algorithms, the main modi cations include de ning functions to compute  $\bar{Q}$  (standard SVM only requires computing the subma@) adding the gupdate subroutine and replacing all gradient formulas used in the implementation. In addition, we keep the same caching operations provided in the standard algorithms, and extend the shrinking steps in LIBLINEAR to handle bottmandg.

On both datasets, we test fair SMO for a RBF kernel and fair DCD for a linear kernel. We consider three fairness settings: (a) demographic parity with covariance parity con\$tail0)s(b) equalized odds with true positive parity constrai(3ts11), and (c) both conditions wit(8.10) and (3.11). For simplicity, we use the same threshelds all fairness constraints in one instance. To compare performances at fairness constraints of different strengths, we2test 0:01; 0:05; 0:1; 1g in each set of experiments. Since it is well known that the training performance of SVM is in uenced by hyper-parameters, for linear SVM instances, we test different values for the penalty parameter C 2 f 0:1; 1; 10g. For nonlinear SVMs with a RBF kernel, we @= 1 and use different values in the kernel functions 2 f 0:02; 0:1; 1g. In each experiment instance, we use 5-fold cross validation with a 80/20 train-test split and report the average statistics for accuracy and fairness.

#### 3.5.3 Experiment Results

We focus on the training time, namely the runtime of an algorithm before returning an optimal solution, as the main performance measure to compare the computational costs of specialized algorithms with and without fairness constraints, and off-the-shelf QP solver. For completeness, we also report accuracy and fairness statistics to verify that the tested fair SVM models provide reasonable





performances. We provide brief discussions about accuracy and fairness performances on both datasets in the captions of Tables 3.1-3.12.

#### Our Algorithms vs. Solver

In Figs. 3.1a and 3.1b, we compare the runtime of fair SMO versus the CVXOPT solver for training a fair non-linear SVM with a RBF kernel on both datasets. Figs. 3.2a and 3.2b provide the same comparisons between fair DCD versus the CVXOPT solver for training fair linear SVMs. We show the runtime separately for different fairness settings and thresholds. We observe that fair SMO noticeably outperforms the solver on all instances. In fact, when using the CVXOPT solver, we pre-compute and input the complete kernel matr**Q** to the solver, whereas we only input the feature vectors to fair SMO; therefore, the effective advantages of fair SMO are greater than those shown in illustration. For training linear SVMs, we observe that fair DCD has a shorter runtime than the CVXOPT solver on all but the German dataset instances **at** 10 in Fig. 3.2a. Moreover, even in these instances, the runtime gaps are smaller than the time needed to pre-generate the kernel matrix for the solver, so fair DCD still has a shorter total runtime.

Hyperparam.	Fair. Const.	No Fair	ness	e=	: 0	e =	0:01	e =	0:05	e =	0:1	e =	: 1
C:1:0;s:0:02	Only DP.	0.74 (	0.01	0.73	0.00	0.78	0.01	0.71	0.01	0.73	0.00	0.76	0.00
	Only TPP.	0.74 (	0.01	0.73	0.00	0.77	0.00	0.72	0.00	0.73	0.00	0.77	0.00
	DP. and TPP.	0.74 0	0.01	0.73	0.00	0.78	0.01	0.71	0.01	0.73	0.00	0.76	0.00
C:1:0;s:0:10	Only DP.	0.76 0	0.01	0.76	0.00	0.78	0.00	0.73	0.00	0.75	0.00	0.78	0.00
	Only TPP.	0.76 0	0.01	0.76	0.00	0.78	0.00	0.74	0.00	0.75	0.00	0.77	0.00
	DP. and TPP.	0.76 0	0.01	0.75	0.00	0.78	0.00	0.73	0.00	0.75	0.00	0.78	0.00
C:1:0;s:1:00	Only DP.	0.71 (	0.01	0.71	0.00	0.76	0.00	0.69	0.00	0.69	0.00	0.70	0.00
	Only TPP.	0.71 (	0.01	0.71	0.00	0.76	0.00	0.69	0.00	0.69	0.00	0.69	0.00
	DP. and TPP.	0.71 (	0.01	0.71	0.00	0.76	0.00	0.68	0.00	0.69	0.00	0.70	0.00

Table 3.1 Predictive accuracy on test data using non-linear SVMs with a RBF kernel on German credit dataset

Hyperparam.	Fair. Const.	No Fairness	e = 0	e = 0:01	e = 0:05	e = 0:1	e = 1
C:1:0;s:0:02	Only DP.	0.15 0.01	0.08 0.03	0.13 0.02	0.09 0.02	0.11 0.03	0.11 0.02
	Only TPP.	0.15 0.01	0.13 0.02	0.16 0.02	0.11 0.01	0.14 0.03	0.14 0.02
	DP. and TPP.	0.15 0.01	0.09 0.03	0.13 0.02	0.09 0.02	0.10 0.03	0.11 0.02
C:1:0;s:0:10	Only DP.	0.14 0.03	0.08 0.02	0.12 0.02	0.08 0.00	0.07 0.02	0.19 0.02
	Only TPP.	0.14 0.03	0.11 0.01	0.16 0.01	0.08 0.00	0.08 0.01	0.22 0.01
	DP. and TPP.	0.14 0.03	0.09 0.02	0.12 0.02	0.08 0.00	0.07 0.01	0.19 0.02
C:1:0;s:1:00	Only DP.	0.03 0.01	0.01 0.00	0.01 0.00	0.02 0.01	0.04 0.00	0.04 0.01
	Only TPP.	0.03 0.01	0.01 0.00	0.01 0.00	0.02 0.00	0.04 0.00	0.06 0.00
	DP. and TPP.	0.03 0.01	0.01 0.00	0.01 0.00	0.02 0.00	0.04 0.00	0.04 0.01

Table 3.2 Demographic parity between renters vs. homeowners on test data using non-linear SVMs with a RBF kernel on German credit dataset

Hyperparam.	Fair. Const.	No Fairness		e = 0		e = 0:01		e = 0:05		e = 0:1		e = 1	
C:1:0;s:0:02	Only DP.	0.19	0.07	0.31	0.06	0.07	0.03	0.15	0.02	0.09	0.01	0.16	0.02
	Only TPP.	0.19	0.07	0.40	0.03	0.03	0.01	0.18	0.01	0.10	0.02	0.18	0.02
	DP. and TPP.	0.19	0.07	0.32	0.05	0.07	0.03	0.15	0.02	0.10	0.01	0.17	0.02
C:1:0;s:0:10	Only DP.	0.16	0.07	0.33	0.04	0.07	0.03	0.06	0.01	0.18	0.05	0.20	0.01
	Only TPP.	0.16	0.07	0.39	0.02	0.04	0.02	0.06	0.00	0.12	0.04	0.20	0.00
	DP. and TPP.	0.16	0.07	0.34	0.04	0.07	0.03	0.06	0.01	0.18	0.05	0.20	0.01
C:1:0;s:1:00	Only DP.	0.07	0.02	0.05	0.00	0.13	0.00	0.09	0.01	0.09	0.00	0.03	0.01
	Only TPP.	0.07	0.02	0.05	0.00	0.13	0.00	0.07	0.00	0.09	0.00	0.01	0.00
	DP. and TPP.	0.07	0.02	0.05	0.00	0.13	0.00	0.08	0.01	0.09	0.01	0.02	0.01

Table 3.3 True positive rate parity between renters vs. homeowners on test data using non-linear SVMs with a RBF kernel on German credit dataset

Hyperparam.	Fair. Const.	No Fairnes	s  e	= 0	e =	0:01	e =	0:05	e =	0:1	e=	: 1
C:1:0;s:0:02	Only DP.	0.83 0.00	0.83	0.01	0.83	0.01	0.83	0.00	0.84	0.01	0.84	0.01
	Only TPP.	0.83 0.00	0.84	0.01	0.84	0.01	0.84	0.01	0.84	0.01	0.84	0.01
	DP. and TPP.	0.83 0.00	0.83	0.00	0.83	0.01	0.83	0.00	0.67	0.17	0.84	0.01
C:1:0;s:0:10	Only DP.	0.82 0.00	0.83	0.01	0.83	0.01	0.83	0.01	0.83	0.01	0.83	0.01
	Only TPP.	0.82 0.00	0.83	0.01	0.83	0.01	0.83	0.01	0.83	0.01	0.83	0.01
	DP. and TPP.	0.82 0.00	0.83	0.01	0.83	0.01	0.83	0.01	0.83	0.01	0.83	0.01
C:1:0;s:1:00	Only DP.	0.76 0.00	0.76	0.00	0.76	0.00	0.76	0.00	0.76	0.00	0.76	0.00
	Only TPP.	0.76 0.00	0.76	0.00	0.76	0.00	0.76	0.00	0.76	0.00	0.76	0.00
	DP. and TPP.	0.76 0.00	0.76	0.00	0.76	0.00	0.61	0.15	0.76	0.00	0.76	0.00

Table 3.4 Predictive accuracy on test data using non-linear SVMs with a RBF kernel on Adult income dataset

Hyperparam.	Fair. Const.	No Fa	airness	e =	= 0	e =	0:01	e =	0:05	e =	0:1	e=	: 1
C:1:0;s:0:02	Only DP.	0.19	0.00	0.11	0.01	0.11	0.01	0.14	0.01	0.15	0.01	0.18	0.01
	Only TPP.	0.19	0.00	0.13	0.01	0.18	0.01	0.18	0.01	0.18	0.01	0.18	0.01
	DP. and TPP.	0.19	0.00	0.12	0.01	0.11	0.01	0.14	0.01	0.12	0.03	0.18	0.01
C:1:0;s:0:10	Only DP.	0.21	0.00	0.15	0.01	0.16	0.01	0.17	0.01	0.18	0.01	0.19	0.01
	Only TPP.	0.21	0.00	0.16	0.01	0.19	0.01	0.19	0.01	0.19	0.01	0.19	0.01
	DP. and TPP.	0.21	0.00	0.15	0.01	0.16	0.01	0.17	0.01	0.18	0.01	0.19	0.01
C:1:0;s:1:00	Only DP.	0.07	0.00	0.02	0.00	0.02	0.00	0.05	0.00	0.06	0.00	0.06	0.00
	Only TPP.	0.07	0.00	0.06	0.00	0.06	0.00	0.06	0.00	0.06	0.00	0.06	0.00
	DP. and TPP.	0.07	0.00	0.02	0.00	0.02	0.00	0.04	0.01	0.06	0.00	0.06	0.00

Table 3.5 Demographic parity between female vs. male test data using non-linear SVMs with a RBF kernel on Adult income dataset

setting with true label§y<sub>i</sub>g and predicted label§y<sub>i</sub>g. They de ne the utility function as a function of  $y_i$ ;  $\hat{y}_i$ , and the speci c format is chosen to re ect whether risk averse, neutral or seeking, and how close the predicted outcoméries to i's desirable outcome. They then de ne a utilitarian sum of these individual utilities as the social welfare measure, and propose to add a constraint on this social welfare value to standard ML models as an in-processing fair ML approach. Hu and Chen (2020) study a similar utility de nition without the risk component in a classi cation setup. They evaluate the overall welfare associated with classi cation decisions through comparing a vector of welfare values, which measure the utilitarian welfare by group. Also in a classi cation setting, Corbett-Davies and Goel (2018) suppose each group has xed bene ts and costs associated with classi cation outcomes, and these values are used as parameters in the utility functions. A group's utility aggregates the bene ts and costs that individuals of the group incur from their classi cation outcomes. A more re ned view of utility is studied in Heidari et al. (2019): they partition one's actual utility into an effort-based component and an advantage component. Utilizing this partition, they group individuals by effort-based utilities and propose a fairness measure equivalent to the expected advantage utility of the worst-off group.

# 4.2 The Basic Optimization Problem

The general problem of maximizing social welfare can be stated

$$\max_{\mathbf{x}} W \mathbf{U}(\mathbf{x}) \quad \mathbf{x} \ge S_{\mathbf{x}} \tag{4.1}$$

where  $\mathbf{x} = (x_1; \dots; x_n)$  is a vector of resources distributed across stakeholders; n, and  $S_{\mathbf{x}}$  is the set of feasible values of permitted by resource limits and other constraites  $(U_1; \dots; U_n)$  is a vector of utility functions where  $U_i(\mathbf{x})$  de ness the utility experienced by stakeholders a result of the resource distribution. We can normally write  $U_i(\mathbf{x})$  as  $U_i(x_i)$ , since a stakeholder's utility typically depends only on the resources allotted to that stakeholder. FW4(W), is a social welfare function that measures the desirability of a vector of utilities. Problem (4.1) maximizes social welfare over all feasible resource allocations.

In practice, it is often convenient to model the utility functidulsusing constraints, because this results in problems better suited for optimization solvers. One therefore writes (4.1) as

$$\max_{\mathbf{x};\mathbf{u}} W^{0}(\mathbf{u}) \quad (\mathbf{x};\mathbf{u}) \ge S_{\mathbf{x}\mathbf{u}} \tag{4.2}$$

where **u** is a vector of utilities, an  $\mathfrak{S}_{xu}$  is de ned so that  $(\mathbf{x}; \mathbf{u}) \ge S_{xu}$  implies  $\mathbf{x} \ge S_x$  and  $\mathbf{u} = \mathbf{U}(\mathbf{x})$ . The function  $W^0$  is a possibly simplied version of W that yields an equivalent optimization problem due to constraints de nin  $\mathfrak{S}_{xu}$ .

To simplify exposition, we assume that the original problem constraints that **(B) De**nsist of (or can be approximated by) a system of linear inequalities and equations. Thus, for example,

when we say that (4.1) is a linear programming (LP) problem for a given, we mean that (4.1) can be formulated as an LP problet (4.2) when  $S_x$  is defined by a linear system. The linearity assumption actually allows a great deal of modeling exibility, because an be approximated by linear constraints wheneves is convex and U(x) is a concave function of. The latter occurs in the common situation where is linear or represents decreasing returns to scale.

All of the SWFs considered here can be formulated as linear, nonlinear, or mixed integer programming problems for whichdvancedsolution technology exists. An LP model optimizes a linear function over continuous variables, subject to linear inequality constraints. The problem reisnely well solved. Nonlinear programming (NLP) models timize a nonlinear function over continuous variables, subject to linear or nonlinear inequality constraints. All the NLP models considered here are relatively easy to solve. Mixein teger/linear programming (MILP) models are LP problems except that some variables must take integer values. They are combinatorial in nature, but state-of-the-art software frequently solves industrial instances with thousands of discrete variables.

If some of the original problem variables are discrete, an otherwise LP problem becomes an MILP problem, and an NLP problem becomes a mixed integer/nonlinear programming (MINLP) problem. The latter can be quite hard to solve. An MILP problem of course remains an MILP problem.

#### 4.2.1 Advantages of Optimization

The optimization of social welfare functions offers several advantages as a framework for incorporating fairness into AI.

- Social welfare functions providebaroader perspective on fairnestan can be achieved by
  focusing exclusively on bias and concepts of parity across groups. They not only have the
  exibility to represent a wide range of fairness concepts, but they encourage modelers to take
  into account the overall welfare of those affected. While AI-based decision making already
  strives to maximize predictive accuracy, a welfare perspective allows it to consider explicitly
  the more general bene ts that accurate predictions can deliver, as well as whether the bene ts
  are distributed justly.
- Social welfare functions allow one totalance equity and ef ciencin a principled way. Where
  equity is an issue, there is often a desire for ef ciency as well. A social welfare approach
  obliges one to consider how equity and utilitarian goals should be represented and balanced
  when one chooses the function to be maximized. One can of course maximize ef ciency subject
  to a constraint on some measure of inequity, but this provides no principled way of regulating
  the trade-off between the two.
- Optimization models allow one to harnessswerful optimization methods which have been developed and re ned over a period of 80 years or more. A wide variety of social welfare functions can be formulated for solution by highly advanced linear, nonlinear, and mixed integer programming solvers. We provide examples in Section 4.5.

 Optimization models offer enormous exibility toclude constraints on the problerDecisions are normally made in the context of resource constraints or other limitations on possible options. These can be represented as constraints in the optimization problem, as nearly all state-of-the-art optimization methods are designed for constrained optimization. Also, a complex social welfare function can often be simpli ed by adding constraints to the optimization problem, resulting in a problem that is easier to solve.

# 4.3 Example: Mortgage Loan Processing

We use mortgage loan processing as a running example, as it is a much-discussed application of AI-based decision making. Issues of fairness arise when an AI system is more likely to deny loans to members of certain groups, perhaps re ecting minority status or gender. A frequently used remedy is to apply statistical bias metrics to detect the problem and adjust the decision algorithms in an attempt to solve it.

Yet bias is only one element of a broader decision-making context. For one thing, there is a clear utilitarian imperative. The reason for automating mortgage decisions in the rst place is to predict more accurately who will default, because defaults are costly for the bank and devastating to home buyers. The desire for accurate prediction is, at root, a desire to maximize utility. Furthermore, bias is regarded as unfair in large part because it reduces the welfare of a segment of society that is already disadvantaged. An aversion to bias is, to a great degree, grounded in a desire for distributive justice in general. All this suggests that loan decisions should be designed to achieve what we really want: ef ciency and distributive justice, rather than focusing exclusively on predictive accuracy and group parity.

The social welfare function (4.1) should be selected to balance efficiency and equity in a suitable fashion; we consider some candidate SWFs in Section 4.5. The stake holders might include the loan applicants, the bank, the bank's stockholders, and the community at large. For simplicity, we focus on the loan applicants as stakeholders. The utility fund ticon verts a given set of loan decision  $d = (d_1; ...; d_n)$  to a vector of expected utilities  $= (u_1; ...; u_n) = U(d)$  that the stakeholders experience as a result. Since granting a loan is a yes-or-no decision, we can de ne to be a binary variable with = 1 if applicanti receives a loan. The utility measure  $U_i(d_i)$  for applicanti could depend on the applicant's nancial situation as well as the amount of the loan, as for example when the marginal value of a loan dollar is greater for an applicant who is less well-off. The SWF can re ect a preference for granting loans to disadvantaged applicants even when they have a somewhat higher probability of default, so as to ensure a more just distribution of utility. This could have the effect of avoiding bias against minority groups, but as part of a more comprehensive assessment of social welfare.

A bank can rely solely on a fully speci ed social welfare optimization model to determine the optimal loan decisions; this is the standard approach in the optimization literature. The welfare

optimization model can also be integrated with other AI methods to support loan decision-making. We specify several possible integration scenarios.

The loan decisional and the associated utilities may depend on predictions from machine learning. Utilizing historical data on loan applications, approvals and defaults, the bank can train machine learning models to predict whether an applicant is quali ed for a loan or the likelihood that an applicant will default if a loan were granted, which respectively require classi cation and regression models. In the former case, suppose applicast true labely i = 1 if he/she is qualied and  $y_i = 1$  otherwise. The applicant's predicted labely i = 1 if he/she is classified as a qualities applicant, an $\phi_i = 1$  otherwise. One way to use these classi cation outcomes is to consider granting loans only for those labeled as 'quali ed', namely the welfare optimization model requires if  $\hat{y}_i = 1$  and  $\hat{y}_i = 1$ . These constraints on also affect the set of feasible utilities thus the corresponding welfare-optimizing decisions. In the latter case, we denote applicate default probability aspi and the predicted probability as. Suppose the bank associates different utilities for when an applicant repays or defaults a loan, then the utility can be de negleas  $\hat{\mathbf{p}}_i u_i^1 + (1 \hat{\mathbf{p}}_i) u_i^0$ , whereu<sup>1</sup> is the utility that results if repays the loan and if i defaults. As shown, the ML-based predictions are used to fully specify the bank's social welfare optimization problem. When confronted with a batch of loan decisions, the bank will use machine learning models to generate the needed predictions, then maximized (U(d)) subject to feasibility constraints, such as, a constraint В on the funds available (where is the requested loan amount).

Another option is for the bank to solve the optimization problem in advance, before particular applicants are considered. It would maxim $i\underline{A} \notin U(d)$  over a set of hypothetical applicants corresponding to various nancial pro les, again using ML-based default probabilities as input. In this case, the utility  $U_k(d_k)$  that accrues to a potential applicant type would depend in part on the estimated number of applicants in the population that have the corresponding pro le. Then when someone with nancial pro le k applies for a loan, the bank would award the load it is 1 in the optimal solution of the welfare maximization problem.

We will later refer to these mentioned options as example **sost**-processing integration of social welfare optimization and AI, because the welfare-based fairness in injected after the learning phase with AI tools. An alternative integration perspective is to incorporate welfare-based fairness considerations into the learning components, and we refer to the sequecessing integration. We specify an example of in-processing integration for the case where the bank uses a classi cation model to predict whether an applicant is quali ed, namely whether a loan application should be approved. The standard training algorithm minimizes the predictive loss, or equivalently maximizes the classi cation accuracy. To incorporate welfare-based fairness, utilities need to be de ned with respect to classi cation outcomes, name  $\mathbf{y}_{,=} \mathbf{U}(\hat{\mathbf{y}})$ . We consider a linear utility formate  $(\hat{y}_i) = \mathbf{b}_i + \mathbf{g}_i \hat{y}_i$ , where  $\mathbf{b}_i$ ;  $\mathbf{g}_i$  respectively denotes starting utility and additional utility gain/loss from classi cation. As we later illustrate in our case study, we can as**b** igg based on a person's true laby eand other features including's group membership. Note that  $\mathbf{g}_i$  values should t the classi cation contexts.

For example, we expect a true positive outcome to have a higher utility than a false negative outcome, that is, when  $y_i = 1$ , the selected;  $g_i$  should satisfy  $u_i(1) > u_i(-1)$ .

# 4.4 Welfare-based Fairness: A General Framework

Drawing motivation from the running example, we now formalize a general framework for designing AI systems with welfare-based fairness guarantees.

# 4.4.1 Step 1: Specify decision problem

We begin by specifying the needed components of the decision problem. This step is critical for the success of later steps as it ensures we have a precise understanding of the problem scope and context. We highlight some key components that commonly exist in problem instances. Note that additional factors may be needed in speci c problems.

- Task: the task speci es the downstream actions and the involved resources to allocate. In our running example, the bank's task is to decide whether to grant loans to applicants.
- Stakeholders: stakeholders are individuals or groups directly or indirectly affected by the decisions, namely, they are the utility recipients in the problem.
- Goals: we characterize the desirable outcomes as goals of the decision problem. These goals serve as the guiding principles for de ning the social welfare objective).
- Constraints: these are restrictions in the problem context that limit which actions are feasible, namely, we specify constraints to de ne the dom&in A main source of restriction is the scarcity of resources, for example, the bank is subject to a budget constraint. In addition, the decision contexts may impose constraints on actions, for instance, the loan allocated to an applicant should not exceed the requested amount.

# 4.4.2 Step 2: De ne utility and social welfare functions

With a clear problem statement, we continue to de ne utility functions and social welfare functions. As we mention in the running example, it is useful to distinguish the de nitions for different components of an integrated decision-making framework. For the prediction component relying on machine learning, we de ne utility and social welfare with respect to the prediction results. For the decision component utilizing optimization, utility and social welfare should re ect the involved stakeholders' well-beings from the decision outcomes.

# 4.4.3 Step 3: Develop decision models

Depending on decision contexts, we can integrate social welfare optimization with rule-based AI systems and machine learning.

#### Full Information: Integration with Rule-based AI

When there is full information on the social welfare optimization problem, it is suf cient to solve the fully speci ed optimization model for the optimal decisions. The needed information may be available from past data, or provided by experts utilizing their domain knowledge. This full information context is suitable for integration by means of rule-based AI, which utilizes a set of rules to encode knowledge relevant to the decision and to produce pre-de ned outcomes. Rule-based systems are increasingly recognized for their capacity to support principled and transparent AI in various application domains. For instance, Brandom (2018) observes the trend in autonomous vehicle industry whereby "companies have shifted to rule-based AI, an older technique that lets engineers hard-code speci c behaviors or logic into an otherwise self-directed system." Moreover, Kim et al. (2021) demonstrate that ethical principles can be precisely represented as rules to include in an AI system. In fact, they suggest that a rule-based formulation is necessary for making ethical decisions.

We highlight two possible schemes to implement the integration of rule-based AI and optimization. The rst method is to use the optimization problem to guide the selection of rules to encode into the AI system, then rely on the rule-based system to make decisions. As we have illustrated in the mortgage example, the bank may pre-specify applicant classes and determine decision-rules for these classes using a social welfare optimization model. Such a rule-based system is straightforward to use: for a new loan applicant, the bank would rst identify which class the applicant belongs to, then approve the loan if the corresponding rule for the class says so and reject otherwise.

Alternatively, in an Al rule base, we can include rules that provide instructions for formulating the optimization problem and for choosing actions based on the optimal solution. This is consistent with the proposal from Bringsjord et al. (2006) that one could constrain Al systems with ethical principles formalized as logic statements, such as if-then statements. For example, the bank may consider rules that require applicants with certain features to receive reasonable prioritization, and these rules can be captured as constraints or incorporated into the objective function in the optimization model. Furthermore, when making the nal loan decisions, the bank may de ne rules about implementing the allocation solution obtained from the optimization problem.

#### Partial Information: Integration with Machine Learning

A second type of context arises when a limited amount of information is required to formulate the relevant optimization problem. The partial information case motivates a natural integration of optimization with machine learning. In particular, we focus on supervised learning methods that train predictive models from labelled data. Suppose a training data Betist ( $x_i$ ;  $y_i$ ) $g_{i=1}^n$  where  $x_i$  is the feature vector angli is the true label, then a supervised learning method trains a predictor function with the accuracy in the predicted labéts( $x_i$ )g as the primary goal. The ML literature has studied a large number of formats for, ranging from a simple functional form in logistic regression and support vector machine to more complex structures like decision tree and neural network. Optimization, as a technique, is broadly used to train ML models, but our emphasis is to integrate optimization as the fairness-seeking strategy.

We formalize two integration approaches that differ in the usage of social welfare optimization. First, in apost-processingiew, ML models are trained to predict the information needed to formulate the welfare optimization model, then the social welfare optimization model is solved for decisions. As shown below, the prediction step focuses solely on accuracy through the standard loss minimization, then the decision step incorporates fairness considerations via the social welfare objective. It is notable that all supervised learning methods are suitable for such post-processing integration, and the decision maker has the exibility to choose the ML methods tting for the problem context and computational requirement.

> Prediction step:  $h = \operatorname{argmin}_{h}L(h; D);$ Decision step: $dd = \operatorname{argmax}_{i}fW(U(d)) : d_{i} = d(\mathbf{x}_{i}; h(\mathbf{x}_{i}))g:$

The other approach is in-processing integration where social welfare optimization is directly embedded into a machine learning model. We can consider this approach as a type of in-processing fair ML method, and the key distinction with the majority of literature is that we encode fairness in a social welfare function. More precisely, we implement the integration by modifying the standard accuracy objective in a training model with a social welfare function. As we display below, one convenient modi cation is to use a weighted sum of the training loss and the negative social welfare. Note that the choice dW clearly affects the complexity of the learning model, hence the success of this in-processing integration is contingent on whether the resulting training model could be solved ef ciently.

 $\begin{array}{ll} \text{Prediction step: } h = \mbox{argmin}_h f \ L \ (h; D) & I \ W(\textbf{U}(h))) : u_i = U(h(\textbf{x}_i); \textbf{x}_i; y_i) g; \\ \text{Decision step} \textbf{x} d = \mbox{argma}_{\textbf{x}} f \ W(\textbf{U}(\textbf{d})) : d_i = d(\textbf{x}_i; h \ (\textbf{x}_i)) g; \\ \end{array}$ 

Remark 1. We brie y discuss the integration potentials with two other core machine learning methods, unsupervised learning and reinforcement learning (RL). Fairness has been studied in both methods, but the progress is much more limited compared to fair supervised learning. Within unsupervised learning, we focus on clustering methods. We can easily apply post-processing integration to clustering methods and utilize the trained clusters as input to specify the optimization problem. For instance, in the loan example, the bank can use clustering algorithms to decide a categorization of nancial pro les that will play a role in the optimization formulation. Recent works in fair clustering, e.g.Abraham et al. (2019); Deepak and Abraham (2020), have explored an in-processing strategy to extend K-means clustering to include fairness considerations by adding a fairness component to the usual K-means objective function. This indicates the potentials of in-processing integration, that is, we can de ne social welfare based fairness component to modify the usual clustering objective functions. In reinforcement learning, the goal is to search for a reward-maximizing policy in a dynamic environment that is typically modelled as a Markov Decision Process. De ning and achieving fairness in RL is more

challenging due to the sequential and dynamic structure. Siddique et al. (2020); Weng (2019) propose a novel framework for fair multi-objective reinforcement learning based on welfare optimization. The key component of their proposal is to replace the standard reward objective with a particular social welfare function on the reward distribution. This exactly captures the perspective of in-processing integration, hence demonstrates the potentials of social welfare optimization for seeking fairness in RL.

# 4.5 A Sampling of Social Welfare Functions

We brie y review a collection of SWFs to illustrate how they can embody various conceptions of equity. For each, we indicate the type of optimization model it yields, and whether it is appropriate for our running example of mortgage loan processing. We classify the SWFs as pure fairness metrics, functions that combine fairness and ef ciency, and statistical fairness metrics.

# 4.5.1 Pure fairness measures

Social welfare functions that measure fairness alone, without an element of ef ciency, are of two basic types: inequality metrics and fairness for the disadvantaged.

Inequality metrics abound in the economics literature. Some simple ones are represented by the following SWFs (which negate the inequality measure):

$$W(\mathbf{u}) = \begin{pmatrix} 8 \\ & (1=\bar{u})(u_{max} \ u_{min}) \\ & (1=\bar{u}) \stackrel{\circ}{a} ju_i \ \bar{u}j \\ & h^i \\ & (1=\bar{u}) \stackrel{\circ}{a} (u_i \ \bar{u})^2 \stackrel{i}{\stackrel{1}{2}} \\ & for the relative mean deviation \\ & h^i \\ & (1=\bar{u}) \stackrel{\circ}{a} (u_i \ \bar{u})^2 \stackrel{i}{\stackrel{1}{2}} \\ & for the coef cient of variation \\ & h^i \\ & (1=\bar{u}) \stackrel{\circ}{a} (u_i \ \bar{u})^2 \stackrel{\circ}{\stackrel{1}{2}} \\ & for the coef cient of variation \\ & h^i \\ & (1=\bar{u}) \stackrel{\circ}{a} (u_i \ \bar{u})^2 \stackrel{\circ}{\stackrel{1}{2}} \\ & for the coef cient of variation \\ & h^i \\ & (1=\bar{u}) \stackrel{\circ}{a} (u_i \ \bar{u})^2 \stackrel{\circ}{\stackrel{1}{2}} \\ & for the coef cient of variation \\ & for the coef cient \\ & for the coef cient \\ & for$$

There is also the well-know Gini coef cient, which is proportional to the area between the Lorenz curve and a diagonal line representing perfect equality. It corresponds to the SWF

$$W(\mathbf{u}) = 1 \quad \frac{1}{2\bar{u}n^2} \mathop{a}\limits_{i;j}^{a} ju_i \quad u_j j$$

Although these SWFs are nonlinear, all but the coef cient of variation have LP models. The coef cient of variation has a convex quadratic programming model with linear constraints, for which there are very ef cient specialized solvers.

Other fairness-based SWFs are concerned with the lot of the disadvantagedbother index measures the fraction of total utility that would have to be transferred from the richer half of the population to the poorer half to achieve perfect equality. The SWF is

$$W(\mathbf{u}) = \frac{1}{2n\bar{u}} \mathop{a}_{i}^{a} ju_{i} \bar{u}_{j}$$

The Hoover index is proportional to the relative mean deviation and can therefore be optimized using the same LP model.

The McLoone indexcompares the total utility of individuals at or below the median utility to the utility they would enjoy if all were brought up to the median utility. The index is 1 if nobody's utility is strictly below the median and approaches 0 if there is a long lower tail. The SWF is

$$W(\mathbf{u}) = \frac{1}{jI(\mathbf{u})j\tilde{u}} \mathop{a}_{i2I(\mathbf{u})}^{a} u_{i}$$

whereũ is the median of utilities inu and I (u) is the set of indices of utilities at or below the median. The McLoone index can be optimized in an MILP model.

The Hoover and McLoone indices measure only the relative welfare of disadvantaged parties, and not their absolute welfare. The aximincriterion addresses both. It is based on the Difference Principle of John Rawls, which states that inequality should exist only to the extent it is necessary to improve the lot of the worst-off (Freeman (2003); Rawls (1999); Richardson and Weithman (1999)). It can be plausibly extended to a lexicographic maximum principle. The SWF is simply

and has an LP model.

Purely fairness-oriented SWFs can be used when equity is truly the only issue of concern. In particular, they are unsuitable for the mortgage problem, where overall utility is a prime consideration.

#### 4.5.2 Combining fairness and ef ciency

Several SWFs combine equity and ef ciency, sometimes with a parameter that regulates the relative importance of each. Perhaps the best knowalpis a fairness for which the SWF is

$$W_{a}(\mathbf{u}) = \begin{array}{c} 8 \\ \gtrless \\ 1 \\ 1 \\ a \\ i \end{array} \overset{a}{}_{i} u_{i}^{1 a} \text{ for } a \quad 0; a \in 1 \\ A \\ i \\ i \\ i \\ i \end{array}$$

Larger values of imply a greater emphasis on equity, with= 0 corresponding to a pure utilitarian criterion å i ui, anda = ¥ to a pure maximin criterion. An important special case is 1, which corresponds to proportional fairness also known as the lash bargaining solution is widely used in telecommunications and other engineering applications. Both proportional fairness and alpha fairness have been given axiomatic and bargaining justi cations (Binmore et al. (1986); Harsanyi (1977); Lan et al. (2010); Nash (1950); Rubinstein (1982)). The alpha fairness SWF is irreducibly nonlinear, but because it is concave for **al**, it can be maximized with reasonable ef ciency by NLP methods.

Alpha fairness is conceptually a reasonable choice for the mortgage problem, because the bank can obtain any desired balance between utility and fairness by adjusting the it is dif cult to

justify to stakeholders any particular choice for the value of perceived bias against minorities can always be addressed by increasingOn the other hand, the presence of 0–1 variatelessoduces an MINLP model, which can be hard to solve. Thus alpha fairness may be practical only for problems with at most a few hundred applicants.

The Kalai-Smorodinsk (K–S) bargaining solution, proposed as an alternative to the Nash bargaining solution, minimizes each person's relative concession. That is, it provides everyone the largest possible utility relative to the maximum one could obtain if other players are disregarded, subject to the condition that all persons receive the same fraction fraction that all persons receive the same fraction fraction that been defended by the bargaining justi cation of Kalai and Smorodinsky (1975), this approach has been defended by Thompson (1994) and is implied by the contractarian philosophy of Gauthier (1983). The SWF can be formulated (

 $W(\mathbf{u}) = \begin{pmatrix} & \\ & a_i u_i; & \text{if } \mathbf{u} = b \mathbf{u}^{max} \text{ for someb with } 0 & b & 1 \\ & 0; & \text{otherwise} \end{pmatrix}$ 

where  $u_i^{max} = max_{(x;u)2S_{xu}} u_i$  for eachi. It can be optimized by maximizing subject to  $u = b u^{max}$  and b = 1, an easy LP problem.

The K–S criterion cannot be used for the mortgage problem, because the role of 0–1 variables in the problem almost ensure that the optimization model will be infeasible.  $\mathbf{S}_{i}^{\text{mide}} \mathbf{\bar{u}}_{i}^{1}$ , we must have  $U_{i}(x_{i}) = b u_{i}^{1}$  for all i. But  $U_{i}(x_{i}) = \bar{v}_{i} + (\bar{u}_{i} - \bar{v}_{i})x_{i}$ , which means that there must be a rational, for each i, is equal to eithe  $\bar{v}_{i} = u_{i}^{1}$  or  $\bar{u}_{i} = u_{i}^{1}$  (which correspond to setting = 0 or  $x_{i} = 1$ , respectively). It is very unlikely that the problem data will have this property.

Williams and Cookson (2000) suggest to be shold criteria for combining maximin and utilitarian objectives in a2-person context. One uses maximin until the cost of fairness becomes too great, whereupon it switches to utilitarianism, and the other does the opposite. Hooker and Williams (2012) generalize the former to persons by proposing the following SWF:

$$W_D(\boldsymbol{u}) = (n \quad 1)D + \overset{n}{\underset{i=1}{\overset{n}{a}}} \max \ u_i \quad D; u_{min}$$

where  $u_{min} = min_i f u_i g$ . The parameteD regulates the equity/ef ciency trade-off in a way that may be easier to interpret in practice than the parameter: parties whose utility is with D of the lowest utility receive special priority. Thus the disadvantaged are favoredD ated nes who is disadvantaged. As with the a parameter D = 0 corresponds to a purely utilitarian criterion a  $D e \neq 1$  to a maximin criterion. Hooker and Williams provide an MILP model of the SWF and show that it is sharp (i.e., its continuous relaxation describes the convex hull of its feasible set). Partly for this reason, they found that the model solves rapidly in computational tests.

This threshold approach is a reasonable choice for the mortgage problem. Since the problem has discrete variables regardless of the SWF used, the MILP-based threshold formulation adds relatively little complexity to the problem. In addition, loan of cers can specify in a meaningful way when an applicant is to be considered disadvantaged, by selecting an appropriate value of

One weakness of the model is that the actual utility levels of disadvantaged parties other than the very worst-off have no effect on the measurement of social welfare, as long as those utilities are within D of the lowest. As a result, the socially optimal solution may not be as sensitive to equity as one might desire. Chen and Hooker (2020a,b) address this is coertby ining utilitarianism with a leximaxrather than a maximin criterion. A leximax (lexicographic maximum) solution is found by rst maximizing the lowest utility, then while holding it xed, maximizing the second lowest utility, and so forth. Chen and Hooker combine leximax and utilitarian criteria by maximizing a sequence of threshold SWFs that have tractable MILP models. Their approach may yield more satisfactory solutions of the mortgage problem.

#### 4.5.3 Statistical bias metrics

While we argue that bias metrics afford an overly narrow perspective on fairness, they nonetheless can be expressed as SWFs if desired. The utility vactor comes simply a binary vector in which  $u_i = 1$  if individual i is selected for some bene t, and = 0 otherwise. In the mortgage example, the bene t is a mortgage loan. We set constant 1 when person actually qualities for selection (as for example when person the mortgage training set repaid the loan), and 0 otherwise. Two groups are compared, respectively indexedNbandN<sup>0</sup>. One is a protected group, such as a minority subpopulation, and the other consists of the rest of the population.

For exampledemographic parityhas the SWF

$$W(\mathbf{u}) = 1 \quad \frac{1}{jNj} \mathop{a}\limits^{a}_{i2N} u_i \quad \frac{1}{jN^{q}} \mathop{a}\limits^{a}_{i2N^{0}} u_i$$

Equalized odds an be measured in two ways, one of whick is ality of opportunity

$$W(\mathbf{u}) = 1 \quad \frac{\dot{a}_{i2N} a_{i} u_{i}}{\dot{a}_{i2N} a_{i}} \quad \frac{\dot{a}_{i2N^{0}} a_{i} u_{i}}{\dot{a}_{i2N^{0}} a_{i}}$$

Another SWF representaccuracy parity

$$W(\mathbf{u}) = 1 \quad \frac{1}{jNj} \mathop{a}\limits_{i \ge N}^{a} a_{i}u_{i} + (1 \quad a_{i})(1 \quad u_{i}) \quad \frac{1}{jNq} \mathop{a}\limits_{i \ge N^{0}}^{a} a_{i}u_{i} + (1 \quad a_{i})(1 \quad u_{i})$$

and still anothepredictive rate parity

$$W(\mathbf{u}) = 1 \quad \frac{\dot{a}_{i2N} a_i u_i}{\dot{a}_{i2N} u_i} \quad \frac{\dot{a}_{i2N^0} a_i u_i}{\dot{a}_{i2N^0} u_i}$$

The computational challenge varies widely across the various bias-oriented SWFs. The rst three SWFs above give rise to linear models (which become MILP models due to the 0–1 restriction on while the last produces an extremely dif cult nonconvex MINLP model.

Bias measures are inappropriate as social welfare objectives for the mortgage problem, because they take no account of efficiency. One can, of course, maximize predictive accuracy subject to constraints on the amount of bias, but this has a number of drawbacks:

- As previously argued, it provides a very limited perspective on the utility actually created by decisions. Indeed, the utility vector consists only of 0–1 choices.
- There is no consensus on which bias measure is suitable in a given context, if any. Bias measures were developed by statisticians to measure predictive accuracy, not to assess fairness.
- There is no principle for balancing equity and ef ciency. If equity is one of the bjectives it should be part of the bjective function. The choice of that function obliges one to justify the equity/ef ciency trade-off mechanism in a transparent manner.
- Bias measurement forces one to identifyriori which individuals in a training set should be selected for bene ts (as indicated b). In a social welfare approach, no prior decisions of this kind are necessary.
- Bias measurement forces one to designate "protected groups" (as indicated by the inhore set There is no clear principle for selecting which groups should be protected, unless one is content simply to recognize those mandated by law.

# 4.6 Case Study: Mortgage Loan Approval

# 4.7 Conclusion

We formalize a general framework for using optimization to incorporate welfare-based fairness into Al applications. The framework provides a guideline for formulating a decision task into a social welfare optimization problem. In particular, we illustrate how optimization can be integrated with rule-based Al systems and machine learning models. By expanding the fairness problem to the optimization of social welfare functions, one can achieve a broader perspective on fairness that are driven by the well-beings of stakeholders and characterize the broader fairness concepts in a principled way. Optimization models also provide the exibility of adding constraints on resources and other problem elements, while harnessing the power of highly advanced optimization solvers.

We conclude the paper by outlining a research program to explore some key questions related to the framework.

• There is a wide gap between the presented general formalization of integration strategies and practical implementations of integrated methods. For integration with rule-based AI, one important direction is to investigate how to build ethics-sensitive rule bases to t into different social welfare optimization scenarios. Previous works on formulating ethics principles into rules,

e.g. Bringsjord et al. (2006); Kim et al. (2021), may provide guidance for this direction. For integration with machine learning, future research could explore the in-processing perspective and study how to de ne social welfare functions to use as the objective in machine learning models. The modi ed objective functions need to have a format that can be efficiently trained, and the trained models need to provide the desirable fairness and welfare guarantees.

- Although optimization solvers have been developed over decades, not all classes of optimization models are readily solvable by state-of-the-art software. Among all classes, linear programming and convex programming problems can be considered tractable up to reasonably large sizes, but non-convex formulations including some mixed integer programming problems are more restricted. For practical use of social welfare optimization models, one may need to apply available computational strategies or design problem-speci c heuristics to speed up solving the optimization problems.
- The social welfare functions we consider are of a static nature, that is, a SWF does not attempt to capture potential dynamics in the utilities. A SWF takes utility values as the input, and the function values characterize the associated static utility distributions. While such a static view is often suf cient and reasonable for a one-shot decision problem, a dynamic perspective may be required in sequential decision problems where decisions need to be made repeatedly and the selected actions have incremental impacts on the long term social welfare. Future research could explore how to extend the presented optimization based framework to t a dynamic view of welfare and fairness. Although this is not a trivial task, there are many well-developed techniques to utilize, such as, stochastic optimization, Markov decision process, etc.

# Chapter 5

# Online Convex Optimization Perspective for Learning from Dynamically Revealed Preferences

# 5.1 Introduction

Preferences of an agent implicitly dictates his/her actions, and in uence for example what a company should offer as its products or how a company should personalize recommendations to an individual customer (agent). This creates incentives for the company/central decision maker to learn the preferences of their agents. Nevertheless, in reality, the true preferences of the agents are often private to the individual agents and are only implicitly revealed in the form of their behaviors/actions to the central decision maker. Such typical interactions for example include a streaming platform suggesting a number of videos to a user and tracking whether the user watches or likes them. As evident from such scenarios, inferring the agents' preference information through agent interactions and observations of their behaviors is a critical task for the decision makers in such settings.

A common assumption adopted to formalize the problem of learning from revealed preferences is that rational agents attility maximizers, that is, they choose actions to maximize their utility functions subject to a set of restrictions. The central decision maker interacting with the agents is the learner. An important learner-centric goal is to design schemes for the learner to extract useful information on the agents' utility functions. This learning point of view of revealed preferences has been explored in a broad range of literature from economics (e.g., Beigman and Vohra (2006); Varian (2006)), machine learning (e.g., Balcan et al. (2014); Dong and Zeng (2020); Dong et al. (2018b)) and operations research (e.g., Ahmadi et al. (2020); Bärmann et al. (2017); Mohajerin Esfahani et al. (2018)). Based on the type of learner-agent interactions and information objective and learning complexity.

## 96 Online Convex Optimization Perspective for Learning from Dynamically Revealed Preferences

In this chapter, we focus on a speci c setup where the learner seeks to learn the utility function of a non-strategic agent while receiving information about the agent's actions in an online fashion. This setup ts naturally in applications where the agents bene t from effective learning of their true preferences. For example this is the case when a streaming platform interacts with its users to learn their preferences. In this example, in a typical interaction, the platform recommends videos to a user and the user takes actions based on the recommendations. User actions, e.g., clicks, movie streaming, etc., are fully observable to the platform and they re ect the user's true preferences.

# 5.1.1 Our Approach and Results

We propose a novel modeling framework for learning from dynamically revealed preferences via online data-driven inverse optimization. In our setup, the learner monitors a sequence of data signals and observes the respective rational decisions of a non-strategic agent without noise over a nite time horizon of T time steps. The learner operates and receives information in an online fashion, and updates an estimate of  $q_{true}$  using newly available information at each time step.

Our framework is enabled by the identi cation  $\partial f^{im}$ , a loss function that is both simple in structure and has practical connections with conventional loss functions from the inverse optimization literature. Online inverse optimization based  $\partial f^{in}$  is a online convex optimization (OCO) problem, which enjoys the exibility to be handled by any deterministic OCO algorithm.

- In Section 5.2, we present a formal description of our problem setting. Section 5.2.1 introduces the agent's problem and discusses a rather broad decomposable structure assumption on the agent's utility functions that is capable of representing all of the utility functions studied in the data-driven inverse optimization literature as well as other key utility functions. Section 5.2.2 describes the learner's inverse optimization problem that minimizes a lossefunction () to obtain an accurate estimate the hidden parameter. We then state the online inverse optimization framework in Section 5.2.3: we describe the sequence of events and de ne regret as the performance measure.
- In Section 5.3, we utilize our structural assumption on the agent's utility function to design a new convex loss function, namesymple loss<sup>sim</sup>.
- We establish in Section 5.4 that in the noiseless setting, a bounded regret with respect to also guarantees a bounded regret with respect to all the other loss functions; see Proposition 5 and Corollary 1. We also brie y discuss the noisy setting.
- Convexity and simplicity of <sup>sim</sup> enables us to use an online convex optimization (OCO) framework (see Section 5.5) that offers the exibility to use different OL algorithms, such as, online Mirror Descent (MD) utilizing a rst-order oracle (Section 5.5.1) and implicit OL based on a solution oracle (Section 5.5.2). In the noiseless setup, our framework equipped with online MD coversall of the problem classes studied in the online data-driven inverse optimization

literature, and matches the corresponding state-of-the-art regret bounds with respect to the loss functions in **a**ni ed manner. In particular, our results immediately generalize the customized algorithms from Bärmann et al. (2017) and completely bypass the requirement to verify the rather technical assumptions of Dong et al. (2018a) and the need to use their expensive MISOCP-based solution oracle; see Section 5.5.3 for a detailed comparison discussion.

Our numerical study in Section 5.6 highlights that when compared to the based implicit OL with an MISOCP solution oracle approach of Dong et al. (2018), based OL algorithms equipped with a rst-order oracle or a solution oracle, particularly online MD, demonstrate signi cant advantages in terms of both the learning performance (i.e., regret bounds) and the computation time. This is directly in line with our theoretical results. Moreover, these results seem to be fairly robust with respect to changes in the structure of the agent's domain as well as the noise in observations.

## 5.1.2 Related Literature

Varian (2006) is one of the earliest and most celebrated work for learning from revealed preferences in the economics literature. They study constructing utility functions of the agent to explain a sequence of her/his observed actions. Nevertheless, this approach has a main shortcoming-a utility function capable of explaining past actions not necessarily also guarantees accurate predictions of the future actions. Consequently, Beigman and Vohra (2006) have initiated a new line of research to learn utility functions capable of predicting future actions with statistical performance guarantees. Beigman and Vohra (2006) examine a statistical setup where the learning algorithm takes as input a batch of observations and is evaluated by its sample complexity guarantees. Zadimoghaddam and Roth (2012) focus on the setting where the agent has a linear or linearly separable concave utility function, and propose learning algorithms with polynomially bounded sample complexity. Balcan et al. (2014) identify a connection between the problem of learning a utility function and the structured prediction problem of D-dimensional linear classes. Through this connection, Balcan et al. (2014) suggest an algorithm for learning utility functions that is superior (in terms of sample complexity) than the method from Zadimoghaddam and Roth (2012) in the case of linear utility functions and is also applicable for learning separable piecewise-linear concave functions and CES functions with explicit sample complexity bounds.

As an alternative to this statistical view, Balcan et al. (2014) study a query-based learning model, where the learner aims to recover the exact utility function by querying an oracle for the agent's optimal actions. The query-based models consider an online feedback mechanism where the learner receives one observation of the agent's action at a time. When the learner has the power to choose which observation to receive from the query oracle, Balcan et al. (2014) give exact learning algorithms for several classes of utility functions. There is a recent research streteroring to optimize he learner's objective function based on information from revealed preferences of the agents. In this stream it is often assumed that the learner has similar power on the selection of the observations. For

example, Amin et al. (2015) and Ji et al. (2018) propose algorithms for nding the pro t-maximizing prices for a seller, who has price controlling power and learns buyer preferences by observing the buying behavior at different price levels. Roth et al. (2016) and Dong et al. (2018b) consider the learning task as Stackelberg games, where the leader player is the learner and a follower player is a strategicagent with incentive to manipulate actions and hide information.

We note two restrictions with the problem setup in these fore-mentioned papers. First, the assumption that the learner can choose observations is not always achievable in practice. A more realistic setup is accommodated by the data-driven inverse optimization view where the learner does not control the sequence of observations. Second, when the learner is optimizing an objective function that does not explicitly measure how well s/he is learning about the agent, the approaches that are effective for choosing the learner's objective-optimizing action provide no guarantees on the quality of the learned agent information.

Inverse optimization generalizes the query-based view and offers a natural abstraction of learning from revealed preferences. This approach is typically used in settings with non-strategic agents, in which the agents have no incentive to hide information from the learner, and thus an agent's decisions reveal her/his true preferences. In this setting, the learner's goal is to recover unknown parameters of an agent's utility function from the observations of her/his true optimal solution. Chan et al. (2021) provides a comprehensive review of inverse optimization. We next summarize key developments in the literature, with an emphasis on two topics that are more relevant to this paper: data-driven inverse optimization and online inverse learning.

Early studies on inverse optimization examine the setting where the agent's optimization problem is xed, see e.g., Ahuja and Orlin (2001); Heuberger (2004); Iyengar and Kang (2005); Schaefer (2009). Unfortunately, this classical setup is limited in its practical applicability as it ignores uncertainty in the environment. A new thread of research on data-driven inverse optimization studies a exible setup, where the learner observes the agent's optimal or sub-optimal decisions corresponding to varying external data signals. In the noiseless case, that is, when observations of optimal solutions/agent actions are available, Keshavarz et al. (2011) show that data-driven inverse optimization of convex programs is polynomial time solvable. In the case of noisy observations, Aswani et al. (2018) proves that such problems are NP-hard in general.

Data-driven inverse optimization is further categorized based on whether observations are given as a batch upfront or in an online manner. In the batch setup, Keshavarz et al. (2011) study the inverse optimization of identifying the unknown af ne weights in a convex objective function that is an af ne combination of pre-selected basis convex functions. Recent work of Aswani et al. (2018) and Mohajerin Esfahani et al. (2018) in the batch setup investigates the inverse optimization of general convex programs without the basis function structure. Aswani et al. (2018) adq**prettie**tion loss `<sup>pre</sup>, which measures the difference between the observed agent action and the predicted agent action through squared norm distance, as the inverse optimization objective. They formulate the inverse problem into a bilevel program using Lagrangian duality, and present two heuristic algorithms with approximation guarantees for solving the bilevel formulation. Mohajerin Esfahani et al. (2018) use
suboptimality loss<sup>sub</sup>, which is de ned as the difference between objective values at the observed agent action and the predicted action, as their loss function and provide a distributionally robust formulation of the inverse problem. Batch setup requires that the learner receives observations of the agent's actions all at once. However, obtaining a large batch of observations all at once as well as learning from such a batch often presents operational and computational challenges. In practice, such strong batch feedback is rare as the learner often interacts with the agent repetitively in a dynamic environment.

A recent stream of research Bärmann et al. (2017); Dong et al. (2018a) adopts a dynamic information acquisition setup and studies the online data-driven inverse optimization where the learner observes a stream of the agent's actions one by one in an online fashion. Both Bärmann et al. (2017) and Dong et al. (2018a) suggest OL algorithms and measure their performance regardly e.e., the difference between the losses incurred from online estimates of the unknown parameters in the agent's utility function and the of ine optimal estimate. Bärmann et al. (2017) consider the problem of learning the linear utility function of an agent given the noiseless online observations of the agent's actions in a dynamic environment. They propose two specialized OL algorithms with rst-order oracles that both achieve a bound  $\mathfrak{A}(\overline{T})$  on the sum of the suboptimality loss<sup>ub</sup> and the estimate loss`est after T periods but lacks regret guarantees with respect to the prediction ploss Dong et al. (2018a) consider the setup, where the learner wishes to learn an unknown linear component of an agent's quadratic objective function from noisy observations. By utilizing the implicit OL framework of Kulis and Bartlett (2010) equipped with a Mixed Integer Second Order Cone Program (MISOCP)-based solution oracle, they provide a regret bound  $\delta (\overline{T})$  with respect to the prediction loss` pre after T periods whenever is convex. Dong et al. (2018a) present a number of rather technical assumptions that guarantee the convexit of however these assumptions are not only dif cult to verify but also quite restrictive. In fact, in Dong et al. (2018a), these were shown to hold only for a speci c class of convex quadratic problem.

Notation. We let  $\mathbb{R}^n_+$  be the set of nonnegativedimensional vectors. For a given vectorwe usev<sub>i</sub> to denote itsi-th element. We let  $[n] := f 1; \ldots; ng$ , and we us  $e_{a_i}g_{i2[n]}$  to represent a collection of entries, such as vectors, functions, etc., indexed in 2t [n]. For a differentiable function  $\tilde{N} f(x)$  to denote the gradient of at x. For a nondifferentiable function, we use  $\tilde{N} f(x)$  to denote the gradient of at x.

# 5.2 Problem Setting

In our setting, the learner monitors a sequence of external sign  $a_{[\Sigma]}$  R<sup>k</sup> and observations f ytg<sub>t2[T]</sub> R<sup>n</sup> of the agent's respective optimal decisions over a nite time horizon to the steps.

the agent's problem is given by  $\inf \min_x f x \colon x \ge [1;1]; x \in 0g; 0+q$  (1)g = minf 1; qg. Then, we deduce that the agent's optimal solution will satisfy the following:

- Whenq < 0, agent's optimal solution iss(q) = 1;
- Whenq > 0, x(q) = 1 if q < 1, andx(q) = 0 if q 1 (note that in the case of alternative optima, we assume that the solver will break ties by selecting the solution with smaller norm);
- When q = 0, x(q) = 1.

To summarize, in the given doma $\mathfrak{Q}$ , when  $2 \begin{bmatrix} 3; 1 \end{bmatrix}$ , x(q) = 1 and c(x(q)) = 0; when  $2 \begin{bmatrix} 1; 3 \end{bmatrix}$ , x(q) = 0 and c(x(q)) = 1.

We next show that implicit OL based  $\partial h^{re}$  with a solution oracle may lead to an unbounded regret. Suppose we chook  $e = \frac{1}{t}$  for all t. Then, at time step, based on the implicit OL based on `pre, we updated  $e_{t+1}$  by solving the following optimization problem: in this example,  $(q) = kx(q_{true}) = x(q)k^2 = (1 + x(q))^2$ , hence

$$q_{t+1} = \underset{q2[3;3]}{\operatorname{argmin}} \frac{1}{2} (q - q_t)^2 + \frac{1}{t} (-1 - x(q))^2$$
:

If we initialize  $q_1 = 3$ , then the above update will generate =  $\arg\min_{q \ge [-3;3]} \frac{1}{2}(q - 3)^2 + \frac{1}{t}(-1 - x(q))^2$ . To decide the optimal solution, we need to compare three scenarios:qwthen, the objective value  $i\mathfrak{D} + \frac{1}{t}(-1 - x(q_1))^2 = \frac{1}{t}(-1 - 0)^2 = \frac{1}{t}$ ; when q < 1, the objective value is  $\frac{1}{2}(q - q_1)^2 + \frac{1}{t}(-1 - x(q))^2 = \frac{1}{2}(q - 3)^2 + 0 - \frac{1}{2}(1 - 3)^2 = 2 > \frac{1}{t}$  (where we used (q) = -1 for q = 1); when  $1 - q < q_1$ , the objective value  $i\frac{1}{2}(q - q_1)^2 + \frac{1}{t}(-1 - x(q))^2 = \frac{1}{2}(q - 3)^2 + \frac{1}{t}(-1 - 0)^2 > \frac{1}{t}$ . Therefore, we have  $q_2 = q_1$ , and by the same derivation, later iterations will always stay at the same estimated  $\mathbf{t} = q_1$ . This means the implicit OL algorithm will generate  $\mathbf{s}$  for all t, and each iteration the learner incurs prediction loss  $a^{\text{pre}}(q_1) = kx(q_{\text{true}}) - x(q_1)k^2 = kx(0) - x(3)k^2 = 1$ . Therefore, the associated regret with respect to is unbounded  $a\mathbf{s}$ !

$$\mathsf{R}_{\mathsf{T}}(\mathsf{f}\,\check{}_{t}^{\mathsf{pre}}\mathsf{g}_{t2[\mathsf{T}]};\mathsf{f}\,\mathsf{q}_{t}\mathsf{g}_{t2[\mathsf{T}]}) = \overset{\circ}{\mathsf{a}}\,\check{}_{t}^{\mathsf{pre}}(\mathsf{q}_{t}) \quad \overset{\circ}{\mathsf{a}}\,\check{}_{t}^{\mathsf{pre}}(\mathsf{q}_{t\mathsf{rue}}) = \mathsf{T}:$$

We note that this example does not invalidate the regret convergence in Theorem 20. With the contrived de nition ofc(x), the loss function<sup>pre</sup> does not satisfy the Lipschitz continuity assumption needed for regret convergence guarantees given in Theorem 20. To be more specified,= $(x(q_{true}) x(q))^2$ : considere > 0,  $q_1 = 1$ ;  $q_2 = 1 + e > 1$ , we conclude<sup>pre</sup><sub>t</sub> $(q_1) = (1 (1))^2 = 0$  and  $\sum_{t=1}^{t} (q_2) = (1 0)^2 = 1$ . As e! 0, there is no niteG as a valid Lipschitz constant fo<sup>pre</sup><sub>t</sub>.

We also examine the use of online Mirror Descent (MD) based<sup>50</sup> in the same setup. Let Euclidean distance be the distance generating function in Bregman distance, then online MD simplies to projected gradient descent. We again chdqse  $\frac{1}{t}$  for all t, then at time step we updateq<sub>t+1</sub> via

$$q_{t+1} = proj_{[3;3]} q_t \frac{1}{t} c(x(q_{true})) c(x(q_t))$$
 :

frameworks for solving robust convex optimization in Ben-Tal et al. (2015); Ho-Nguyen and K I nç-Karzan (2018); Ho-Nguyen and K I nç-Karzan (2019). For our particular application, in Section 5.3, we design linear functions  $t_{t}^{sim}(q)g_{t2[T]}$  to which we apply OCO algorithms. While using non-arbitrary classes of loss functions does not invalidate any guarantees from deterministic OCO methods such as online MD, more caution is needed for OCO algorithms which involve randomness and provide guarantees **ex**pected regrestuch as stochastic gradient descent. This is due to the fact that the design of the loss functions in speci c applications may create undesirable dependence among random variables and invalidate certain steps used in the analysis of stochastic OCO algorithms. Hence, here we will focus on deterministic OCO algorithms.

# 5.3 Loss Functions

Loss function (q) plays a key role in the formulation of the inverse problem 8). Since the learner's goal is to mimic the agent's true action with the estimated preferent the appropriate loss functions should re ect how close the prediction (q; u) is to  $x(q_{true}; u)$  at a given signal. The following are common loss functions used in inverse optimization context (recall that agent's forward objective in (5.1) and  $(q; u_t)$  is the optimal solution to (5.2) for given,  $u_t$ ):

- Prediction loss:  $pre(q; x(q; u_t); y_t; u_t) \coloneqq ky_t \quad x(q; u_t)k^2$ ,
- Suboptimality loss:  $sub(q; x(q; u_t); y_t; u_t) := f(y_t; q; u_t) f(x(q; u_t); q; u_t)$ , and
- Estimate loss:  $est(q; x(q; u_t); y_t; u_t) \coloneqq f(x(q; u_t); q_{true}; u_t) \quad f(y_t; q_{true}; u_t).$

These functions use the observation a proxy of the true  $actiox(q_{true}; u_t)$ . Under perfect information with  $y_t = x(q_{true}; u_t)$ , `pre directly compares the distance between the true action and the predicted action.<sup>sub</sup> and `est utilize the agent's objective function.<sup>sub</sup> measures how much would affect the agent's optimal objective value at estimate<sup>est</sup> measures how much(q; u\_t) would change the observed agent's objective value. Under imperfect information, due to the noises present in  $y_t$ , the loss values can be split into two components, one re ects the difference bet(vaper) andx(q<sub>true</sub>; u<sub>t</sub>), and the other one is dependent on the noise shift(qQue; u<sub>t</sub>) to  $y_t$ .

Within data-driven inverse optimization in a batch seture, is used in Aswani et al. (2018) and <sup>sub</sup> is used in Mohajerin Esfahani et al. (2018) (see Section 5.1.2). In the online inverse optimization setup, <sup>sub</sup> and <sup>est</sup> are studied by Bärmann et al. (2017) under the assumption **f tisatinear** inx, and <sup>pre</sup> by Dong et al. (2018a) when the forward problem is a quadratic program with a special structure. The OL algorithms from these latter two papers are customized for the chosen loss functions and forward problem structure, indicating the lack of a uni ed general framework.

We introduce the following shorthand notation.

 $\hat{x}_{t}^{\text{pre}}(q) \coloneqq \hat{y}_{t}^{\text{pre}}(q; x(q; u_{t}); y_{t}; u_{t}); \quad \hat{y}_{t}^{\text{sub}}(q) \coloneqq \hat{y}_{t}^{\text{sub}}(q; x(q; u_{t}); y_{t}; u_{t}); \quad \hat{y}_{t}^{\text{est}}(q) \coloneqq \hat{y}_{t}^{\text{sub}}(q; u_{t}); y_{t}^{\text{sub}}(q; u_{t})$ 

Assumption 2 ensures that<sup>im</sup> is a convex function of, and thus any deterministic OCO algorithm will be applicable for regret minimization with respect  $t_{p}^{pim}g_{t2[T]}$ . Then, as a consequence of Proposition 5 (and Corollary 1), such algorithms will also be minimizing regret with respect to the loss functions  $t_{r}^{sub}$ ,  $t_{r}^{est}$  (and  $t_{r}^{pre}$ ), as well.

Remark 4. The regret bounds with respect to these loss functions have the following implications under perfect information. A sublinear regret bound with respect<sup>in</sup>timplies that the average loss incurred by the estimateq<sub>t</sub>g approaches the of ine optimal loss over time. Sublinear regret bounds with respect to<sup>sub</sup> and<sup>• est</sup> indicate that the learner is able to generate estimates that lead to vanishing errors in the predicted agent's objective function values. A sublinear regret bound with respect to<sup>pre</sup> additionally indicates that the average<sub>2</sub>-norm distance between the predicted agent's action and her/his true action decreases to zero over time. Note that none of these regret guarantees in particular ensures that the generated from the online learning process are good approximations of<sub>true</sub>. In general, this is an overly ambitious task as (Bärmann et al., 2017, Example 3.2) has shown a simple case where the exact recovery action be guaranteed. We note that stronger performance guarantees, such  $a_{2[T]}$  kq<sub>t</sub> q<sub>true</sub>k ! 0 may be possible for special cases, for example, wherk(q; u<sub>t</sub>) has a closed form the forward problem may be non-unique, our framework is not aiming to predict the chosen action (q<sub>true</sub>; u<sub>t</sub>), instead, we measure the learning performance with regret values based on the objective value of the agent.

#### 5.4.2 Imperfect Information

The case when the learner has access to only imperfect information about the agent's actions is of natural interest as well. Mohajerin Esfahani et al. (2018) identify two types of noisy information as of interest: (i)measurement noist that is, for allt 2 [T], the learner observes  $= x(q_{true}; u_t) + e_t$  with  $e_t$  denoting a random noise, and (tip)unded rationality which means for all 2 [T], the agent may choose a sub-optimal action instead ( $\alpha_{true}; u_t$ ). Such imperfect information does not affect the convexity property of loss functions, and so botth and  $\sin^{10}$  and  $\sin^{10}$  remain convex (see Lemma 3 and 4) still enabling the use of OCO algorithms for regret minimization with respect to these loss functions. For instance, we can still apply  $a^{im}$ -based OCO algorithm to minimize regret with respect for the section.

However, since  $y_t$  is no longer guaranteed to be a minimize ( $\mathfrak{S}(f)$ ) with  $u = u_t$ , Proposition 5 does not hold in general, and conseque  $\mathbb{R}(y) = u_t$ ,  $\mathbb{R}(g_{t_2[T]})$  is not guaranteed to bound the regrets with respect to the other loss functions. In addition, due to the noise  $\mathfrak{s}(\mathfrak{s}(n))$  and  $\mathfrak{s}(n)$  and

- w-center  $q_w := \underset{q^{2Q}}{\operatorname{argminw}}(q)$ .
- Set width W:=  $\max_{q \ge Q} V_{q_w}(q) = \max_{q \ge Q} w(q) = \min_{q \ge Q} w(q)$ .

When functions  $t_{t}^{sim}(q)g_{t2[T]}$  are convex in q, online MD as stated in (Ho-Nguyen and K I nç-Karzan, 2019, Algorithm 1) is applicable to guarantee a sublinear regret bound  $dh_{t}^{sim}g_{t2[T]}$ ; f  $q_{t}g_{t2[T]}$ , which further bounds regrets with respect to the other loss functions as discussed in Section 5.4.

Theorem 18. (Ho-Nguyen and K I nç-Karzan, 2019, Theorem 1) Supt is a convex and t : Q 7! R is a convex function for 2 [T]. Suppose there exists 2 (0;¥) such that all the subgradients of t are bounded, i.e.max<sub>s2¶`t</sub>(q) ksk G for all q 2 Q and t 2 [T]. Let the step size the chosen as  $h_t = \frac{2W}{G^2T}$ . At time step t, using the online Mirror Descent algorithm, we generate the step size to be chosen as  $h_t = \frac{2W}{G^2T}$ .

$$q_{t+1} \coloneqq \operatorname{Prox}_{q_t}(h_t s_t) = \underset{q \ge Q}{\operatorname{argminfhh}_t s_t}; q_i + V_{q_t}(q)g; \qquad (5.5)$$

where \$2  $\P^t(q_t)$ . Then the sequend  $e_t g_{t2[T]}$  satis es  $R_t(f_t g_{t2[T]}; fq_t g_{t2[T]}) = \frac{p}{2WG^2T}$ .

In applying Theorem 18 to the loss function  $\hat{s}^{im}g_{t^2[T]}$ , we have the subgradie  $a_t = c(y_t) c(x(q_t; u_t))$ . So, it suffices to seG 2 max  $kc(x)k : x \ge X$  g. The set width W depends on Q only and can be computed for a given and Bregman distance explicitly.

# 5.5.2 Implicit Online Learning with a Solution Oracle

We next review the implicit OL with a solution oracle from Dong et al. (2018a). This algorithm was rst introduced in its general form in Kulis and Bartlett (2010).

The implicit online learning algorithm computes

$$q_{t+1} \coloneqq \operatorname{argmin}_{d2Q} L_t(q); \tag{5.6}$$

where  $L_t(q) = V_{q_t}(q) + h_t t(q)$  and  $V_q(q^0)$  is the Bregman distance is a step size. This approach does not rely on the rst-order oracle  $\delta q$  but rather assumes the existence **soba**tion oracleto solve (5.6). Kulis and Bartlett (2010) establish the following regret bound on the OL using implicit update (5.6).

Theorem 19. (Kulis and Bartlett, 2010, Theorem 3.2.) Supp**Qsis** convex, and<sub>t</sub> : Q7! R is a convex and differentiable function fdr2 [T]. Let q be the of ine optimal solution tonin<sub>q2Q</sub>  $a_{t2[T]}$  (q). For any  $0 < a_t = \frac{L_t(q_{t+1})}{L_t(q_t)}$  for t 2 [T], for any step size t > 0, an implicit OL algorithm with the update rule (5.6) attains

$$R_{T}(f_{t}^{*}g_{t2[T]}; fq_{t}g_{t2[T]}) = \overset{a}{\underset{t2[T]}{a}} \frac{1}{h_{t}} (1 a_{t})h_{t}^{*}(q_{t}) + V_{q_{t}}(q_{t}) V_{q_{t+1}}(q_{t}) :$$
(5.7)

When `t is a convex and Lipschitz continuous function paind the domain a has a nite width with respect to the selected Bregman divergence, the regret  $b(\overline{b} u \overline{r}) durther results$  in  $a O(^{p} \overline{T})$  bound on  $R_{T}(f \ t_{g_{12}[T]}; f q_{t} g_{t_{2}[T]})$ .

Theorem 20. Suppose is convex, and for each 2 [T], `t : Q 7! R is a convex function of that is uniformly Lipschitz continuous with parameter and suppose  $\max_{q_1;q_2 2 Q} V_{q_1}(q_2)$  W. Then, by choosing  $h_t = \frac{W}{G} p_{t}^1$  for t 2 [T], an implicit OL algorithm with the update rule 6.6) attains

$$\mathsf{R}_{\mathsf{T}}(\mathsf{f}_{\mathsf{t}}^{\mathsf{g}}\mathsf{g}_{\mathsf{t}^{2}[\mathsf{T}]};\mathsf{f}\,\mathsf{q}_{\mathsf{t}}^{\mathsf{g}}\mathsf{g}_{\mathsf{t}^{2}[\mathsf{T}]}) \quad 2^{\mathsf{p}} \, \overline{\not{\mathsf{W}}\mathsf{G}^{2}\mathsf{T}}: \tag{5.8}$$

To apply Theorem 20, we can choose the Lipschitz paran**Gebey** de nition. For instance, with the loss functions  $t_t^{sim}g_{t_2[T]}$ , we have  $t_t^{sim}(q) = t_t^{sim}(q^0) = t_t^{o}(y_t) = t_t^{o}(y_t) = t_t^{o}(x_t(q_t; u_t))$  is a  $q^0 kkc(y_t) = c(x(q_t; u_t))k$ , hence g = 2max kc(x)k = x 2 X g suffices. Alternatively, with  $t_t^{sim}g_{t_2[T]}$ ,  $t_t^{sim}(q) = t_t^{sim}(q^0; u_t) = t_t^{sim}(q^0; u_t) = t_t^{o}(q^0; u_t) = t_t^{sim}(q^0; u_t) = t_t^$ 

#### 5.5.3 Comparison with the Existing Approaches

Bärmann et al. (2017) study online inverse optimization under perfect information where the agent's objective f is a bilinear function of and x, i.e., f(x;q) = hq;xi. They suggest using the online gradient descent and the Multiplicative Weights Update (MWU) algorithms to gerfere  $p_{P[T]}$  estimates and show via separate analysis that the resulting estimates have vanishing average losses with respect to <sup>est</sup> and <sup>sub</sup> (at the rateO(1=  $^{p}T$ )) but do not present their regret bounds or analyze <sup>pre</sup> loss. Note that both online gradient descent and MWU algorithm are simply special cases of the online MD algorithm customized to the geometry of the problem domain. Moreover, the setting studied in Bärmann et al. (2017) clearly satis es our Assumption 2 and the perfect information assumption, hence we can utilize  $\delta b^{pr}$ -based OL framework equipped with online MD and directly derive average regret boundsQ(1= T) on <sup>sim</sup>, <sup>sub</sup> and <sup>est</sup>. In addition, as opposed to the simple bilinear form of f considered in Bärmann et al. (2017), our framework can handle more general functions f in the forward problem wheth (x; u) and/orf<sub>2</sub>(q; u) are nontrivial. In this respect, it is of interest to study the case of strongly contexes, u), where through Corollary 1, our framework also leads to regret bound with respect to<sup>6</sup>.

Dong et al. (2018a) study the following problem where linear inq and strongly convex in

$$\min_{x} \frac{1}{2}x^{>} Px h q_{true}; xi : x 2 X (u) :$$
 (5.9)

Here, P is a positive de nite matrix and (u) is the agent's feasible domain determined by the external signal xed as. In this setting, Dong et al. (2018a) propose a regret minimization algorithm utilizing the implicit OL method (Kulis and Bartlett (2010)) with a nonconvex MISOCP oracle. They

focus on the prediction loss<sup>pre</sup>, and establish  $\mathfrak{O}({}^{p}\overline{T})$  bound on  $\mathbb{R}_{T}(f)^{pre}_{t} g_{t2[T]}; f q_{t} g_{t2[T]})$  whenever  $f_{t}^{pre}(q)$  is a convex function of a main limitation of their approach is that the convexity  $\mathfrak{O}_{T}^{pe}$ does not hold in general. Although they identify a technical sufficient condition (Dong et al., 2018a, Assumption 3.3) that can guarantee convexity  $\mathfrak{O}_{T}^{pe}$ , they also remark that this condition is restrictive and very hard to verify in practice even for the simplest form of problem classes. In fact, the only example they identify as satisfying their assumption is when the agent's optimization problem is (5.9) and the seX (u) mustalwayscontain the minimizer of the unrestricted objective minimization problem, i.e.  $p^{-1}q_{true} 2 X$  (u) for all possibleu.

When the agent's problem has the speci c form( $\mathfrak{Sd}9$ ), the algorithm from Dong et al. (2018a) updates $\mathfrak{a}_{t+1}$  as the optimal solution of the following bilevel program:

$$q_{t+1} \coloneqq \underset{q2Q}{\operatorname{argmin}} \quad \frac{1}{2} kq \quad q_t k^2 + h_t ky_t \quad x(q; u_t) k^2 : x(q; u_t) 2 \underset{x}{\operatorname{argmin}} \quad \frac{1}{2} x^> Px \ h \ q; xi : x \ 2 \ X \ (u_t) \qquad :$$

It was shown in Dong et al. (2018a) that when the feasible dot  $\mathbf{M}_{a}(\mathbf{n}_{t})$  is polyhedral, this bilevel program can be represented as a MISOCP. Consequently, the implicit OL algorithm of Dong et al. (2018a) utilizes an MISOCP based solution oracle to genérate  $\mathbf{p}_{2[T]}$ . The main convergence result (Dong et al., 2018a, Theorem 3.2) proves that under their assumptions by choosing the step size  $h_{t} \mu = \mathbf{1}^{p} \bar{\mathbf{t}}$ , the sequence of estimate  $\mathbf{p}_{2[T]}$  generated with the above update yield  $\mathbf{G}(\mathbf{a}^{T})$  bound on the regre  $\mathbf{R}_{T}(\mathbf{f})^{pre}_{t} \mathbf{g}_{t2[T]}; \mathbf{f} \mathbf{q}_{t} \mathbf{g}_{t2[T]})$ .

Note that the format of in (5.9) satis es our Assumption 2, and consequentin is guaranteed to be convex for an X (u). Therefore, our OCO framework based on minimizing regret for loss functions  $\int_{t}^{sim} g_{t2|T1}$  is applicable to (5.9). In addition, in the perfect information setting, through Proposition 5 and Corollary 1, our framework can provide regret bounds with respectod to sim, `est, `sub and` pre, without further structural assumptions on the agent's domain. In contrast, the implicit OL approach of Dong et al. (2018a) for minimizing regret with respect for equires additional conditions on the agent's domain (see (Dong et al., 2018a, Assumptions 3.1, 3.2, 3.3)) in order to guarantee a regret bound. In particular, it is speci cally focused brand provides no insight on other performance measures of interest captured signation of the sub either. Moreover, online MD in our framework uses a much simpler (and computationally faster) rst-order oracle in contrast to the expensive MISOCP oracle in the implicit OL approach of Dong et al. (2018a). One aspect that Dong et al. (2018a) emphasize is the noise in observations: they prove theoretical expected regret bounds with respect to pre when  $y_t$  is a noisy observation of  $(q_{true}; u_t)$ , but the theoretical guarantees of our framework based oinsim do not readily extend to the imperfect information setup. We caution that their analysis for the noisy setup is subject to the restrictive conditions needed to ensure convexity of  $t_{t}^{pre}(q)$ , so it might not provide meaningful performance guarantees in general.

# 5.6 Computational Study

We perform numerical experiments on a practical application that is motivated by a company (learner) seeking to learn about its customer's (agent's) preferences in a changing market. We assume the customer is a rational decision maker, and in any given market situation, her/his action re ects accurately her/his optimal preferences. These experiments do not aim to provide structural insights on speci c instances, rather, our main purpose is to demonstrate the performànicabatsed OCO algorithms from various aspects and the comparison with an alternative aspecta approach in Dong et al. (2018a).

We rst focus on the case when perfect information is available, i.e., there is no noise in learner's observations of the agent's optimal actions, and address three main questions. First, are there notable performance differences among OL algorithms based on different oracles? Second, how do the algorithm performances vary in terms of different loss functions? Third, does the structure of the agent's feasible region affect complexity of the learning problem and the algorithm performances? While discussing these questions, we also compare against existing algorithms from the literature.

In the second part of our numerical study, we examine the robustness of these OL algorithms under imperfect information, i.e., when there is random noise to the learner's observations of the agent's optimal actions. Recall that in the imperfect information setup, our OL based approach is not guaranteed to provide low regret guarantees, so these experiments essentially shed light to their empirical performance in the noisy setup. Despite the lack of theoretical bounds, we can analyze the OL algorithms to explain the observed empirical trends in losses and regrets (see Appendix B.4).

All algorithms are coded in Python 3.8, and Gurobi 8.1.1 with default settings is used to solve the mathematical programs needed for the subproblems associated with the corresponding oracles. We limit the solution time of each mathematical program to be at most 3600 seconds. We have not hit this imposed time limit in any of our experiments. All experiments are conducted on a server with 2.8 GHz processor and 64GB memory.

#### 5.6.1 Problem Instances

We consider a market with products that evolves over a nite time horizon e.g., the product prices change. These changes consequently impact the agent's feasible actions; in this case, agents are customers interested in purchasing the products. Fort *@a[Dh]*, we let ut denote the market parameters relevant to the agent's decisions at period/dhen constraint parameters are xed as ut, an agent's action(q<sub>true</sub>; ut) is an optimal solution to an optimization problem parametrized by ut andq<sub>true</sub>, whereq<sub>true</sub> captures the agent's preferences over the products. We model the agent's optimization problem as a maximization of her/his utility function subject to feasibility constraints. The learner knows the agent's decision problem up to the parameter *q*<sub>RCE</sub> and the learner's goal is to estimateq<sub>true</sub> using observations of the agent's action is negotiated to the market condition at each period/2 [T].

We study two different forms for the agent's utility function.

(a) For direct comparison with Dong et al. (2018a), we examine the case where the agent's utility function has the quadratic form (5.9), i.e., the agent's act(oppue; ut) is given by

$$x(q_{true}; u_t) \coloneqq \underset{x}{\text{argmax}} \quad \frac{1}{2}x^{>} Px + hq_{true}; xi : x 2 X (u_t) ; \qquad (5.10)$$

where P 2  $S_{++}^n$  is a xed positive de nite matrix known by both the learner and the agent and X (u<sub>t</sub>) represents the domain for the agent's feasible actions determined by the market parameters<sub>t</sub>.

(b) We also examine a second setup where the agent has a CES utility function ₩i2h Hence, in periodt, the agent's action(q<sub>true</sub>; u<sub>t</sub>) is given by

$$x(q_{true}; u_t) \coloneqq \underset{x}{argmax} \overset{a}{\underset{i2[n]}{a}} (q_{true})_i x_i^2 \colon x 2 X (u_t) : \quad (5.11)$$

Note that this setup with a CES utility has not been previously studied in an OL framework.

These particular forms of utility functions ((5.10) and (5.11) imply that the dimensions of and x are the same, i.ep = n. Moreover, observe that both of the objective function ((5.11)) and ((5.11)) satisfy Assumption 2, and thus in both cases (q) is convex inq.

To identify the impact of agent's feasible region on the complexity of the problem and on the performance of the learning algorithms, we experiment on a variety of settings (or).

- (i) Continuous knapsadkomain: in this setting, we impose only a budget constraint on the agent:
   X (u<sub>t</sub>) = X <sup>ck</sup>(p<sub>t</sub>; b<sub>t</sub>) := f x 2 R<sup>n</sup><sub>+</sub> : hp<sub>t</sub>; xi b<sub>t</sub>g, where the parameteps 2 R<sup>n</sup><sub>+</sub> correspond to the product prices ankle 2 R<sub>+</sub> is the budget available to the customer during time petriod Note that bothp<sub>t</sub> and b<sub>t</sub> can vary in each time period [T].
- (ii) Continuous polytopedomain: here, we generalize the continuous knapsack domain and model general resource constraints resulting in a polytope as the feasible Xegian = X <sup>cp</sup>(A<sub>t</sub>; c<sub>t</sub>) = f x 2 R<sup>n</sup><sub>+</sub> : A<sub>t</sub>x c<sub>t</sub>g, where all the parameters are nonnegative.
- (iii) Binary knapsackdomain: in this case, we again impose a single budget constraint, but also require that the agent's action is a binary vector (u<sub>t</sub>) = X <sup>bk</sup>(p<sub>t</sub>; b<sub>t</sub>) := f x 2 f 0; 1g<sup>n</sup>: hp<sub>t</sub>; xi b<sub>t</sub>g.
- (iv) Equality constrained knapsa**c** komain: that is  $X(u_t) = X^{eck}(p_t; b_t) := f x 2 R_+^n : hp_t; x_i = b_t g.$

We ran experiments with the utility function  $(\mathbf{5}, 10)$  where we choose the matrix be a positive de nite diagonal matrix and generate each of its diagonal entry is drawing a number from [1;21] uniformly and then normalizing the drawn vector  $(\mathbf{6}_{11}; \ldots; \mathbf{P}_{nn})$  to have a unit 1-norm, and we

also set the domain to  $be^{ck}(p_t; b_t)$ ,  $X^{cp}(A_t; c_t)$ , or  $X^{bk}(p_t; b_t)$ . In the case of CES utility function (5.11), for implementation simplicity, we use instances with the dom ( $aff^{ck}(u_t)$ ).

In all of our experiments, we consider a market with 50 goods. We compare OL algorithms by runningT = 500 iterations on a batch do randomly generated instances for each setting. The domainQ is set be a unit simplex, i.eQ =  $q \ge R_{+}^{n} : a_{i2[n]}q_i = 1$ . We follow the same instance generation methodology used in (Bärmann et al., 2017, Section 4.1) for generating the true parameter  $q_{true}$  and the agent's domaix (u<sub>t</sub>). In each instance  $q_{true}$  is obtained by drawing a random sample from a uniform distribution ove[1;100Q<sup>n</sup> and then normalizing the sampled vector to have a unit `1-norm. In the case of  $c^{k}(p_t;b_t)$ ; X  $b^{k}(p_t;b_t)$ , and X  $e^{ck}(p_t;b_t)$ , for all t 2 [T], the constraint parameters  $p_t; b_t$  are generated randomly as follows: is set  $asq_{true} + 100 \ 1_n + r$ , wherer is an integer vector sampled from discrete uniform distribution over the collection of integer vectors in [  $10; 10]^n$  (numpy.random.randiffunction is used). The budget is selected uniformly random from the range[1;  $a_{i=1}^n(p_t)_i$ ]. In the case of continuous polytope dom2irf<sup>cp</sup>(A<sub>t</sub>; c<sub>t</sub>), we choose  $t_t$  as an m n matrix with m = 10, where each row of the generated in the same way pasand each coordinate in the vectore is drawn uniformly random from from  $from[1; a_{i=1}^m(A_t)_{ii}]$ .

In the OL setup, at time step the learner observes the signal and the agent's action, and uses the information revealed so far to construct the estimate Under perfect information, we havey<sub>t</sub> =  $x(q_{true}; u_t)$  for all t; under imperfect information, we assume  $x(q_{true}; u_t) + e_t$ , where  $e_t$  denotes the random noise. In the noisy setup, each coordinatesinandomly drawn from a uniform distribution over a given range. We consider two ranges to simulate small and large noises, and we choose the range bounds based on the average  $a_{t2[T]} kx(q_{true}; u_t)k$ : the small noises are generated with the ranged=n; d=n], and the large noises are generated with (d].

#### 5.6.2 Implementation Details

In order to compute the estimates  $g_{t2[T]}$ , we implement three OL algorithms and compare their performances. By taking advantage of the convexity of, we design two OL algorithms minimizing regret with respect to<sup>sim</sup>: one equipped with a rst-order oracle and one with a solution oracle. In addition, for comparison with the literature, we implemented another implicit OL algorithm with a solution oracle aimed to minimize the regret associated with i.e., the one from Dong et al. (2018a) that utilizes an MISOCP solution oracle. We provided precisely the same dynamic observations, i.e., the realizations of signals along with the agent's optimum action  $(q_{true}; u_t)$  in each iteration 2 [T] to all of these OL algorithms.

In the case of OL with the rst-order oracle, beca (B) is a unit simplex, we use the negative entropy function  $(q) = a_{i=1}^{n} q_i \ln(q_i)$  as the distance generating function in the de nition of Bregman distance  $V_{q_t}(q)$ . Then, the update rul (6.5) for the OL with rst-order oracle is given explicitly by the following formula, where  $(q_t)$  is a subgradient  $\delta_t^{sim}(q)$  at  $q_t$ .

$$(q_{t+1})_i = \frac{(q_t)_i \exp(-h_t(s_t(q_t))_i)}{a_{j=1}^n (q_t)_j \exp(-h_t(s_t(q_t))_j)}; \text{ for all } i \ge [n]:$$

The main challenge in the implementation of implicit OL algorithm with a solution oracle is whether one can design a computationally tractable solution oracle. When the loss fur(qt) arsed in the implicit OL involvesx(q; ut), as is the case in all loss functions from Section 5.3 except(q), (5.6) is a bilevel program. Bilevel programs are dif cult to solve in general, but can be reformulated into a single level problem using KKT conditions of the inner level problem whenever the inner level is a convex problem. In contrast to this, whether is used as the loss function in an implicit OL algorithm,(5.6) becomes a single level optimization problem quant thus the solution oracle becomes much simpler. Consequently, we study two variants of the implicit OL algorithm based on `sim and` pre that are necessarily equipped with different solution oracles.

In the rst variant, we design an implicit OL algorithm to minimize the regret with respect<sup>in</sup>to Using the squared Euclidean norm as the d.g.f., we arrive at the implicit OL algorithm with a solution oracle that updates  $\mu_{1,1}$  as the optimal solution to

$$q_{t+1} \coloneqq \underset{q2Q}{\operatorname{argmin}} \frac{1}{2} kq \quad q_t k^2 + h_t t^{sim}(q):$$

Under Assumption  $2_{t}^{sim}(q)$  is convex inq, and when the domain is convex, the above problem is a convex program. Therefore, the implementation requires only a convex solution oracle; see Appendix B.3 for the explicit formulations of these oracles.

For comparison purposes, we implement a second variant of the implicit OL algorithm minimizing regret with respect to the loss function<sup>PIe</sup>. By following the same approach taken in Dong et al. (2018a), we use the squared Euclidean norm as the d.g.f., and the resulting solution oracle updates  $q_{t+1}$  with the following bilevel program, where the inner level comput( $e_{T}$ ,  $u_t$ ) used in  $t_t^{pre}(q)$ :

$$q_{t+1} \coloneqq \underset{q2Q}{\operatorname{argmin}} \frac{1}{2} kq \quad q_t k^2 + h_t t^{pre}(q)$$

When the agent is maximizing a concave objective function over a polyhedral dom(ain), we can reformulate the above bilevel program into a mixed integer program (MIP).

Consequently, at time this  $e^{pre}$ -based implicit OL algorithm requires a nonconvex oracle given by the MIP formulation to obtain  $u_{t+1}$ . In the case o(5.10), it was demonstrated in Dong et al. (2018a) that when the domation  $(u_t)$  is polyhedral, the MIP reformulation admits a nice MISOCP structure due to the quadratic objective. For completeness, we provide the MISOCP reformulation of this solution oracle in Appendix B.3. Note that due to the advanced capabilities of modern MIP solvers, the resulting MISOCP still remains computationally tractable whenever the scale of the agent's problem is relatively small.

On the other hand, when the domain of the inner problet  $(\mathbf{u}_t)$  is nonconvex, e.g., when we consider  $\mathbb{X}^{bk}(p_t; b_t)$  that involves binary variables, or when the agent maximizes a nonconcave function over a convex domaix  $(\mathbf{u}_t)$  as in the case  $(\mathbf{5}.11)$  for  $q \ge \mathbb{R}^n_+$ , we no longer have access to KKT based optimality certi cates for the inner problem. Consequently, in such cases, we do not know

how to design a computationally tractable solution oracle, and this is an open question. Therefore, we did not experiment with the pre-based implicit OL algorithm in these cases.

#### 5.6.3 Perfect Information Experiments

In this section, we discuss our numerical results along with plots that highlight our key observations pertinent to the questions of interest to the perfect information case listed at the beginning of Section 5.6.

#### Learning a Quadratic Utility Function

In this case, we assume that the agent's utility function is of f(5m10). We rst compare the performance of the three OL approaches in terms of both average regret performance and the solution time. Figures 5.1 and 5.2 display the means of average (expected) regret performance of the iterates f q<sub>r</sub>q<sub>r2/T1</sub> returned by the OL algorithms with respect to all ve loss functions of interest for the instances where the agent's domain is of continuous knapsack and polytope type, respectively; the shaded areas indicate 95% con dence interval for the means. These means are computed based on all 50 random instances generated in the experiment. In Figures 5.1 and 5.2 (and all the later ones as well), the scale of loss functions naturally differ because the associated regret and loss values are evaluated with respect to different terms present in their corresponding loss de nitions. In terms of the rate at which the average regret converges, in the case of the continuous knapsack instances, Figure 5.1 shows that regardless of the loss function used to evaluate the performance, all three OL algorithms have quite similar performances. This empirical observation is in line with the theoretical regret guarantees given in Section 5.5; recall that this particular domain type was the focus of Dong et al. (2018a), and their analysis presents some restrictive assumptions guaranteeing convergence of their approach on this type of instances. For the continuous polytope instances, Figure 5.2 demonstrates similar performances from the two OL algorithms based soft but highlights the drastically different performance of the implicit OL with the<sup>pre</sup>-minimizing solution oracle, which now leads to average regrets converging to non-zero values. Recall from Section 5.5.2, the regret convergence of an implicit OL algorithm with a solution oracle requires the convexity of the selected loss function. In fact, Dong et al. (2018a) adopt further strong assumption  $(u_t)$  to guarantee that pre is a convex function of q when the agent's problem is of for( $\mathbf{5}$ .10) with X ( $u_t$ ) = X  $^{ck}(p_t; b_t)$ . Our empirical results indicate that these assumptions are indeed hard to satisfy in general and our randomly generated continuous polytope instances do not necessarily satisfy their required assumption. In contrast, since `<sup>sim</sup> is guaranteed to be a convex function of the agent's problem is of for (5.10) regardless of the structure of the agent's domain (ut), the average regrets of their-based implicit OL algorithm with the solution oracle converge to zero for instances with polytope domain as well. Furthermore, we note that the regret convergence of the-based implicit OL algorithm with the solution oracle is slightly better than the OL with the rst-order oracle in both types of instances.



Fig. 5.1 Means of average regret with respect to different loss functions for continuous knapsack instances; the shaded region is 95% con dence interval for the means.



Fig. 5.2 Means of average regrets with respect to different loss functions for continuous polytope instances; the shaded region is 95% con dence interval for the means.

In our numerical study, we observe almost no variation in terms of the solution time of the OL algorithms across different random instances generated from the same setting. We thereby report the time spent by all three OL algorithms on a randomly selected instance from our problem set. When computing the solution time at iterationwe always ignore the time taken to nd(d<sub>true</sub>; u<sub>t</sub>). In iteration of the OL with the rst-order oracle, we account for the time to compute; u<sub>t</sub>) and generated<sub>t+1</sub> using the rst-order oracle. Lastly, in each iteration of both of the and `pre-based implicit OL algorithms with a solution oracle, we account for the time used by the corresponding solution oracles in updating<sub>t+1</sub>. For an arbitrary instance with the continuous knapsack domain, OL with the rst-order oracle nishes in about 0.08 seconds<sup>in</sup>, based implicit OL with the solution oracle takes 2.03 seconds, and based implicit OL with the solution oracle takes 146 seconds. These highlight that, by a signi cant margin, our OL algorithms minimizing regret with respect to the loss function<sup>sim</sup> and utilizing the rst-order oracle and the solution oracle are much more computationally ef cient than the<sup>pre</sup>-based implicit OL with the MISOCP solution oracle one from Dong et al. (2018a).

We next analyze whether the agent's domain structure has any visible effect on the overall regret performance of the OL with the rst-order oracle. From Figures 5.1 and 5.2, we observe that the superiority of the OL with rst-order oracle in terms of the average regret is slightly more obvious in the continuous knapsack setting than in the polytope setting. In Figure 5.3, we compare the means of average regrets for the continuous knapsack instances versus the binary knapsack instances. The regret performances with respect to the loss functions and `est seem to vary only slightly when



Fig. 5.3 Means of average regret with respect to different loss functions contrasting continuous knapsack instances with binary knapsack instances, when OL with the rst-order oracle is used.



Fig. 5.4 Means of average regret (on a logarithmic scale) with respect to different loss functions over T = 500 iterations for (a) continuous knapsack instances, (b) continuous polytope instances, (c) binary knapsack instances, when OL with rst-order oracle is used.

the agent's domain type changes from continuous knapsack to binary knapsack; yet these differences are slightly more noticeable in the case of loss functions and sim.

Lastly, we examine the regret performance of OL with the rst-order oracle with respect to different loss functions(g). From the experiment results from continuous knapsack and continuous polytope instances (respectively Figures 5.1 and 5.2), we observe that the average regret with respect to any of the four loss functions convergences roughly at the same rate, but the corresponding regret bounds differ in their scales. This is not surprising, as the corresponding regrets are based on different terms, e.g., norms of solutions or objective function values, etc. Moreover, recall from Section 5.4 that in the perfect information case the following relationship among the regret bounds with respect to different loss functions (here for simplicity in notation, we den  $\mathbf{R}_{t}^{\text{sim}} := \mathbf{R}_{T}(\mathbf{f}_{t}^{\text{sim}} \mathbf{g}_{t2|T1}; \mathbf{f} \mathbf{q}_{t} \mathbf{g}_{t2|T1}),$ R<sup>sub</sup>+ R<sup>est</sup> <sup>g</sup>R<sup>pre</sup> 0; whereg is the strong convexity parameter of the quadratic etc.) holds:R<sup>sim</sup> objective function in(5.10). Recall that our instance generation guaran  $\operatorname{Re2S}_{++}^n$ , i.e., its smallest eigenvalue  $_{min}(P) > 0$ , and then by the de nition of strong convexity, we ded get  $I_{min}(P)$ . Figure 5.4 displays (on a logarithm scale) the means of the average regrets for different loss functions for estimates generated from the OL with the rst-order oracle on instances in which the agent's domain is either a continuous knapsack, polytope, or a binary knapsack type. These results also con rm the theoretical relationship among the regrets for different loss functions we have established in Section 5.4.

#### 122 Online Convex Optimization Perspective for Learning from Dynamically Revealed Preferences



Fig. 5.5 Means of average regrets with respect to different loss functions for equality constrained knapsack instances; the shaded region is 95% con dence interval for the means.

#### Learning a CES Utility Function

Here, we examine the case when the agent's utility function is of ((5r.ml) and summarize our ndings on the average regrets in Figure 5.5. We note that the OL with the rst-order oracle has a quite noticeable advantage over the implicit OL with a solution oracle in terms of the regret convergence. In this case, on a typical instance, OL with the rst-order oracle takes seconds to complete and `sim-based implicit OL with the solution oracle takes 2 seconds.

#### 5.6.4 Imperfect Information Experiments

We next study the performance of the twein-based OL algorithms when the observations are corrupted with random noise. We test this imperfect information setup on two types of instances where (1) the agent is maximizing a concave quadratic utility function on a continuous knapsack domain, and (2) the agent is maximizing a CES utility function over an equality constrained knapsack domain. We observed that the impact of the noises on the solution time of the OL algorithms was negligible in both of these instance types.

We plot the outcomes differently from the perfect information case. We still show the average regret with respect to<sup>sim</sup> to illustrate that the<sup>sim</sup>-based OL algorithms remain valid under noises. For `<sup>sub</sup> and`<sup>est</sup>, we report the difference between the average loss incurredff**q**pr**x**(q<sub>t</sub>; u<sub>t</sub>)g and the average loss from the difference between the latter evaluates the loss due to imperfect information. Lastly, we plot the average squared norm distance bet**fixe(ep**; u<sub>t</sub>)g and f x(q<sub>true</sub>; u<sub>t</sub>)g as a measure of prediction loss.

We report the results for when the agent's problem has the ( $\mathbf{5}$ r $\mathbf{n}$  $\mathbf{0}$ ) with the domainX ( $u_t$ ) = X <sup>ck</sup>( $u_t$ ) in Figures 5.6 and 5.7. First, we observe that under small noises, the average regret with respect to <sup>sim</sup> has a similar convergence trend as in the perfect information experiment results plotted in Figure 5.1, and the convergence appears to be slower under large noises. For the three loss based measures, both OL algorithms lead to decreasing trends; in particular, the decreasing average squared norm distance between the predicted actions and the true actions indicates that these OL algorithms' predictions of the agent's actions are becoming more accura**T**ei**nsr**eases. Not surprisingly,



Fig. 5.6 Learning a quadratic utility function under small noises: means of selected performance measures oveT = 500 iterations for continuous knapsack instances; the shaded region is 95% con dence interval for the means.



Fig. 5.7 Learning a quadratic utility function under large noises: means of selected performance measures oveT = 500 iterations for continuous knapsack instances; the shaded region is 95% con dence interval for the means.

the effects of large noises on performances are more noticeable, and it is worth noting that the OL algorithm with rst order oracle outperforms the one with solution oracle under large noises.

We next discuss results in the CES setup, i.e., when the agent's problem has the formwith the equality constrained knapsack domain,  $\aleph e_{i}(u_{t}) = X^{eck}(u_{t})$ , under small noises in Figure 5.8, and under large noises in Figure 5.9. Similar to the instance of learning a quadratic utility function, both <sup>sim</sup>-based OL algorithms lead to converging average regrets with resperied and decreasing average loss gaps, which are consistent with our above theoretical analysis. We note that the magnitudes of noises appears to have much smaller effects on the learning performance, as both gures show similar output.

We study the online preference learning task, where a learner wishes to learn a non-strategic agent's private utility function through observing the agent's utility-maximizing actions in a changing environment. We adopt an online inverse optimization setup, where the learner observes a stream of agent's actions in an online fashion and the learning performance is measured by regret associated with a loss function. Due to the inverse optimization component, attaining or proving convexity is dif cult for all of the usual loss functions in the literature. We address this challenge by designingloss function that is convex under relatively mild assumptions. We establish that the regret with respect to `sim also bounds the regret with respect to the three classical loss functions is functionally used in the inverse optimization literature. This allows us to design a exible OL framework that

#### 124 Online Convex Optimization Perspective for Learning from Dynamically Revealed Preferences



Fig. 5.8 Learning a CES utility function under small noises: means of selected performance measures over T = 500 iterations for equality constrained knapsack instances with 50; the shaded region is 95% con dence interval for the means.



Fig. 5.9 Learning a CES utility function under large noises: means of selected performance measures over T = 500 iterations for equality constrained knapsack instances with 50; the shaded region is 95% con dence interval for the means.

Mirror Descent) has signi cant advantages in terms of regret performance and solution time over other OL algorithms from the literature.

# Chapter 6

# Modeling and Inferring Dynamic Ethical Judgments Around High-stakes Allocations

# 6.1 Introduction

Al and ML tools are permeating various facets of life in modern society. These powerful technologies are applied in numerous policy domains to inform or make consequential decisions impacting people's lives. In particular, they increasingly inform or automate high-stakes ation decisions in domains such as lending, employment, education, and healthcare. The past decade has witnessed an overwhelming body of evidence establishing the need for AI and ML to re ect social values and moral ideals, such as justice and fairness. However, translating these principles into computationally tractable and veri able forms has proven to be challenging. As an example, consider outcome fairness through ML-based decisions. Existing mathematical formulations of fairness for ML (e.g., statistical parity) have primarily treated it as static, one-size- t-all concept-captured in terms of ad hoc predictive parity conditions across socially salient groups. The algorithmic interventions designed to guarantee these notions are similarly opicandone-shot These mathematical treatments fail to account for the ontext-dependermature of moral ideals and the incremental, local remedies often needed to ensure justice.

A growing body of work has called on the AI-ethics community to bring stakeholders' judgments into the process of formulating moral values and principles for AI. Following this human-centric view toward ethics, we rst observe that peoplets oral judgments are seldom invariable across situations and contexts. In particular, they are often informed by numerous consideration, including but not limited to:

<sup>&</sup>lt;sup>1</sup>Unless otherwise stated, we use the term `people' to refistrate holders not the public.

- Who stands to receive harm/bene t as a result of the allocation and what is the extent of the harm/bene t?
- What is the policy's historical backdrop (e.g., does the decision impose additional burdens on historically under-served individuals and groups?)
- What are the future implications of various policy alternatives (e.g., will incurring an extra burden today improve the society's overall welfare and prosperity tomorrow?)

Even absent of disagreement regarding the true answer to the above questions, how they inform allocations could be a genuine point of normative disagreement. So an ethically-minded social planner in charge of the allocation will need to understand and aggregate stakeholders' moral judgments. Additionally, allocational policies are often sequential in nature. Through the decision-making policies the planner enacts today, he/she can impact and shift the above determinants of moral judgments, giving rise to a different state tomorrow. Understanding such dynamics allows the planner to design effective interventions that stir society toward a desirable social state over time.

We consider a setting in which a social planner or policymaker has to make make a consider a setting in which a social planner or policymaker has to make a construct the social planner has t regarding the llocation of scarce resources in high-stakes social domains. Our goal is to understand stakeholders' moral judgments regarding sadbcational policies In particular, we aim to evaluate the sensitivity of these judgments to the context/history of allocations and their perceived future impact on various socially salient groups. As a concrete example, consider a central agency in charge of allocating scarce medical resources (e.g. vaccines, hospital beds) to patients during a viral epidemic. The allocation decisions shape the overall environment by affecting the welfare of patients, and the urgency and the demand for the particular resources. A patient is associated with many attributes which describe his/her demographic backgrounds, medical characteristics, lifestyle, etc. Depending on the social context, particularly how severe the health issue is and how scarce the needed healthcare resource is, people may have different opinions about which of these attributes should be considered in the policy making process. As we have observed during the current COVID-19 pandemic, age and underlying medical conditions are critical factors to consider to prioritize more vulnerable people. We can imagine that as vaccines roll out and treatment improves, people's emphasis on these factors may reduce; this is an example of external factors changing the social context, which further in uences people's moral judgments.

#### 6.1.1 Our Approach and Results

We propose a mathematical model to capture and infer stakeholders' potentially-dynamic moral preferences. Our model utilizes a Markov Decision Process (MDP) to represent the sequential decision making context, more speci cally, the sequential allocation of scarce resources. We suppose the reward function of the MDP re ects a stakeholder's moral preferences on the resource distribution decisions, then an allocation policy viewed as more ethically/morally acceptable will have a higher reward. Based on this moral preference model, we can infer a stakeholder's ethical/moral judgments

Moral principles from bioethics. There is a robust history of debate in bioethics about normative principles for the distribution of scarce medical resources across a variety of contexts Cohen et al.[2009], Bayer et al.[2011], Wertheimer and Emanuel[2006]. These contexts can range from non-emergency contexts, such as organ donation and hospital triage, to emergency contexts, such as natural disasters and pandemics (where the health of an entire population is impacted). One solution to this problem is to simply distribute resources randomly by either a lottery or other randomization policy Peterson[2008]. There is an Egalitarian justi cation for these principles, and they are commonly implemented in non-emergency contexts under " rst come, rst served." However, in emergency contexts, the policies of hospitals and governments will almost always favor certain groups over others. As Savulescu et al [2020] declare: "there are no egalitarians in a pandemic." We are interested in the normative principles which are used to justify these types of emergency policies.

We group these normative principles for fair distribution into three broad categories:

• Prioritarian:

· Consequentialist:

This approach attempts to maximize the overall bene ts and minimize overall losses, which might mean favoring those who have "most to lose" (high expected value, given treatment), with these impacts existing at a local (family) or global (society) level.

• Desert:

This approach favors those who are owed compensation or reward because of their eligible features, usually some form of qualifying past behavior like lifestyle choices and effort.

The rst two categories will often converge on allocating resources to the most vulnerable, so long as the consequences of patient deaths are roughly comparable, and the probability of survival without treatment is usually inversely proportional to the probability of survival with treatment. However, these assumptions can and do come apart. For instance, in an in uential article in the New England Journal of Medicine from early in the COVID-19 pandemic, Emmanuel et al. [2020] provide Consequentialist recommendations which go against a purely Prioritarian focus on vulnerability:

Operationalizing the value of maximizing bene ts means that people who are sick but could recover if treated are given priority over those who are unlikely to recover even if treated and those who are likely to recover without treatment. Because young, severely ill patients will often comprise many of those who are sick but could recover with treatment, this operationalization also has the effect of giving priority to those who are worst off in the sense of being at risk of dying young and not having a full life ...

are relevant does change over time. In the famous case-study of Memorial Medical Center in New Orleans, when not all patients could be evacuated from the storm-ravaged hospital, the staff changed their normative principles over the course of several days from a Prioritarian to a Consequentialist approach. It is crucial to understand the ways in which not only static but also dynamic features of a situation play a role in the development of policies for resource allocation.

# 6.2 Problem formulation

We consider a Markov Decision Process (MDP) model representing the sequential allocation of scarce resources. For simplicity, we work with a single type of resource, such as, hospital beds during an epidemic. We suppose the allocation policy proceeds in phases: in each phase, the policy prioritizes people with certain features, such as speci c age ranges, medical histories, and occupations. A stakeholder's moral preference is re ected via her/his opinions on who should be prioritized next. As the allocation unfolds over time, stakeholders' moral preferences may shift in line with the evolving societal context.

# 6.2.1 MDP Model

A standard MDP is de ned with a tuple A = hS; A; P, Ri, where S is the set of states, is the set of actions P: S A S ! [0; 1] is the transition probability function, with  $F(s^0; s; a)$  denoting the probability of transitioning to state from states by taking actiona, R: S ! R is a state-based reward function.

A states 2 S in our running example describes the current state of affairs—that is, the current pro le of immunity and vulnerability to the virus across variogsoups Groups are de ned as follows:. We considen ethically relevant features, and uses these features to assign people to groups. We de nes with the group features  $\mathbf{s}_{t} = (\mathbf{s}_{t,1}; \ldots; \mathbf{s}_{t;n})$ , where  $\mathbf{s}_{t;i}$  is a vector representing the pro le of groupi at stept. Note that we do not require mutually exclusive features, namely, a person may have multiple features, thus belonging to multiple groups simultaneously. In our running example,  $\mathbf{s}_{t;i} = (\mathbf{x}_i^t; \mathbf{v}_i^t; \mathbf{d}_i^t)$ , which respectively are the proportion of group that, at step, have received the resource and thus immune to the virus, are susceptible to infecting the virus, and have passed away from contracting the virus. As we are not aiming to use the most of cient state representations, it may be possible to reduce the size  $\mathbf{S}$  if with alternative state de nitions.

An action  $a_t 2 A$  represents the current time step's resource allocation decision, that is, how resources will be distributed across the groups at time. Similar to the state notation, we can de ne  $a_t$  with the group based action feature ( $a_1^t$ ;...; $a_n^t$ ), where  $a_i^t$  is the proportion of the current step's available resources allocated to group to group the special case when the policy assigns all available resources to group we will have  $a_i^t = 1$  and  $a_i^t = 0$  for all j  $\in$  i.

Next, to de ne the transition probabilities, we suppose the MDP model is deterministic, namely  $P(s^0_1s; a) \ge f \ 0; \ 1g \ for \ all \ s^0, \ s \ and \ a.$  The deterministic assumption ts the resource allocation setting



Fig. 6.2 Piecewise linear rewards for group 1 (left) and group 2 (right)

importance ranking among groups remain unchanged throughout state shifts. Therefore, the simple linear reward is insuf cient for modeling changing priorities across groups over time.

A more exible alternative is to use a spline reward function (i.e. a piece-wise polynomial function) to model changes in priorities as the result of changes to the underlying state of the world. Suppose a spline reward function function for pieces  $R_1; \ldots; R_m$ . Each piece represents a different type of moral preference in its corresponding domain. The domain boundary points therefore indicate shifts in preferences. The complete feature space considered in de ning rewards is split into these disjoint domains. In our running example, we de ne a piece-wise linear reward functivith respect to the resource allocation features for each group, the interval[I<sub>i</sub>; u<sub>i</sub>] for possiblex<sub>i</sub> values is covered by ordered, disjoint subinterva[ $\xi = c_i^{(0)}; c_i^{(1)}; [c_i^{(1)}; c_i^{(2)}]; \ldots; [c_i^{(m-1)}; u_i = c_i^{(m)}]$ . In addition, we have slope vectors  $w^{(1)}; \ldots; w^{(m)}$  characterizing all pieces. For  $= 1; \ldots; m$ , when  $(x_1; \ldots; x_n) \ge [c_1^{(k-1)}; c_1^{(k)}] : \ldots : [c_n^{(k-1)}; c_n^{(k)}]$ , the reward de nition follows the there and prioritizes the groups with weighted<sup>(k)</sup>. We next show an example of a 3-piece reward function with two groups: we illustrate the reward of each group separately for clarity.

Based on this general spline reward, we focus on a simple two-piece case where the rst piece is linear and the second piece is constant. We select this particular reward de nition to reduce the needed number of parameters.

$$\mathsf{R}(\mathsf{s};\mathsf{w};\mathsf{c}) \coloneqq \mathsf{R}(\mathsf{x}_1;\ldots;\mathsf{x}_n;(\mathsf{w}_1;\ldots;\mathsf{w}_n);(\mathsf{c}_1;\ldots;\mathsf{c}_n)) \coloneqq \overset{\mathsf{n}}{\underset{i=1}{\overset{\mathsf{n}}{\overset{\mathsf{n}}{\overset{\mathsf{w}_i}}}} \mathsf{w}_i \min \mathsf{x}_i;\mathsf{c}_i \mathsf{g} \tag{6.1}$$

In this reward function, we skip the piece indexwite;  $c^{(k)}$  to simplify notations. We can interpret the reward as follows. Throughout the allocation process, give further increases as the group receives more resources. In the beginning, allocations to give further with the weight of w<sub>i</sub>. After x<sub>i</sub> increases to exceed, the cumulative reward will stay xed and ic, namely, further allocations to group gain 0 additional rewards. We can interpret as the resource level that is considered suf ciently high for groupso that additional resources given to the group will not be rewarded. Using the two-piece reward function, we can model changes in the importance ranking among groups.

Based on the MDP model, we use a trajectory  $(s_0; a_0; \ldots; s_{T-1}; a_{T-1}; s_T)$  to denote a sequence of allocation decisions and the resulting state changes. The trajectory reward is a discounted sum of the immediate reward gained at each state of the trajectory, naR(elyw, c) =  $a_{t=1}^T g^t r(s_{t+1}; s; w, c) = a_{t=1}^T g^t (R(s_{t+1}; w, c) - R(s; w, c))$  with g > 0 as the discount factor R(t; w, c) represents a stake-holder's perceived gain from shifting the societal context from  $s_{T+1}$  through a sequence of allocation policies  $a_1; \ldots; a_T$  overtime.

#### 6.2.2 Fairness Preference Model

We take the perspective of a policy planner who wishes to infer a stakeholder's preference by learning her/his reward function. We follow the framework of preference-based reward learning studied in literature, e.g. B y k et al. (2019); Sadigh et al. (2017). The reward learning involves an interactive process where the planner asks a stakeholder to answer a sequence of queries, and uses the query answers to iteratively update the estimate**s**voafndc. Each query is a comparison between two trajectories both starting from the same initial state and of equal length/number of phases. Again, example from the tutorial could help. We next show a sample query in the setting of our running example.





We use the Bradley-Terry model, a standard human choice model Luce (2012), to represent a stakeholder's moral preference model. Suppose a query asks to compare trajectamiets. If a stakeholder preferts, to  $t_2$ , we denote it with  $t_1 = t_2$ . A stakeholder with reward function (s; w, c) will choose  $t_1$  as more preferable, namely, more ethically acceptable with probability:

$$P(t_1 \quad t_2 j w, c) = \frac{expR(t_1; w, c)}{expR(t_1; w, c) + expR(t_2; w, c)}:$$
(6.2)

approximation.

$$\begin{split} J_{t} &= \arg \max_{Q \geq Q} H(W;C) \quad E_{u \mid U_{w \mid c} \mid (Q)} H(W;Cju) \\ &= \arg \max_{Q \geq Q} \quad E_{(w;c) \mid P(W;C)} \log P(w;c) + E_{(w;c) \mid P(W;C);u \mid U_{w \mid c} \mid (Q)} \log P(w;cju) \\ &= \arg \max_{Q \geq Q} E_{(w;c) \mid P(W;C);u \mid U_{w \mid c} \mid (Q)} [\log P(w;cju;Q) \mid \log P(w;c)] \\ &= \arg \max_{Q \geq Q} E_{(w;c) \mid P(W;C);u \mid U_{w \mid c} \mid (Q)} [\log P(u;Qjw;c) \mid \log P(u)] \\ &= \arg \max_{Q \geq Q} E_{(w;c) \mid P(W;C);u \mid U_{w \mid c} \mid (Q)} \log P(u;Qjw;c) \mid \log P(u;Qjw;c) \mid \log P(u;Qjw_{i};c_{i}) \\ &= \arg \max_{Q \geq Q} E_{(w;c) \mid P(W;C);u \mid U_{w \mid c} \mid (Q)} \log P(u;Qjw;c) \mid \log \frac{1}{M} \bigotimes_{(w_{i};c_{i}) \geq W} P(u;Qjw_{i};c_{i}) \\ &= \arg \max_{Q \geq Q} \frac{1}{M} \bigotimes_{(w_{j};c_{j}) \geq W} E_{u \mid U_{w \mid c} \mid (Q)} P(u;Qjw_{j};c_{j}) \log M \frac{P(u;Qjw_{j};c_{j})}{a_{(w_{i};c_{i}) \geq W} P(u;Qjw_{i};c_{i})} : \end{split}$$

# 6.4 Experimental Design

We carry out human-subject experiments on Amazon Mechanical Turk to understand people's moral assessments of allocation policies in the hypothetical scenario. We apply the active learning frame-work introduced in Section 6.3 to elicit our participants' moral preferences regarding the sequential allocation of scarce medical resources in a hypothetical epidemic.

#### 6.4.1 Scenario and Setup

The participant rst receives basic information about the study, including study background and goals, the hypothetical decision-making context, and the type of questions they need to respond to if they decide to participate. Then the participant answers a questionnaire consisting of 20 queries. Each question/query consists of comparison between two trajectories each representing a different allocation policy.

The following information is presented to participants who wish to participate in our study.

Background and Task Description

The goal of this survey is to understand your moral judgments regarding the sequential allocation of scarce medical resources.

#### Hypothetical Scenario

Imagine a viral epidemic that has infected millions of people around the world leading to a disease with a very high mortality rate. There is currently only a single highly effective for the disease—those who receive the cure will fully recover (if currently infected) and become immune to the virus in the future. Unfortunately, the number of cure doses that can be produced and administered every month is **iso** it ed public health of cials need to decide which groups should be prioritized at any given time. In the questionnaire that follows, we will present you with additional information about several possible states of the epidemic and ask you to choose your preferred allocational policy between two cure allocation policies.


Fig. 6.5 Rewards based on estimates from survey responses

viewed as suf ciently cured. The medically vulnerable and essential workers have similar reward trends as the elderly. Lastly, the military personnel group has slightly higher rewards than the medically vulnerable and essential workers in the beginning of cure allocation, but drops to the least prioritized for majority of the allocation process.

We also analyze the aggregate preferences on a more re ned level. In Fig. 6.5b, we focus on the top priority group throughout the cure allocation process. Each stacked bar shows the distribution of responses prioritizing each group when all groups have the same given cured proportion. We observe the same preference dynamics as in Fig. xx. Caregivers are the most common top priority at lower cured levels. As more people are cured, it becomes increasingly ethical to prioritize essential workers, public-health compliant and elderly groups.

To explore further, we compare groups pairwise to understand the ethical position of each group against another group at different cured levels. We highlight observations from three pairs: elderly vs. medically vulnerable, caregivers vs. essential workers, public health compliant vs. military personnel. As we discuss in Section prioritizing either group in these groups respectively re ect prioritarianism, restorative justice and distributive justice. Comparing each pair could provide further insights on how these ethical values play out in people's ethical judgments. Between the elderly and the medically vulnerable, the respondents are roughly equally split in preferring either group and the prioritization trend is stable at all cured levels. This indicates that when stakeholders care about favoring the worse off in allocation, they may be indifferent to the different reasons leading people to the disadvantaged positions. Between caregivers and essential workers, more respondents prioritize the former throughout the cure allocation process, but the gaps are small. After the cured proportion reaches 40% in both groups, essential workers receive slightly stronger prioritization. Lastly, the public-health compliant are noticeably prioritized over military personnel at all cured levels, and the gaps grow wider as more people receive cures.



Fig. 6.6 Number of responses prioritizing each group in a pair

that the distinctions among what contribute to people's instrumental values and desertness could impact a stakeholder's moral judgments.

Discussion.In a number of surveys, respondents mention at least once that they prefer a policy because it bene ts multiple groups at once. Recall that we illustrate the impacts of each three-phase allocation policy with a stacked bar plot, and each phase the proportion of people receiving cures is highlighted in all the groups. Due to the overlaps among groups, as we explain previously, the cure allocation to one group effectively affects all groups through the overlaps. In our experiment instance, the six groups have different overlapping pro les. For example, the public-health compliant group has a larger overlap with the medically vulnerable group than the essential worker group. Suppose a respondent views curing the medically vulnerable as the most ethical action, he/she is likely to nd curing the public-health compliant in a phase more ethically acceptable than curing the essential worker, as the former would lead to a more prominent cured effect for the medically vulnerable.

Our current experiment design does not prompt people to state their thresholds explicitly. When asked to explain their choices, respondents could easily indicate their weight preferences for the groups by indicating which groups they believe should be prioritized. On the contrary, it is more dif cult for respondents to specify the exact cure thresholds that would make a group suf ciently protected so that further prioritization is unnecessary. Nevertheless, we receive answers mentioning that a policy is or is not preferred due to a certain group already has a relatively high cured level, which demonstrate that respondents use thresholds implicitly when formulating ethical judgments.

#### 6.6 Conclusion

- Chen, V. and Hooker, J. N. (2020a). Balancing fairness and ef ciency in an optimization model. preprint, 2006.05963.
- Chen, V. and Hooker, J. N. (2020b). A just approach balancing Rawlsian leximax fairness and utilitarianism. InAAAI/ACM Conference on AI, Ethics, and Society (Alpa) es 221–227.
- Chiarandini, M., Fagerberg, R., and Gualandi, S. (2019). Handling preferences in student-project allocation. Annals of Operations Resear 275(1):39–78.
- Chouldechova, A. (2017). Fair prediction with disparate impact: A study of bias in recidivism prediction instruments big data 5(2):153–163.
- Chouldechova, A., Benavides-Prado, D., Fialko, O., and Vaithianathan, R. (2018). A case study of algorithm-assisted decision making in child maltreatment hotline screening decision on Fairness, Accountability and Transparenpages 134–148. PMLR.
- Chouldechova, A. and Roth, A. (2020). A snapshot of the frontiers of fairness in machine learning. Communications of the ACM3(5):82–89.
- Christiano, P., Leike, J., Brown, T. B., Martic, M., Legg, S., and Amodei, D. (2017). Deep reinforcement learning from human preferencesXiv preprint arXiv:1706.03741
- Cohen, W. W., Schapire, R. E., and Singer, Y. (2009). Principles for allocation of scarce medical interventions.Lancet 373:423–431.
- Corbett-Davies, S. and Goel, S. (2018). The measure and mismeasure of fairness: A critical review of fair machine learningarXiv preprint arXiv:1808.00023
- Cowell, F. A. (2000). Measurement of inequality. In Atkinson, A. B. and Bourguignon, F., editors, Handbook of Income Distribution 1, pages 89–166. Elsevier.
- Dalton, H. (1920). The measurement of the inequality of incorfies nomic Journal 30:348-461.
- Deepak, P. and Abraham, S. S. (2020). Representativity fairness in clusteril Mgeb Bci pages 202–211.
- Dieterich, W., Mendoza, C., and Brennan, T. (2016). COMPAS risk scales: Demonstrating accuracy equity and predictive parity. Technical report, Northpointe Inc. Research Department. 8 July.
- Dong, C., Chen, Y., and Zeng, B. (2018a). Generalized inverse optimization through online learning. In Proceedings of the 31st Conference on Neural Information Processing Systems S'18, pages 86–95.
- Dong, C. and Zeng, B. (2020). Expert learning through generalized inverse multiobjective optimization: Models, insights, and algorithms. In III, H. D. and Singh, A., editensceedings of the 37th International Conference on Machine Learninguiume 119 of Proceedings of Machine Learning Researchpages 2648–2657. PMLR.
- Dong, J., Roth, A., Schutzman, Z., Waggoner, B., and Wu, Z. S. (2018b). Strategic classi cation from revealed preferences. Proceedings of the 2018 ACM Conference on Economics and Computation (EC '18), pages 55–70. ACM.
- Donini, M., Oneto, L., Ben-David, S., Shawe-Taylor, J., and Pontil, M. (2018a). Empirical risk minimization under fairness constraints. **Proceedings of the 32nd International Conference on Neural Information Processing System** PS'18, page 2796–2806, Red Hook, NY, USA. Curran Associates Inc.

- Donini, M., Oneto, L., Ben-David, S., Shawe-Taylor, J. S., and Pontil, M. (2018b). Empirical risk minimization under fairness constraints. Advances in Neural Information Processing Systems pages 2791–2801.
- Dwork, C., Hardt, M., Pitassi, T., Reingold, O., and Zemel, R. (2012a). Fairness through awareness. In Proceedings of the 3rd innovations in theoretical computer science conferences 214–226.
- Dwork, C., Hardt, M., Pitassi, T., Reingold, O., and Zemel, R. S. (2012b). Fairness through awareness. In Symposium on Innovations in Theoretical Computer Science (JTages 214–226.
- Eisenhandler, O. and Tzur, M. (2019). The humanitarian pickup and distribution pro**Opernations** Research67:10–32.
- Emanuel, E., Persad, G., Upshur, R., Thome, B., Parker, M., Glickman, A., Zhang, C., Boyle, C., Smith, M., and Phillips, J. (2020). Fair allocation of scarce medical resources in the time of covid-19. New England Journal of Medicin@82:2049–2055.
- Fan, R.-E., Chang, K.-W., Hsieh, C.-J., Wang, X.-R., and Lin, C.-J. (2008). Liblinear: A library for large linear classi cationthe Journal of machine Learning research1871–1874.
- Fan, R.-E., Chen, P.-H., and Lin, C.-J. (2005). Working set selection using second order information for training support vector machines ournal of machine learning researcos (Dec):1889–1918.
- Frankfurt, H. G. (2015)On Inequality Princeton University Press.
- Freeman, S., editor (2003) The Cambridge Companion to Raw Gambridge University Press.
- Gabriel, I. (2020). Arti cial intelligence, values, and alignme Minds and Machines 30:411-437.
- Gauthier, D. (1983)Morals by AgreementOxford University Press.
- Gini, C. (1912). Variabilità e mutabilità P. Cuppini, reprinted 1955 in E. Pizetti abd T. Salvemini, eds., Memorie di metodologica statistic Rome: Libreria Eredi Virgilio Veschi.
- Glasmachers, T., Igel, C., Bennett, K. P., and Parrado-Hernández, E. (2006). Maximum-gain working set selection for symslournal of Machine Learning Researcit(7).
- Hardt, M., Price, E., and Srebro, N. (2016a). Equality of opportunity in supervised learning. In Proceedings of 30th International Conference on Neural Information Processinges 3323–3331.
- Hardt, M., Price, E., and Srebro, N. (2016b). Equality of opportunity in supervised learning. In Advances in neural information processing systemages 3315–3323.
- Harsanyi, J. C. (1977) Rational Behavior and Bargaining Equilibrium in Games and Social Situations Cambridge University Press.
- Hazan, E. (2019). Introduction to online convex optimizatiarXiv preprint arXiv:1909.0520.7
- Heidari, H., Ferrari, C., Gummadi, K., and Krause, A. (2018). Fairness behind a veil of ignorance: A welfare analysis for automated decision makingAdvances in Neural Information Processing Systemspages 1265–1276.
- Heidari, H., Loi, M., Gummadi, K. P., and Krause, A. (2019). A moral framework for understanding fair ml through economic models of equality of opportunity. **Proceedings** of the Conference on Fairness, Accountability, and Transpare, **p**ages 181–190.

# Appendix A

# Appendix to Chapter 3

## A.1 Pseudocode

### A.1.1 Fair Sequential Minimal Optimization for Fair General SVM

## Algorithm 1 Fair SMO

Input: Iraining dataset $D = f x_i; y_i; z_i g_{i_2[n]}$ , fairness constraints parameteps; $c_i; c_i g_{i_2[k]}$ , kernel function K, stopping
criteria tolerance.
Output: mg satisfying KKT conditions (3.26) up to tolerance
1: Initialize m <sup>(0)</sup> ;g <sup>(0)</sup> to all 0
2: Initialize gradientsG, with $G_m = \tilde{N}_m g(m, g) = 1$ for all i 2 [n] and $G_g = \tilde{N}_g g(m, g) = (e_i - e_i)$ for all I 2 [k].
3: for t = 0; 1; 2; ::: do
4: if m(m <sup>(t)</sup> ;g <sup>(t)</sup> ) M(m <sup>(t)</sup> ;g <sup>(t)</sup> )+t and jg <sup>(t)</sup> G <sub>g</sub> j t for all I 2 [k] then
5: Returnm <sup>(t)</sup> ; $g^{(t)}$ and terminate.
6: else
7: Choose working $se(i; j) = Select B(m^{(i)}; g^{(t)}; I; G)$ .
8: if j G 1 then
9: Update working set variable $\mathfrak{s}(t^{(t+1)}; \mathbf{m}^{(t+1)}) = Updat_{\mathfrak{s}(\mathbf{m}^{(t)}; fi; jg; G)}$ .
10: UpdateG using (3.33).
11: end if
<sup>12:</sup> Updateg variables and obtain corresponding gradige( $ts^{1}$ ; G = Updateg( $g^{(t)}$ ; G).
13: end if
14: end for

# Appendix B

# Appendix to Chapter 5

#### B.1 Convexity Status of pre and `est

In this appendix, we examine the convexity stat@\$ and`<sup>est</sup> under our Assumption 2. Recall that we already establish in Lemma 3 that under Assumption<sup>1</sup>(2q) is convex inq. On the other hand, `<sup>pre</sup> and`<sup>est</sup> are not guaranteed to be convexqueven under Assumption 2 and even when agent's problem is a one dimensional optimization problem.

Example 1. Suppose Q = [1; 1] and  $q_{true} = 1$ , and an agent's forward problem  $risin_x f qx : x 2 X g$ . We consider a convex domation = [1; 1]. Let  $x(q) := \arg \min_x f qx : x 2 X g$ , i.e., the set of optimizers of the agent's problem for given. Then, we easily deduce that the agent's optimal action(s) at a given are: x(q) = 1 if q > 0, x(q) 2 X if q = 0, and x(q) = 1 if q < 0. Speci cally, this implies  $x(q_{true}) = 1$ .

Consider  $q_1 = 1$ ,  $q_2 = 1$  and  $l = \frac{1}{4}$ , then  $\tilde{q} = lq_1 + (1 l)q_2 = \frac{1}{2}$ . By the format of x(q) in this problem and the de nition  $\tilde{o}^{\text{rest}}$ , we observe that

$$e^{st}(q_1) = q_{true}(x(1) \quad x(q_{true})) = 0;$$
  
 $e^{st}(q_2) = q_{true}(x(1) \quad x(q_{true})) = 2;$   
 $e^{st}(q) = q_{true}(x(1=2) \quad x(q_{true})) = 2:$ 

Therefore, we deduce<sup>st</sup>( $\tilde{q}$ ) > I `<sup>est</sup>( $q_1$ )+(1 I)`<sup>est</sup>( $q_2$ ) that shows that<sup>est</sup> is not a convex function of q. Similarly, in the case of <sup>pre</sup>, we arrive at

$$r^{\text{pre}}(q_1) = (x(1) \quad x(q_{\text{true}}))^2 = 0;$$
  
 $r^{\text{pre}}(q_2) = (x(1) \quad x(q_{\text{true}}))^2 = 4;$   
 $r^{\text{pre}}(\tilde{q}) = (x(1=2) \quad x(q_{\text{true}}))^2 = 4:$ 

Similarly, we arrive at  $pre(\tilde{q}) > 1$   $pre(q_1) + (1 1)^{re}(q_2)$  and hence conclude is not convex.



(a) quadratic utility function, con(b) quadratic utility function, con(c) CES utility function, equality tinuous knapsack constrained knapsack

Fig. B.1 Means of average SP regrets otver 500 iterations; the shaded region is 95% con dence interval for the means.

average regret with respect  $\delta^{ub}$  follows a similar convergence trend as the average SP regret. Even though the online SP algorithm appears to be suf cient to guarantee the convergence of the average regret with respect  $\delta^{ub}$  in this case, we observe that there is still a lack of convergence guarantees for the other regrets, in particul  $\mathbf{R}_{T}^{sim}$  and  $\mathbf{R}_{T}^{pre}$  are increasing throughout all iterations.

#### B.2.2 OCO with Solution Oracle

To apply the implicit OL algorithm with respect  $t\delta^{ub}$ , we need a solution oracle for solving  $\min_{q2Q} V_q(q_t) + h_t \overset{sub}{t}(q)$ . Even though  $\overset{sub}{t}(q)$  is convex inq, the required solution oracle may be computationally intractable, because the format  $\overset{sub}{t}(q)$  makes the underlying optimization model a bilevel program with both and x as the decision variables. Recall from Appendix B.3, the solution oracle for  $\overset{sim}{s}$ -based implicit OL algorithm is naturally a convex program with the convenient structure of  $\overset{sim}{s}$ . The  $\overset{pre}{s}$ -based solution oracles are bilevel programs, which may not be tractable in general, as we have discussed for the case of learning a CES utility function.

Similar computational dif culties occur for using<sup>ub</sup>-based solution oracles. We again suppose that the squared Euclidean norm is used as the distance generating function in the solution oracle, then the <sup>sub</sup>-based implicit OL algorithm updates, 1 by solving the following bilevel program:

$$q_{t+1} = \underset{q2Q}{\operatorname{argmin}} \frac{1}{2} kq \quad q_t k^2 + h_t (hq; c(y_t) \quad c(x(q; u_t))i \quad f_1(x(q; u_t); u_t))$$

where

$$x(q;u_t) 2 \operatorname{argminf}_{v} f(x;q;u_t) : g(x;u_t) = 0; x 2 X g:$$

To obtain a single-level reformulation, we can again apply the KKT conditions to the inner problem to represent  $(q; u_t)$  in the constraints. We demonstrate the following reformulations for our experiment instances to further illustrate where the computational challenges rise. First, when the agent's problem has the form (5.10) with a continuous polytope domaix,  $(u_t) = X cp(A_t; c_t)$ , the subbased solution

oracle simpli es to the following single level model.

$$\begin{split} q_{t+1} &= \underset{q;x;w,v;y;z}{argmin} \quad \frac{1}{2} kq \quad q_t k^2 + h_t (hq;x \quad yi \quad \frac{1}{2} x^T P x) \\ &\quad s.t. A_t x \quad c_t; \ x 2 \ R_+^n \\ &\quad w_i \quad M y_i; \ i 2 \ [n] \\ &\quad x_i \quad M(1 \quad y_i); \ i 2 \ [n] \\ &\quad v_j \quad M z_j; \ j 2 \ [m] \\ &\quad (A_t)_j^3 x \quad (c_t)_j \quad M(1 \quad z_j); \ j 2 \ [m] \\ &\quad P x \quad q + A_t^s v \quad w = 0 \\ &\quad v 2 \ R_+^m; \ w 2 \ R_+^n; \ y 2 \ f \ 0; \ 1g^n; \ z 2 \ f \ 0; \ 1g^m \\ &\quad q 2 \ Q: \end{split}$$

For the other case, when the agent's problem has the (forth ) with an equally constrained knapsack domain, X  $(u_t) = X e^{ck}(p_t; b_t)$ , the `sub based solution oracle is equivalent to:

$$\begin{split} q_{t+1} &= \underset{q;x;w;v;y}{\operatorname{argmin}} \quad \frac{1}{2} kq \quad q_t k^2 + h_t(hq;x \quad yi \quad \frac{1}{2} x^T \mathsf{P} x) \\ &\text{s.t. } p_t x = b_t; \; x \; 2 \; \mathsf{R}^n_+ \\ &w_i \quad My_i; \; i \; 2 \; [n] \\ &x_i \quad M(1 \quad y_i); \; i \; 2 \; [n] \\ &2q_i x_i + v(p_t)_i \quad w_i = \; 0; \; i \; 2 \; [n] \\ &v \; 2 \; \mathsf{R}; w \; 2 \; \mathsf{R}^n_+; \; y \; 2 \; f \; 0; \; 1g^n \\ &q \; 2 \; Q; \end{split}$$

We note that the reformulated objective functions contain the bilinear hermin, which means the objectives are not guaranteed to be convex in general. As a result, where a sed implicit OL algorithm would require an expensive general purpose nonconvex solution oracle.

# B.3 Formulations for the Solution Oracles Used in the Implicit OL Algorithms

In this section, we give the solution oracles used in implicit OL algorithms based of orand pre for two forms of agent's utility functions corresponding to the ones used in our numerical experiments. We include pre-based implicit OL in our discussion for the sake of comparison between based OL framework and the previous work Dong et al. (2018a). In our computational study, we implemented the following solution oracles when they can be readily solved by standard optimization software.

#### B.3.1 Solution Oracle for `sim-based Implicit OL Algorithm

Suppose that the squared Euclidean norm is used as the distance generating function in the implicit OL algorithm with the solution oracle. Recall from De nition 1 that has the following form under Assumption 2:

$$\sum^{sim}(q; x(q_t; u_t); y_t; u_t) \coloneqq hq; c(y_t) \quad c(x(q_t; u_t))i + hq_{true}; c(x(q_t; u_t)) \quad c(y_t)i:$$

Since the constant term  $i\hbar^{im}(q)$  has no impact when  $\hbar^{sim}(q)$  is used in the objective function of an optimization problem, it can be ignored in the solution oracle formulation. Then, we deduce that the solution oracle for the sim-based implicit OL algorithm updates 1 as

$$q_{t+1} = \underset{q2Q}{\operatorname{argmin}} \frac{1}{2} kq \quad q_t k^2 + h_t hq; c(y_t) \quad c(x(q_t; u_t))i:$$

In particular, when the agent's problem has the form 0) we have  $f(x;q;u) = \frac{1}{2}x^{>}Px + q;xi$ , i.e., c(x) = x. Thus, in this case, the solution oracle for the based implicit OL algorithm updates  $q_{t+1}$  as

$$q_{t+1} = \underset{q2Q}{\operatorname{argmin}} \frac{1}{2} kq \quad q_t k^2 + h_t hq; \quad y_t + x(q_t; u_t)i:$$

In the case of CES utility function, i.e., when the agent's problem has the (formal), we have  $f(x;q;u) = a_{i2[n]}(q)_i x_i^2$ , and in this case the solution oracle for the based implicit OL algorithm updates  $q_{t+1}$  as

$$q_{t+1} = \underset{q \ge Q}{\operatorname{arg\,min}} \frac{1}{2} kq \quad q_t k^2 + h_t \overset{o}{a}_{i \ge [n]} q_i (y_t)_i^2 \quad x(q_t; u_t)_i^2 :$$

### B.3.2 Solution Oracle for `pre-based Implicit OL Algorithm

Suppose that the squared Euclidean norm is used as the distance generating function in the implicit OL algorithm with the solution oracle. Then, the solution oracle for the based implicit OL algorithm updates  $a_{t+1}$  by solving the following bilevel program:

$$q_{t+1} = \underset{q2Q}{\operatorname{argmin}} \frac{1}{2} kq \quad q_t k^2 + h_t ky_t \quad x(q; u_t) k^2;$$

where

$$x(q; u_t) 2 \underset{x}{\operatorname{argminf}} f(x; q; u_t) : g(x; u_t) = 0; x 2 X g$$

Recall that when the agent's problem has the f( $\mathfrak{m}10$ ) with a continuous polytope domain, i.e., X (u<sub>t</sub>) = X <sup>cp</sup>(A<sub>t</sub>;c<sub>t</sub>), we have

$$x(q;u_t) \coloneqq \underset{x}{\text{argmax}} \quad \frac{1}{2}x^{>}Px + hq; xi : A_tx \quad c_t; x 2 R_+^n \quad ;$$

where P 2  $S_{++}^n$  is a xed positive de nite matrix known by both the learner and the agent. Using the KKT optimality conditions for the inner problem, and then introducing binary variables to linearize the resulting nonlinear relations, it is possible to reformulate this bilevel problem into a single level optimization problem with binary variables. In particular, in this case, following these outlined steps, Dong et al. (2018a) proposed the following reformulation of this bilevel problem into a single level MISOCP:

$$\begin{split} q_{t+1} &= \mathop{argmin}_{q;x;w,v;y;z} \ \ \frac{1}{2} \, kq \quad q_t k^2 + \, h_t \, ky_t \quad xk^2 \\ & \text{s.t. } A_t x \quad c_t; \, x 2 \, R^n_+ \\ & w_i \quad My_i; \, i 2 \, [n] \\ & x_i \quad M(1 \quad y_i); \, i 2 \, [n] \\ & v_j \quad Mz_j; \, \, j 2 \, [m] \\ & (A_t)_j^> x \quad (c_t)_j \quad M(1 \quad z_j); \, \, j 2 \, [m] \\ & Px \quad q + \, A_t^> v \quad w = 0 \\ & v 2 \, R^m_+; w 2 \, R^n_+; \, y 2 \, f \, 0; \, 1g^n; \, z 2 \, f \, 0; \, 1g^m \\ & q \, 2 \, Q: \end{split}$$

Here,M is the so-called bigM constant. The variables2  $R_{+}^{m}$ ; w 2  $R_{+}^{n}$  are the variables corresponding to the Lagrangian multipliers, the binary variables2 f 0; 1g for all i 2 [n] are used to linearize the KKT condition w<sub>i</sub>x<sub>i</sub> = 0, andz<sub>j</sub> 2 f 0; 1g for all j 2 [m] are introduced to linearize the KKT relation  $v_{j}((A_{t})_{j}^{3} x (c_{t})_{j}) = 0$ . Therefore, the bigM constants must be selected so that they upper bound the components in the bilinear expressions,  $e_{i}$ gandw<sub>i</sub> for the complementarity constraim(x<sub>i</sub> = 0 as well as(A<sub>t</sub>)<sub>j</sub><sup>3</sup> x (c<sub>t</sub>)<sub>j</sub> andv<sub>j</sub> for the constraintr<sub>j</sub>((A<sub>t</sub>)<sub>j</sub><sup>3</sup> x (c<sub>t</sub>)<sub>j</sub>) = 0. Because in our instances the agent's domain for is bounded, we can easily obtain boundsxpand(A<sub>t</sub>)<sub>j</sub><sup>3</sup> x (c<sub>t</sub>)<sub>j</sub> terms. It is also possible to derive an upper bound for the Lagrange multipliers under a Slater condition assumption on the primal problem. Nevertheless, it is well known that usin bl/b/g/mulations signi cantly degrade the optimization solver performance, and instead it is encouraged in Gurobi solver that such bigM constraints are encoded as indicator constraints, which is a form of logical constraint feature of the Gurobi solver. Note that this alternative implementation is possible because the bigM constraints essentially represent a complementarity type logical condition.

Note that the continuous knapsack dom  $\operatorname{kin}^{\operatorname{k}}(p_t; b_t)$  is a special case of the continuous polytope domain X  $\operatorname{cp}(A_t; c_t)$ , and thus the same reformulation also holds in that case.

Finally note that when the agent's problem has the form 1) with an equally constrained knapsack domain, i.eX,  $(u_t) = X e^{ck}(p_t; b_t)$ , we have

$$( )$$

$$x(q;u_t) \coloneqq \underset{x}{\operatorname{argmin}} \quad \underset{i2[n]}{\overset{a}{a}} q_i x_i^2 \colon p_t^{>} x = b_t; \ x \ge R_+^n \quad :$$

In this case, the bilevel program corresponding to the solution oracle in the ased implicit OL algorithm has the following single level reformulation.

$$\begin{split} q_{t+1} &= \underset{q;x;w,v,y}{argmin} \quad \frac{1}{2} kq \quad q_t k^2 + h_t \, ky_t \quad xk^2 \\ &\quad \text{s.t. } p_t x = b_t; \; x 2 \; R^n_+ \\ &\quad w_i \quad My_i; \; i \; 2 \; [n] \\ &\quad x_i \quad M(1 \quad y_i); \; i \; 2 \; [n] \\ &\quad 2q_i x_i + v(p_t)_i \quad w_i = 0; \; i \; 2 \; [n] \\ &\quad v \; 2 \; R; w \; 2 \; R^n_+; \; y \; 2 \; f \; 0; \; 1g^n \\ &\quad q \; 2 \; Q; \end{split}$$

Unfortunately, this nonconvex mixed integer program contains the bilinear top myshere both x andq are continuous variables, in a general constraint, not of a complementarity type constraint. Note that the primal domain is equality constrained continuous knapsack, and thus we can ind an upper bound on variables. Moreover, for 2 Q and wherQ is bounded like the Euclidean ball or the simplex case that we focus on in this paper, we can ind a boundary well. However, because this bilinear term of xi is appearing in a general constraint and not in a complementary constraint, there is no technique to reformulate this nonconvexity as linear constraints by introducing new binary variables. Hence, in this case the based implicit OL algorithm requires a computationally expensive general purpose nonconvex solution oracle.

# B.4 `<sup>sim</sup>-based Online Inverse Optimization under Imperfect Information

In this appendix, we look into the learning performance in the imperfect information setup under Assumption 2. Recall that for any of the four loss de nitions, the loss  $va_t(up)$  can be split into a component measuring the difference betwe(ep,ut) and  $x(q_{true}; u_t)$ , and another one re ecting the noise shifting  $x(q_{true}; u_t)$  to  $y_t$ . For instance, we can write  $p^{re}(q_t) = ky_t x_t + x_t x_t k^2 = ky_t x_t k^2 + kx_t x_t k^2 + 2hy_t x_t; x_t x_t i$ . Moreover, this mean  $a_t(q_{true}) = (q_{true}; x(q_{true}; u_t); y_t; u_t)$  can be viewed as an imperfect information loss.

In Section 5.6.4, we consider a different collection of performance measures, consisting of the average regret with respect  $\dot{w}^m$ , the difference between the average loss incurred fround  $(q_t; u_t)g$ 

and the average imperfect information loss frbq<sub>true</sub>;  $x(q_{true}; u_t)g$  with respect to <sup>sub</sup> and <sup>est</sup>, and the average squared norm distance betwfeqt;  $u_t)g$  and f  $x(q_{true}; u_t)g$ . We next show that the average regret with respect to <sup>sub</sup> can be used to bound the other three loss-based measures.

For notation ease, we denote  $x(q_t; u_t)$  and  $x_t := x(q_{true}; u_t)$  for all t.

Proposition 6. Suppose Assumption 2 holds and the observations contain noise, namely, in each time step t,  $y = x(q_{true}; u_t) + e_t$  for some nonzere. For any sequence  $q_t g_{t2[T]}$ , we have

(a) 
$$\overset{a}{a} \overset{sub}{_{t^{2}[T]}} \overset{sub}{_{t^{2}[T]}} (q_{t}) \overset{a}{_{t^{2}[T]}} \overset{sub}{_{t^{2}[T]}} (q_{true}) = R_{T}(f \overset{sim}{_{t}} g_{t^{2}[T]}; f q_{t} g_{t^{2}[T]});$$
  
(b)  $\overset{a}{a} \overset{est}{_{t^{2}[T]}} \overset{est}{_{t^{2}[T]}} (q_{true}) = \overset{a}{_{t^{2}[T]}} f (x_{t}; q_{true}; u_{t}) = f (x_{t}; q_{true}; u_{t})$   
 $R_{T}(f \overset{sim}{_{t}} g_{t^{2}[T]}; f q_{t} g_{t^{2}[T]}) + \overset{a}{_{t^{2}T}} hq_{true} = q_{t}; c(y_{t}) = c(x_{t})i;$ 

Proof. (a) By de nition of `sub, we have

$$\hat{a}_{t2[T]}^{sub}(q_{t}) \quad \hat{a}_{t2[T]}^{sub}(q_{true}) = \hat{a}_{t2[T]}^{sub}(q_{true}) \quad f(y_{t};q_{t};u_{t}) \quad f(y_{t};q_{true};u_{t}) + f(x_{t};q_{true};u_{t})$$

$$= \hat{a}_{t2[T]}^{sub}(q_{t};c(y_{t}) \quad c(x_{t})i \quad f_{1}(x_{t}) \quad h q_{true};c(y_{t}) \quad c(x_{t})i + f_{1}(x_{t})$$

$$= \hat{a}_{t2[T]}^{sub}(q_{t}) \quad q_{true};c(y_{t}) \quad c(x_{t})i + \hat{a}_{t2[T]}^{sub}(q_{true};u_{t}) \quad f(x_{t};q_{true};u_{t})$$

$$= \hat{a}_{t2[T]}^{sub}(q_{t}) \quad \min_{q^{2}Q} \hat{a}_{t2[T]}^{sim}(q) + \min_{q^{2}Q} \hat{a}_{t2[T]}^{sim}(q)$$

$$= \hat{a}_{t2[T]}^{sim}(q_{t}) \quad q_{t2}(q_{t2}) + \hat{a}_{t}^{sim}(q_{true}) = R_{T}(f_{t}^{sim}g_{t2[T]};fq_{t}g_{t2[T]});$$

where the initial three equations follow from de nition and simple algebra, the rst inequality follows from the fact that minimizes f given  $q_{true}$ ;  $u_t$ , and the last equation is due to  $\int_{t}^{sim}(q_{true}) = 0$  for all t.

(b) The equation is a direct consequence of the loss de nition.

$$\overset{\circ}{a} \overset{\circ}{t} \overset{est}{t}(q_t) \qquad \overset{\circ}{a} \overset{\circ}{t} \overset{est}{t}(q_{true}) = \overset{\circ}{a} f(x_t; q_{true}; u_t) \qquad f(y_t; q_{true}; u_t) \qquad f(x_t; q_{true}; u_t) + f(y_t; q_{true}; u_t)$$
$$= \overset{\circ}{a} f(x_t; q_{true}; u_t) \qquad f(x_t; q_{true}; u_t):$$

To derive the inequality, we observe that Lemma 5 still holds under imperfect information, hence we can apply Lemma 5 (b) and (c) to rewrite the loss difference.

$$\hat{a}_{t2[T]}^{i} \hat{e}_{t2[T]}^{i} (q_{t}) + \hat{a}_{t2[T]}^{i} (q_{true}) = \hat{a}_{t2[T]}^{i} \hat{e}_{t2[T]}^{i} (q_{t}) + \hat{a}_{t2[T]}^{i} \hat{e}_{t2[T]}^{i} (q_{true})$$

$$= \hat{a}_{t2[T]}^{i} \hat{e}_{t2[T]}^{i} (q_{t}) + \hat{a}_{t2[T]}^{i} (q_{true}) + \hat{a}_{t2[T]}^{i} (q_{t}) + \hat{a}_{t2[T]}^{i} ($$

where the equations follow from de nition and algebra, then the inequality uses the fact that minimizes f given  $q_t$ ;  $u_t$ .

Under imperfect information, (Mohajerin Esfahani et al., 2018, Proposition 2.5) remains valid for a strongly convext, namely,  $t^{sub}(q) = \frac{g}{2} t^{pre}(q)$  for all t and for all q 2 Q, with g being the strong convexity parameter off. Therefore, we can further derive bounds with respect  $t^{pre}_{t}g_{t_2[T]}$ .

Corollary 2. Suppose Assumption 2 holds and the observation  $g_{2[T]}$  contain noises. Suppose further that f is a strongly convex function affor every q with a constant y > 0, i.e.,  $f(x;q;u) f(y;q;u) h s_y; x yi + \frac{g}{2}kx yk^2$ , where  $s_y$  is a subgradient of (y;q;u) with respect toy. Then, for any sequence  $f_{t}g_{t2[T]} = Q$  we have:

!

(a) 
$$\frac{g}{2} \overset{a}{}_{t2[T]}^{t} \overset{pre}{}_{t2[T]}^{pre}(q_t) \overset{a}{}_{t2[T]}^{t} \overset{pre}{}_{t2[T]}^{pre}(q_{true}) = R_T(f_t^{sim}g_{t2[T]}; fq_tg_{t2[T]}) = g \overset{a}{a}_{t2[T]}^{s} kx_t y_t k^2 \overset{a}{}_{t2[T]}^{s} hs_{y_t}; x_t y_t i;$$
  
(b)  $\frac{g}{2} \overset{a}{}_{t2[T]}^{s} kx_t x_t k^2 = R_T(f_t^{sim}g_{t2[T]}; fq_tg_{t2[T]}) + \overset{a}{a}_{t2T}^{s} hq_{true} q_t; c(y_t) c(x_t) i:$ 

Proof. The strong convexity of has a few implications that we will use to derive the desirable statements.

$$f(y_t;q_t;u_t) = f(x_t;q_t;u_t) + s_{x_t};y_t = x_t i + \frac{g}{2}kx_t - y_t k^2 - \frac{g}{2}kx_t - y_t k^2$$
(B.2)

$$f(x_t; q_{true}; u_t) = f(y_t; q_{true}; u_t) + s_{y_t}; x_t = y_t i + \frac{g}{2} k x_t = y_t k^2$$
 (B.3)

$$f(x_t; q_{true}; u_t) \quad f(x_t; q_{true}; u_t) \quad h \in S_{x_t}; x_t \quad x_t i + \frac{g}{2} k x_t \quad x_t k^2 \quad \frac{g}{2} k x_t \quad x_t k^2:$$
(B.4)

(a) We apply (B.2) and (B.3) to bound the loss gap.

$$\begin{array}{l} \frac{g}{2} \quad \overset{a}{a} \stackrel{\text{`pre}}{t^{2}(T)} (q_{t}) \quad \overset{a}{a} \stackrel{\text{`pre}}{t^{2}(T)} (q_{true}) = \frac{g}{2} \overset{a}{a} kx_{t} \quad y_{t}k^{2} \quad \frac{g}{2} \overset{a}{a} kx_{t} \quad y_{t}k^{2} \\ \overset{a}{t^{2}(T)} f(y_{t};q_{t};u_{t}) \quad f(x_{t};q_{t};u_{t}) \quad \frac{g}{2} \overset{a}{a} kx_{t} \quad y_{t}k^{2} = \overset{a}{a} \stackrel{\text{`sub}}{t^{2}(T)} (q_{t}) \quad \frac{g}{2} \overset{a}{a} kx_{t} \quad y_{t}k^{2} \\ R_{T}(f \stackrel{\text{`sim}}{t} g_{t2[T]}; f q_{t} g_{t2[T]}) + \overset{a}{a} \stackrel{\text{`sub}}{t^{2}(T)} (q_{true}) \quad \frac{g}{2} \overset{a}{a} kx_{t} \quad y_{t}k^{2} \\ = R_{T}(f \stackrel{\text{`sim}}{t} g_{t2[T]}; f q_{t} g_{t2[T]}) + \overset{a}{a} (f(y_{t};q_{true};u_{t}) \quad f(x_{t};q_{true};u_{t})) \quad \frac{g}{2} \overset{a}{a} kx_{t} \quad y_{t}k^{2} \\ R_{T}(f \stackrel{\text{`sim}}{t} g_{t2[T]}; f q_{t} g_{t2[T]}) + \overset{a}{a} (f(y_{t};q_{true};u_{t}) \quad f(x_{t};q_{true};u_{t})) \quad \frac{g}{2} \overset{a}{a} kx_{t} \quad y_{t}k^{2} \\ R_{T}(f \stackrel{\text{`sim}}{t} g_{t2[T]}; f q_{t} g_{t2[T]}) \quad g \overset{a}{a} kx_{t} \quad y_{t}k^{2} h s_{y_{t}}; x_{t} \quad y_{t}i; \end{array}$$

where the rst inequality use(B.2), the second inequality follows from Proposition 6(b), and the last inequality uses (B.3).

(b) We rst use (B.4) to bound the squared distance, then apply Proposition 6 (b).

$$\frac{g}{2} \underset{t2[T]}{\overset{a}{a}} kx_t x_t k^2 \underset{t2T}{\overset{a}{b}} f(x_t; q_{true}; u_t) f(x_t; q_{true}; u_t)$$

$$R_T(f \underset{t}{\overset{sim}{}} g_{t2[T]}; f q_t g_{t2[T]}) + \underset{t2T}{\overset{a}{b}} hq_{true} q_t; c(y_t) c(x_t) i:$$

For `est and `pre, we have weaker results. Proposition 6(b) shows that the average loss difference between  $q_t g$  and  $q_{true}$  decreases overtime to be below a noise-dependent  $rac{1}{2}$   $q_t ; c(y_t) = c(x_t)i$ . Lastly, in the special case of a strongly convex forward objective orollary 2(a) bounds the prediction loss gap with  $(f \cdot_t^{sim}g_{t2[T]}; f q_t g_{t2[T]})$  and additional noise-dependent terms. Corollary 2(b) proves that the same bound on the average loss difference based applies to the average squared norm distance between the predicted actions and the true actions.

Note that the experiment results from learning a quadratic utility function given in Section 5.6.4 illustrate these theoretical bounding relations. Moreover, as we next show, the results from learning a CES utility function demonstrate similar patterns.
