Fairness Modeling

J. N. Hooker

Keywords Fairness, Fairness modeling, Equity modeling, Welfare economics, Social choice theory, Inequality-averse optimization, Distributive justice, Proportional fairness, Nash bargaining solution, Alpha fairness, Kalai-Smorodinsky bargaining solution

1 Introduction

There is growing interest in incorporating fairness or equity into the solution of optimization models. Fairness concerns can arise when costs or benefits are distributed to multiple stakeholders, and particularly when a traditional objective of minimizing total cost or maximizing total benefit results in an inequitable distribution.

For example, if donated organs are allocated in the most economically efficient fashion, some patients may wait far longer for a transplant than others [21]. If earthquake shelters are located so as to minimize average distance from residents, persons living in less densely populated areas may have much further to travel [23]. Electric grid development that maximizes total services delivered may focus on urban customers while neglecting small towns and rural areas [31]. If traffic signals at intersections are timed to maximize traffic throughput, motorists on side streets may have to wait forever for a green light [3].

Scenarios such as these pose the problem of finding a mathematical formulation of fairness or equity that is suitable for the problem at hand and

J.N. Hooker
Tepper School of Business, Carnegie Mellon University, Pittsburgh, USA

can be incorporated into an optimization model. This article surveys a variety of fairness criteria that have been proposed, classified under three broad categories: inequality metrics, fairness for the disadvantaged, and criteria that balance efficiency and fairness.

To standardize terminology, suppose that the solution of an optimization model allocates *utilities* $u_1, \ldots, u_n$ to a collection of n *stakeholders*, and there is interest in the fairness of this allocation. Utility can take the form of wealth, resources, negative cost, health outcomes, or some other type of benefit. Stakeholders can be individuals, organizations, demographic groups, geographic regions, or other entities for which distributive justice is a concern.

2 Inequality Metrics

Inequality metrics reflect the degree to which utility is distributed equally among the stakeholders. Since they are insensitive to the actual level of welfare the stakeholders enjoy, inequality metrics are generally used to place an upper bound on the degree of inequality that results from a conventional utility-maximizing objective function.

The *relative mean deviation*, one of the simplest inequality metrics, is the normalized sum of deviations from the average utility $\bar{u}$. It is given by

$$\frac{1}{n\bar{u}} \sum_i |u_i - \bar{u}| \tag{1}$$

An upper bound UB can be placed on (1) by introducing variables $v_1, \ldots, v_n$ and adding the constraints

$$\begin{aligned} &\frac{1}{n} \sum_i v_i \leq \text{UB}\bar{u} \\ &-v_i \leq u_i - \bar{u} \leq v_i, \ \text{all } i \end{aligned}$$

The *coefficient of variation* is the normalized standard deviation and can be written

$$c_v = \frac{1}{\bar{u}} \Big[\frac{1}{n} \sum_i (u_i - \bar{u})^2 \Big]^{\frac{1}{2}}$$

A related measure is *Jain's index* $(1 + c_v^2)^{-1}$, which is widely used in telecommunications [15]. Larger values of the index indicate greater equality, with 1 corresponding to perfect equality. Since Jain's index is a strictly decreasing function of c_v, bounding it below is equivalent to bounding c_v above. Other metrics developed specifically for networks and telecommunications include QoE fairness [10, 14], TCP fairness [27], G's fairness index, and Bossaert's fairness index [22].

The popular *Gini coefficient* is based on the Lorentz curve, which in this context is the cumulative utility distribution. The Gini coefficient is proportional to the area between the curve and the diagonal line connecting its endpoints, so that a coefficient of zero indicates perfect equality. The coefficient is given by

$$\frac{1}{2\bar{u}n^2} \sum_{i,j} |u_i - u_j| \tag{2}$$

An upper bound UB on (2) can be imposed by the constraints

$$\begin{aligned} &\frac{1}{2n^2} \sum_{i,j} v_{ij} \leq \text{UB}\bar{u} \\ &-v_i \leq u_i - u_j \leq v_{ij}, \text{ all } i, j \end{aligned}$$

The *Hoover index* is also related to the Lorenz curve, as it is proportional to the maximum vertical distance between the curve and the diagonal line representing perfect equality [13]. It can be interpreted as the fraction of total utility that would have to be redistributed to achieve perfect equality. The index is given by

$$\frac{n}{2\bar{u}} \sum_{i} |u_i - \bar{u}|$$

and is therefore proportional to the relative mean deviation.

A general discussion of inequality metrics can be found in [6, 16]. A shortcoming of these metrics is that they provide no guidance as to how much utility should be sacrificed to achieve an acceptable degree of equality. The models presented in subsequent sections below are sensitive to both efficiency and fairness and therefore take a position on how they should be balanced.

3 Fairness for the Disadvantaged

A second category of models strive for equality by improving the lot of stakeholders at the lower end of the distribution. Because this tends to push up the average welfare level, these models can also take some account of overall efficiency.

The *maximin* criterion, based on the famous difference principle of John Rawls, maximizes the welfare of the worst-off individual or social class [28]. It is defended with a social contract argument that has been intensely discussed in the philosophical literature (as surveyed in [8, 29]). 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 is frequently used as a general criterion for distributing utility. It can be adopted as an objective function by maximizing $\min_i\{u_i\}$, which is easily linearized by maximizing the scalar variable v subject to $v \leq u_i$ for all i.

A problematic feature of the maximin criterion is that, in some cases, it can ignore the welfare of stakeholders other than the very worst off. For example, $(u_1, u_2) = (1, 1)$ is an optimal solution of a problem with two stakeholders and constraints $u_1 \leq 1$, $u_2 \leq 10$, even though stakeholder 2 could receive 10 times as much utility in another optimal solution.

One approach to this difficulty is to use *β-fairness*, which maximizes the average utility of the least advantaged $100\beta\%$ of the population [25, 7]. Thus $\beta = 0$ corresponds to a maximin objective and $\beta = 1$ to a purely utilitarian objective. This criterion arguably reflects the spirit of the Rawlsian difference principle, since the lowest $100\beta\%$ of the population could be regarded as the lowest social class. Mathematically, β-fairness is achieved by maximizing

$$\min_{\boldsymbol{\delta}} \left\{ \frac{1}{\beta n} \sum_i u_i \delta_i \;\middle|\; \sum_i \delta_i = \beta n;\;\; \delta_i \in \{0,1\},\ \text{all } i \right\} \tag{3}$$

over feasible distributions $\boldsymbol{u}$. Since the coefficient matrix of the linear programming (LP) relaxation of (3) is totally unimodular, one can solve (3) by solving the dual of the LP relaxation:

$$\max_{\lambda, \boldsymbol{\mu}} \left\{ \beta n \lambda + \sum_i \mu_i \;\middle|\; \lambda + \mu_i \leq \frac{u_i}{\beta n},\ \mu_i \geq 0,\ \text{all } i \right\} \tag{4}$$

Thus one can compute a β-fair solution by solving (4) over all feasible $\boldsymbol{u}$ as well as over λ and $\boldsymbol{\mu}$, which is again a linear programming problem if $\boldsymbol{u}$ is linearly constrained..

While β-fairness considers stakeholders other than the very worst-off, there is no fairness consideration *within* the lowest social class, which is treated in a purely utilitarian manner. This can be addressed by using, instead, a *leximax* or lexicographic maximum fairness criterion, which (roughly speaking) maximizes the lowest utility, then the second lowest, and so forth. More precisely, if $u_{\langle 1\rangle}, \ldots, u_{\langle n\rangle}$ is $u_1, \ldots, u_n$ in nondecreasing order, then $\boldsymbol{u}^*$ is a lexicographic maximum if for any feasible $\boldsymbol{u}$, there is a $k \in \{1, \ldots, n\}$ such that $u^*_{\langle i\rangle} = u_{\langle i\rangle}$ for $i = 1, \ldots, k-1$ and $u^*_{\langle k\rangle} \geq u_{\langle k\rangle}$.

The *McLoone index* compares the total utility of stakeholders 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 otherwise takes a smaller positive value. The McLoone index therefore benefits the disadvantaged by penalizing inequality in the lower half of the distribution, but it is unconcerned by the existence of highly privileged stakeholders. It is frequently used in public sector contexts, particularly education [33]. The index is given by

$$\frac{1}{|I(\boldsymbol{u})|\tilde{u}} \sum_{i \in I(\boldsymbol{u})} u_i$$

where $\tilde{u}$ is the median of utilities in $\boldsymbol{u}$ and $I(\boldsymbol{u})$ is the set of indices of utilities at or below the median, so that $I(\boldsymbol{u}) = \{i \mid u_i \leq \tilde{u}\}$. A mixed integer/linear programming (MILP) model for this criterion is given in [5].

4 Balancing Fairness and Efficiency

Practical applications nearly always require a trade-off between efficiency and fairness. This section presents some objective functions that are designed to strike a reasonable balance between them. These functions are well suited to replace the traditional utility-maximizing objective function in an optimization model, since they incorporate total utility as well as fairness.

A simple way to manage the fairness/efficiency trade-off is to use a convex combination of total utility and some fairness measure, such as the Gini coefficient. However, it is unclear how to select an appropriate multiplier to weight the two objectives, or even understand what the multiplier means, particularly when the objectives are expressed in different units. The Gini coefficient, for example, is dimensionless, while utility is not. The remainder of this section therefore focuses on more sophisticated criteria for balancing fairness and efficiency.

Perhaps the most frequently used of these criteria, especially in engineering fields, is *alpha fairness*. The objective function is

$$\begin{cases} \dfrac{1}{1-\alpha} \displaystyle\sum_i u_i^{1-\alpha}, \text{ for } \alpha \geq 0,\ \alpha \neq 1 \\ \displaystyle\sum_i \log(u_i), \qquad \text{for } \alpha = 1 \end{cases} \tag{5}$$

The parameter α indicates how much emphasis is placed on fairness, with larger values corresponding to greater emphasis. In particular, $\alpha = 0$ corresponds to a purely utilitarian objective, and $\alpha = \infty$ to a maximin objective. An interesting special case is $\alpha = 1$, which yields *proportional fairness*, also known as the *Nash bargaining solution*. A number of axiomatic derivations and rational bargaining arguments have been provided for proportional fairness [24, 11, 30, 1], and an axiomatic derivation exists for general alpha fairness as well [19, 20].

The alpha fairness objective (5) is nonlinear, but its concavity makes it easier to maximize. In addition, the parameter α is not without practical interpretation, partly because the special case $\alpha = 1$ provides a meaningful benchmark. Suppose, for example, that (5) is maximized subject to a budget constraint $\sum_i c_i u_i \leq B$, where B is the total budget. The coefficient c_i is the unit cost of creating utility for stakeholder i, and $1/c_i$ is the stakeholder's *conversion efficiency*. Then proportional fairness ($\alpha = 1$) allots utility to stakeholders in proportion to their conversion efficiency (whence the name), with a greater personal reward for efficiency when $\alpha < 1$ and a smaller one when $\alpha > 1$. Also, if a stakeholder increases its conversion efficiency and the problem is re-optimized, the stakeholder takes benefits from others if $\alpha < 1$ and shares the surplus with others if $\alpha > 1$, thus incentivizing competition in one case and cooperation in the other.

One anomaly of alpha fairness is that in certain cases, a stakeholder's utility allocation can shrink when the set of feasible utility solutions $\boldsymbol{u}$ is enlarged. This is addressed by the *Kalai-Smorodinksy (K–S) bargaining solution*, which is also supported by axiomatic and bargaining arguments [17]. It is defined in terms of each stakeholder's relative utility gain γ, which is the ratio of the allocated utility to the utility the stakeholder would enjoy if given access to all available resources. The K–S solution maximizes the relative utility gain on the condition that it is the same for all stakeholders. Mathematically, it maximizes γ over feasible distributions $\boldsymbol{u}$ subject to $\boldsymbol{u} = \gamma \boldsymbol{u}^{\max}$, where each $u_i^{\max}$ is the maximum of u'_i over all feasible utility distributions $\boldsymbol{u}'$. The K–S solution may be suitable for bargaining situations in which the parties view equal relative concessions as fair, as when a buyer and seller negotiate a price, or labor and management negotiate wages [32]. It is also arguably consistent with a "contractarian" ethical philosophy in the Hobbesian tradition [9].

A more recently proposed approach to balancing fairess and efficiency is provided by *threshold criteria*, which are of two types [34]. One, based on an *efficiency threshold*, imposes a maximin objective until the efficiency cost becomes unacceptably great, at which point some stakeholders are switched to a utilitarian criterion. The other, based on an *equity threshold*, imposes a utilitarian criterion until inequity becomes unacceptably great, at which point a maximin criterion is introduced.

Contours for a two-stakeholder efficiency-threshold criterion appear in Fig. 1. The central region of the figure depicts a maximin objective and the outer regions a utilitarian objective, while the parameter Δ determines the boundaries between the regions. The area under the curve represents the set of feasible utilities (u_1, u_2). The small open circle designates the maximin solution, which may be unsatisfactory because, in this part of the feasible set, a small improvement in stakeholder 1's utility demands a substantial sacrifice in total utility. The efficiency-threshold solution is shown by the black dot, which allows much greater total utility in exchange for a modest reduction in u_2.

This two-stakeholder criterion can be generalized to n stakeholders as follows [12]:

$$(n-1)\Delta + \sum_i \max\left\{u_i - \Delta, u_{\min}\right\}$$

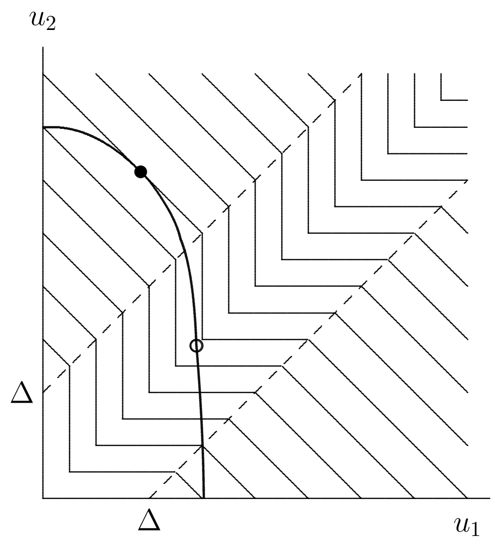

Figure 1 Contours for a two-stakeholder efficiency-threshold criterion.

where $u_{\min} = \min_i\{u_i\}$. This function gives priority to stakeholders whose utility is within $\varDelta$ of the lowest, so that larger values of $\varDelta$ correspond to a greater emphasis on fairness. In particular, $\varDelta = 0$ corresponds to a purely utilitarian criterion and $\varDelta = \infty$ to a maximin criterion. In practice, $\varDelta$ is chosen so that stakeholders within $\varDelta$ of the worst-off are viewed as deserving special priority, which means that $\varDelta$ has an interpretation that is not available for the α parameter in alpha fairness. The n-stakeholder criterion can be embedded in an optimization model by introducing continuous variables v_i and w, 0–1 variables δ_i, and maximizing $n\varDelta + \sum_i v_i$ subject to

$$
\begin{aligned}
& u_i - \varDelta \le v_i \le u_i - \varDelta\delta_i, \text{ all } i \\
& w \le v_i \le w + (M - \varDelta)\delta_i, \text{ all } i \\
& u_i - u_j \le M, \text{ all } i, j \\
& u_i \ge 0, \ \delta_i \in \{0, 1\}, \text{ all } i
\end{aligned}
$$

and subject to the remaining constraints in the model. The second line consists of "big-M" constraints (common in MILP models) that require the constant M to be some large number.

The efficiency threshold is appropriate for situations in which equity is the primary concern, such as task assignments to workers or provision of public services to neighborhoods. An equity threshold is more suitable when total utility is the primary goal, but not at the cost of excessive inequity, as could occur in telecommunications, traffic signal management, or perhaps disaster recovery. Contours of a two-stakeholder equity-threshold model appear in Fig. 2. The small open circle is the utilitarian solution, which may be unsatisfactory because it is much more favorable to stakeholder 2. The black dot indicates the equity-threshold solution, which continues to generate substantial utility while providing greater equity.

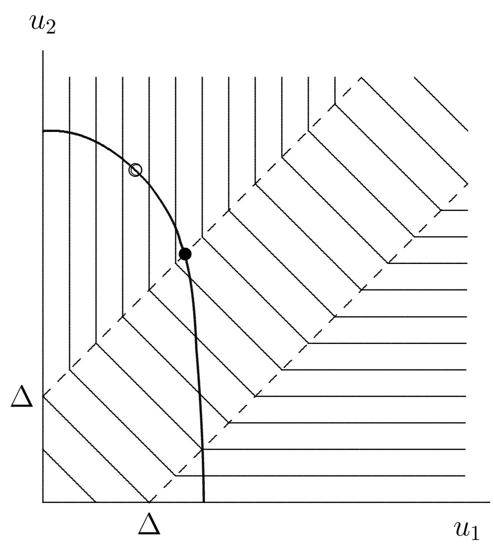

Figure 2 Contours for a two-stakeholder equity-threshold criterion.

An n-stakeholder extension of the equity-threshold criterion is

$$n\Delta + \sum_i \min\left\{u_i - \Delta, u_{\min}\right\}$$

Here a smaller Δ corresponds to a stronger emphasis on equity, with $\Delta = 0$ yielding a maximin solution. In applications, the value of Δ is chosen so that stakeholders with utility exceeding the lowest utility by more than Δ are regarded as privileged, and adding utility to privileged stakeholders is not seen as enhancing social welfare unless the lowest utility is increased by an equal amount. This criterion has a linear programming formulation that maximizes $n\Delta + \sum_i v_i$ subject to

$$\begin{aligned} &v_i \leq w \leq u_i, \;\; v_i \leq u_i - \Delta, \;\; \text{all } i \\ &w \geq 0 \end{aligned}$$

and subject to the remaining constraints in the optimization model.

The tendency of the maximin criterion to ignore the welfare of all but the very worst-off carries over to threshold criteria that blend maximin with a utilitarian objective. A possible remedy is to use a threshold formulation that combines *leximax* and utilitarian objectives, although this presents a significant modeling challenge. A utility-threshold formulation of this kind that requires solving a sequence of MILP models has been proposed [4], and a similar equity-threshold formulation could perhaps be developed.

5 Conclusion

Fairness modeling is a relatively recent research program in operations research that addresses social justice concerns. It may also forge new connections with other fields. Interactions with economics, statistics, and engineering have historically led to advances in optimization, and collaboration with ethicists and welfare economists on the precise formulation of fairness concepts may bring similar benefits to the art and science of modeling.

Comprehensive treatments of fairness modeling can be found in [18] and [5]. In addition, [26] reviews formulations developed for telecommunications and facility location, two major users of fairness models.

See also

Emergency Evacuation: Optimization Modeling
Fractional Programming
Multicriteria Decision Analysis for Corporate Responsibility and Sustainable Investments

References

[1] Binmore K, Rubinstein A, Wolinsky A (1986) The Nash bargaining solution in economic modeling. RAND Journal of Economics 17:176–188
[2] Charnes A, Cooper WW (1962) Programming with linear fractional functionals. Naval Research Logistics Quarterly 9:181–186
[3] Chen LW, Sharma P, Tseng YC (2013) Dynamic traffic control with fairness and throughput optimization using vehicular communications. IEEE Journal on Selected Areas in Communications 31:504–512
[4] Chen V, Hooker JN (2022) Combining leximax fairness and efficiency in an optimization model. European Journal of Operational Research 299:235–248
[5] Chen V, Hooker JN (2023) A guide to formulating fairness in optimization models. Annals of Operations Research 326:581–619

[6] Cowell FA (2000) Measurement of inequality. In: Atkinson AB, Bourguignon F (eds) Handbook of Income Distribution, vol 1, Elsevier, pp 89–166
[7] Filippi C, Guastaroba G, Speranza M (2021) On single-source capacitated facility location with cost and fairness objectives. European Journal of Operational Research 289:959–974
[8] Freeman S (ed) (2003) The Cambridge Companion to Rawls. Cambridge University Press
[9] Gauthier D (1983) Morals by Agreement. Oxford University Press
[10] Georgopoulos P, Elkhatib Y, Broadbent M, Mu M, Race N (2013) Towards network-wide QoE fairness using openflow-assisted adaptive video streaming. In: Proceedings of the 2013 ACM SIGCOMM Workshop on Future Human-centric Multimedia Networking, pp 15–20
[11] Harsanyi JC (1977) Rational Behavior and Bargaining Equilibrium in Games and Social Situations. Cambridge University Press
[12] Hooker JN, Williams HP (2012) Combining equity and utilitarianism in a mathematical programming model. Management Science 58:1682–1693
[13] Hoover EM (1936) The measurement of industrial localization. Review of Economics and Statistics 18:162–171
[14] Hoßfeld T, Skorin-Kapov L, Heegaard PE, Varela M (2018) A new QoE fairness index for QoE management. Quality and User Experience 3
[15] Jain R, Chiu DM, Hawe W (1984) A quantitative measure of fairness and discrimination for resource allocation in shared computer systems. Tech. Rep. TR–301, Eastern Research Laboratory, DEC, Hudson, MA
[16] Jenkins SP, Van Kerm P (2011) The measurement of economic inequality. In: Nolan B, Salverda W, Smeeding TM (eds) The Oxford Handbook of Economic Inequality, Oxford University Press
[17] Kalai E, Smorodinsky M (1975) Other solutions to Nash's bargaining problem. Econometrica 43:513–518
[18] Karsu O, Morton A (2015) Inequality averse optimization in operational research. European Journal of Operational Research 245:343–359
[19] Lan T, Chiang M (2011) An axiomatic theory of fairness in resource allocation. Tech. rep., Princeton University
[20] Lan T, Kao D, Chiang M, Sabharwal A (2010) An axiomatic theory of fairness in network resource allocation. In: Proceedings of the 29th Conference on Information communications (INFOCOM), pp 1343–1351
[21] McElfresh C, Dickerson J (2018) Balancing lexicographic fairness and a utilitarian objective with application to kidney exchange. 32nd AAAI Conference on Artificial Intelligence pp 1161–1168
[22] Mehta R (2020) Recursive quadratic programming for constrained nonlinear optimization of session throughput in multiple-flow network topologies. Engineering Reports 2:1–14
[23] Mostajabdaveh M, Gutjahr WJ, Sibel Salman F (2019) Inequity-averse shelter location for disaster preparedness. IISE Transactions 51(8):809–829

[24] Nash J (1950) The bargaining problem. Econometrica 18:155–162
[25] Ogryczak W, Zawadzki M (2002) Conditional median: A parametric solution concept for location problems. Annals of Operations Research 110:167–181
[26] Ogryczak W, Luss H, Pióro M, Nace D, Tomaszewski A (2014) Fair optimization and networks: A survey. Journal of Applied Mathematics 2014:1–25
[27] Pokhrel SR, Panda M, Vu HL, Mandjes M (2016) TCP performance over Wi–Fi: Joint impact of buffer and channel losses. IEEE Transactions on Mobile Computing 15:1279–1291
[28] Rawls J (1999) A Theory of Justice (revised). Harvard University Press (original edition 1971)
[29] Richardson HS, Weithman PJ (eds) (1999) The Philosophy of Rawls (5 volumes). Garland
[30] Rubinstein A (1982) Perfect equilibrium in a bargaining model. Econometrica 50:97–âĂŞ109
[31] Sackey CV, Nock D, Cao C, Armanios D, Davis A (2023) Incorporating elicited preferences for equality into electricity system planning modeling. Sustainability 15, published online
[32] Thompson W (1994) Cooperative models of bargaining. In: Aumann RJ, Hart S (eds) Handbook of Game Theory, vol 2, North-Holland, pp 1237–1284
[33] Verstegen DA (1996) Concepts and measures of fiscal inequality: A new approach and effects for five states. Journal of Education Finance 22:145–160
[34] Williams A, Cookson R (2000) Equity in health. In: Culyer A, Newhouse J (eds) Handbook of Health Economics, pp 1863–1910