Chapter

Facets of Fairness: Equal Share, Equal Chance and Uniform Mechanisms

Holger H. Hoos[1]

I. Introduction

Fairness is a notion of universal importance – whether we play, compete or trade with each other, we tend to strive for situations in which everyone feels treated fairly. Conversely, lack of fairness burdens human relations, producing tension, discontent and frustration. But what exactly *is* fairness? How can it be created, protected and enforced? What, apart from our innate perceptions, signals its presence or absence? Can we detect and measure it? This chapter attempts to approach this latter question from a technical perspective. While this perspective is necessarily limited in its scope, we believe that it can contribute meaningfully to a debate focussed on social and psychological aspects of a notion that is fundamental to the way we interact with each other, yet elusive to quantitative, analytical approaches.

Moreover, the specific angle we pursue here is of a computational nature. Our point of departure is the observation that notions of fairness play a role in computation; for example, in networking, parallel processing and compute-intensive optimisation. Computational tools can help explore and understand notions of fairness - for example, through simulation - and they can also help ensure fairness - for example, by producing better solutions to resource allocation problems or more compelling explanations of perceived lack of fairness.

The motivating question behind our investigation is this: Can fairness be detected or measured? Or, equivalently, yet perhaps easier to tackle: Can the *absence* of fairness be detected or measured? To facilitate exploration of this question, we will discuss three perhaps simplistic, but in our opinion informative, notions of fairness: equal share, equal chance and uniform mechanisms. While we do not contend that any of these, in isolation, captures the

[1] University of British Columbia, Department of Computer Science and Peter Wall Institute for Advanced Studies, hoos@cs.ubc.ca

concept of fairness to a satisfactory degree, we believe that they reflect general aspects of what fairness, or its absence, means in a broad range of situations, and hence represent what we consider interesting *facets of fairness*.

As we motivate, discuss and illustrate these facets of fairness, we highlight formal concepts related to them, including the Gini coefficient (an often used quantitative measure of inequality), the Pareto principle and Zipf's law (which characterise widely observed distributions that are markedly different from the ideal of equal share), as well as the amplification of unfairness and uncertainty through compounding effects. Since unfairness amplification represents a major challenge even for the most permissive of our notions of fairness, uniform mechanisms, we see considerable value in approaches for recognising unfairness amplification and for limiting its effects.

II. Equal Share

Perhaps the most simplistic notion of fairness is this: Every participant in an interactive endeavour gets an equal share of all relevant resources. Let us consider an example from the area of sports, in which two teams would like to use the same field for training sessions. The principle of equal share would clearly be violated if one team were allocated less time or less desirable times. We note that, while it is straightforward to quantify the amount of time allocated, the desirability of certain times over others depends entirely on the preferences of the two teams and their members: While for some, early Sunday morning may be an ideal time for football training, others might much prefer Saturday late afternoon. Considering such differences in preferences, the notion of equal allocation of desirable times can become quite difficult: Not only may it be challenging to elicit those preferences, particularly in a quantitative form, but there is also the potentially thorny issue of different preferences between the members of a team to be taken into account.[2] Furthermore, even in cases where the preferences are consistent between team members and known precisely, the

---

[2] See, for example, Vincent Conitzer (2006): "Computational Aspects of Preference Aggregation". Ph.D. Dissertation, Computer Science Department, Carnegie Mellon University, Pittsburgh (PA), USA.

problem of finding a time allocation that is optimal with respect to those preferences is known to be computationally very challenging.[3]

Another prominent and timely example is that of income distribution. In this context, the principle of equal share would seem to imply that, at the very least, two workers who perform comparable work with comparable efficiency and quality of outcome should receive the same compensation. Although at the first glance, even in a market-based economy that principle should hold, since if it did not, there would be an economic incentive to give work preferably to the worker who performs it at lower cost. Upon closer examination, however, the situation is more complex. One worker may, for example, receive higher compensation based on seniority, social connections, geographical location or negotiating skills. Is this unfair? Not necessarily, since in the first case, some income may simply be deferred to later career stages, as is the case in compensation schemes including regular increments over time; in the second case, the social connections could confer trust, which in turn may have monetary value, since it is negatively correlated with risk; in the third case, the cost of living may be much lower in some locations compared to others, and hence substantial differences in monetary income may be required to achieve comparable standards of living differences in negotiating skills may correlate with social skills that contribute to the value an employee or business partner brings to an organisation. Of course, it is easy to find examples for these and other factors contributing to income inequality that would be perceived as unjustified and unfair.

It is interesting to note that the study of income distributions has given rise to well-known quantitative measures of "inequality"; technically, these are measures of statistical variation over a population, where the term "population" may apply to a group of people or any other collection of entities or items. Perhaps the most widely used of these, in the context of income inequality, is the *Gini coefficient* (also known as *Gini index).* Developed by the Italian statistician and sociologist Corrado Gini, and published in 1914,[4] the Gini coefficient

---

[3] Technically, the corresponding scheduling problem is known to be NP-hard - see, for example, S. Even, A. Itai and A. Shamir (1976): "On the complexity of timetable and multicommodity flow problems". SIAM Journal on Computing 5 (4): 691–703.

[4] Corrado Gini (1914): "Di una misura della concentrazione indipendente dalla distribuzione del carattere". Atti del R. Istituto Veneto di Scienze, Lettere ed Arti, tomo LXXIII, parte II, pp. 1203-1258.

measures the degree to which the allocation of a resource, such as share of the total income of a group of households, individuals or other entities, deviates from a uniform distribution, which assigns equal shares to everyone. The technical definition involves the calculation of an integral over the so-called Lorenz curve, which relates cumulative shares of a resource to cumulative fractions of a population. Intuitively, a Gini coefficient of zero characterises a distribution in the form of perfectly equal shares, while a Gini coefficient of one corresponds to the maximally unequal distribution where a single member of a very large population is assigned all of the resource in question.[5] Sometimes, the values of the Gini coefficient are measured on a scale of 0-100, where 0 corresponds to perfect equality and 100 to perfect inequality of shares.

The Gini coefficient is prominently used to assess inequality of income distributions over households within a nation, and to compare income inequality between nations. According to data collected by the OECD, income inequality at the national level in the late 2000s, as measured by the Gini coefficient, varied from 0.24 for Slovenia to 0.50 for Chile, with Scandinavian and most Eastern European countries below the OECD average of 0.31, and English-speaking and most mediterranean countries above the average. Interestingly, with few exceptions, the national Gini indices have risen significantly over the last 25 years, indicating an increase in household income inequality. While there is a tendency for richer countries, *i.e.*, those with higher median equalised houshold income, to have lower Gini coefficient and hence less income inequality, there are notable exceptions to this pattern, including the USA, the UK and Canada as examples of countries with high median and high income inequality, and several formerly communist countries as examples for low median income and low income inequality.[6]

Like any statistic that aggregates information over an entire population into a single number, in certain cases, the Gini coefficient can fail to capture important differences between populations. For example, a Gini coefficient of 0.5 is measured for a population in which half of the households have no income and the other half all have the same income, as

5 See also: J.Basulto Santos & J.J. Busto Guerrero (2010): "Gini's Concentration Ratio (1908-1914)". Electronic Journal for History of Probability and Statistics, 6(1).

6 OECD (2011): Society at a Glance 2011: OECD Social Indicators, pp. 66-67, OECD Publishing.

well as for a population in which 75% of the income is equally distributed over 25% of the population and the remaining 75% population equally share 25% of the income.[7] But its main shortcomings as a measure of fairness or unfairness stem from more fundamental problems inherent in the notion of equal share. While it is easy to agree that *ceteris paribus*, resources should be distributed uniformly, in practice, differences between members of a population may justify, perhaps even call for, unequal distribution; furthermore, resources may be not be divisible into equal parts, rendering impossible equal distribution.

Before moving on to the notion of equal chance, which addresses especially the latter of these problems, it is worth noting that we live in a world of decidedly unequal shares. In terms of household income, the inequality is in fact considerably more extreme than suggested by the Gini coefficient as a sole indicator: As first observed by the Italian engineer and economist Vilfredo Pareto, the distribution of income over a population tends to be characterised by a mathematical relationship known as a power law, and as a result, income, and likewise, wealth, is concentrated predominantly on very few households or individuals.[8] In practice, this phenomonen is often characterised by the so-called *Pareto principle*, according to which 80% of a resource belongs to 20% of the population. (In reality, power law distributions may be characterised by variants of this principle where the fractions differ from 80/20.) The mathematical power-law distribution named after Pareto further has the property that rules like the 80/20 principle that characterise them can be applied iteratively: of the 20% that hold 80% of the resources, 20% (*i.e.*, 4% of the overall population) hold 80% (*i.e.*, 64% of the overall resources). An important property of such distributions is that their mean over the population is dominated by the few members with the largest share, in the case of wealth, the few richest individuals - one reason why the income or wealth of populations is summarised by the median rather than the mean. (A more detailed discussion and explanations for power-law income distributions can be found in a 2002 article by William Reed.[9])

---

7 See, for example, http://en.wikipedia.org/wiki/Gini_coefficient (last visited on 15 March 2012).

8 Vilfredo Pareto (1897): "Cours d'Economie Politique", Vol. 2, F. Pichou, Paris, France.

9 William J. Reed (2002): "The Pareto law of incomes—an explanation and an extension". Physica A: Statistical Mechanics and its Applications 319: 469–486.

Intriguingly, similarly skewed distributions of shares have been observed in many different contexts; these include the population of cities in various countries, the sizes of corporations, book sales, the popularity of web sites and even the frequency in which words of a given language occur in a large corpus of texts.[10] To illustrate this last example for the English language, the rule known as *Zipf's law*[11] states that the most widely used word (which, not surprisingly, is 'the') occurs about twice as often as the second word in the frequency-ranked list (which happens to be 'of'), about three times as often as the third word in the list ('and'), *etc*.

Overall, it is remarkable to which extent a wide range of human endeavours appear to produce distributions that are in stark contrast with the 'equal share' principle, and one may well conjecture that the dynamics responsible for this phenomenon would make it difficult under many circumstances to bring about fairness in the simplistic sense of 'equal share'.

III.Equal Chance

An interesting problem arises when an indivisible resource has to be distributed among a group of interested parties. Returning to our example involving the allocation of training time on a sports field to two teams, this situation arises when both teams have an equally high preference for a specific slot – say, Saturday afternoon – and that slot can be given to one team only. A classical solution to this problem is based on a random process, such as a coin toss, which gives each team an equal chance of winning the desired Saturday afternoon training time. We say that this process, and the coin used to implement it, is fair, if each team has the same probability of winning - in our example, ½.

In general, the idea of using *randomisation,* often in the form of a simple game of chance, to fairly allocate indivisible or fundamentally unequal resources is widely used and very powerful. Although the outcomes obtained in this way typically give substantial advantages to some participants, there is an intuitive perception of fairness as long as no one

---

10 Mark E.J. Newman (2005): "Power laws, Pareto distributions and Zipf's law". Contemporary Physics 46(5): 323-351.

11 Named after the American linguist George Kingley Zipf, who proposed it in 1935; see: George K. Zipf (1935): "The Psycho-Biology of Language". Houghton Mifflin, Boston (MA), USA.

has privileged access to those advantages, *i.e.,* everyone has the same chance of coming out ahead. If, on the other hand, it turned out that the equal chance principle had been violated - for example, because a biased coin, a loaded die or a stacked deck of cards had been used – participants would feel treated unfairly.

The equal chance principle as a fundamental notion of fairness arises perhaps in its purest form in games of chance played for the purpose of monetary gains, such as roulette, slot machines or lotteries. From a technical perspective, rational participants in this type of gambling trade a modest amount of money they own against the chance of a larger winning. Games of this nature are tightly regulated to ensure fairness in the sense of equal chances of winning for all participants - at least under the assumption that they play in what they are assured should be equivalent ways, such as betting on any single number in roulette, pulling an arm on any of a number of identical slot machines or selecting any combination of numbers in Lotto 6/49.

Unfortunately, ensuring precisely equal chances in such games is very difficult. This problem is illustrated particularly well for the example of dice – perhaps the most widely used objects expressly made for producing random outcomes with equal chance. In a large-scale study involving 4.38 million throws, Iversen *et al.* demonstrated that the inexpensive dice widely used in board games show a small, but consistent and statistically bias towards even numbers, which were obtained overall in 50.72% rather than the 50% of cases expected of a fair die.[12] This bias is explained primarily by the fact that these dice have shallow holes to indicate the numbers, filled with a drop of paint each, which ever-so-slightly reduces the weight of the faces with even numbers (2,4,6) compared to those with odd numbers (1,3,5).

The deviation of 0.72% from the results produced by an ideal, fair die seems very small. Unfortunately, by repeatedly using such an unfair die, this bias can get amplified rather rapidly. To illustrate this effect, let us consider we are using our imperfect die to determine the winner in a simple game of chance involving two players, Alice and Bernard. Alice wins if a throw of the die produces an even number, Bernard wins otherwise, and the loser pays the winner 1 dollar. Alice and Bernard start with the same amount of money each and continue

---

12 G.R. Iversen, W.H. Longcor, F. Mosteller, J.P. Gilbert and C. Youtz (1971): "Bias and runs in dice throwing and recording: A few million throws". Psychometrika 36(1): 1-19.

playing until one of them has lost everything. Using standard statistical techniques, it can be shown that if both players start with 10 dollars, the probability that Bernard, who is every-so-slightly disadvantaged by the imperfect die, ends up bankrupt is 4/3 that of Alice losing everything (~57% *vs* ~43%).[13] If both players start with 100 dollars, the number of rounds to be played is likely much higher, and the probability of ending up bankrupt is now over 17 times higher for Bernard than for Alice (~94.7% *vs* ~5.3), and with an initial capital of 250 dollars, Bernard ends up bankrupt more than 1300 times as likely as Alice (~99.925% *vs* ~0.075%). Starting with 1000 dollars each, Bernard is virtually certain to end up bankrupt: The actual likelihood for Bernard to *not* end up bankrupt is below 1 in 3 trillion - by comparison, the odds of dying as a result of a catastrophic asteroid or comet impact have been estimated at around 1 in 20 000.[14]

This example demonstrates how, under certain circumstances, small deviations from fairness in terms of equal chance can lead to substantial unfairness. We refer to this phenomenon as *unfairness amplification*. It is based on a compounding effect in a setting where slightly biased decisions are made over and over again. The effect is, in fact, qualitatively the same as that observed in the compounding of interest. The inequality in the end result increases rapidly and in a highly non-linear fashion with the number of rounds and with the magnitude of the initial bias. In our simple game, the larger the initial capital (or, equivalently, the smaller the stakes in each round) and the larger the bias in the die used, the higher the probability of Bernard ending up bankrupt. Although the previously cited study by Iversen *et al.* also found that precision dice, as used in casinos, are much less biased than the inexpensive ones considered in our example, as physical objects, they are still imperfect and therefore subject to bias and unfairness amplification.

Even better mechanisms for the unbiased generation of random events exist in the form of so-called pseudo-random number generators (PRNGs) – computer software based on carefully designed mathematical techniques that can generate sequences of numbers that are statistically indistinguishable from those produced by an unbiased source of randomness such

---

[13] The formula used for computing the probability of bankruptcy for Bernard is [1-(*q*/*p*)^*c*]/[1-(*q*/*p*)^(2**c*)] for a biased die producing outcomes with probability p in Alice's favour and q=1-p in Bernard's favour and initial capital c for both players.

[14] C.R. Chapman & D. Morrison (1994): "Impacts on the Earth by asteroids and comets: assessing the hazard", Nature 367 (6458): 33–40.

as a perfect die.[15] PRNGs are very interesting from a philosophical point of view, since the way in which they generate numbers is entirely deterministic. Therefore, knowing the internal state of the generator, the number it will produce next can be predicted with certainty; yet, when the state of the generator is unknown, such predictions are as impossible as they would be for a perfect die, even if a long history of numbers obtained from that same device (PRNG or perfect die) in the past were known.

Conceptually, the way in which this unpredictability is achieved for a PRNG can be understood as being based on the principle of *uncertainty amplification*, where very small differences at the beginning of the iterative generation process may give rise to arbitrary (small or large) differences later on.

Pseudo-random number generators are used extensively in on-line gambling, as well as in electronic slot machines; they also play a crucial role in the modern cryptographic techniques underlying the transmission of sensitive information over the internet and, in particular, secure financial transaction.

High-quality pseudo-random number generators are readily available to software developers, and generators of somewhat more modest, but still very good quality can be found on-line (*e.g.*, at www.randomizer.org[16]). Furthermore, true random numbers, based on measurements of physical phenomena believed to be random, can be obtained from web-based services (notably, www.random.org[17]). Finally, computational methods can be used to detect even subtle biases in repeated random decisions, and such biases can be quantified easily, for example by means of relative entropy or the p-values obtained from various statistical tests. Thus, the concern of unfairness amplification due to biased random decisions can be addressed cheaply and effectively. This ability greatly facilitates the application of the

---

15 See, for example, M. Matsumoto & T. Nishimura (1998): "Mersenne twister: a 623-dimensionally equidistributed uniform pseudo-random number generator". ACM Transactions on Modeling and Computer Simulation 8(1):3-30.

16 Urbaniak, G. C., & Plous, S. (2011). Research Randomizer (Version 3.0) [Computer software]. Last visited on 10 March 2012.

17 Built and operated by Mads Haahr, School of Computer Science and Statistics at Trinity College, Dublin, Ireland. Last visited on 10 March 2012.

principle of equal chance, in the sense of equal probability of winning a desirable (or avoiding an undesirable) outcome.

Yet it is not difficult to identify cases where the application of the principle of equal chance leads to outcomes perceived as unfair. For example, we would not expect that in a fair football game, both teams have precisely the same probability of winning – unless, of course, we consider both teams to be exactly equal in playing strength. Likewise, we would not find it unfair if between two applicants for a position, the better qualified person ends up landing the job. These situations are characterised by the fact that there are differences between the participants that justify unequal chances. What matters in such situations is that the differences that give rise to unequal chances (or, in fact, outright unequal shares) are relevant to and commensurate with the differences in outcome, and that the mechanisms used to ensure this are transparent and accepted by all participants. This observation gives rise to the final notion of fairness we discuss in this chapter: that of a uniform mechanism.

IV.Uniform Mechanisms

We use the term *uniform mechanism*[18] to refer to a well-defined set of rules that treats equivalent participants in an interactive endeavour in the same way; in this context, *equivalence* refers to equality in all aspects relevant to the given context. Additionally, when participants are not equivalent, the differences in their treatment and in the outcomes of that treatment should be commensurate with the magnitude of the respective deviation from equivalence. In other words, a participant who is only slightly better positioned with respect to the relevant criteria compared to another participant should only have a slight advantage.[1] Different from the notions of equal share and equal chance, which are based on outcomes, in terms of allocation of a resource or probability of allocation, uniform mechanisms represent a

---

[18] The general term mechanism denotes a technical concept in game theory that is roughly consistent with the less formal way in which we use it here. Within game theory, the area of mechanism design is concerned with devising rules and processes that govern the co-operation (and resolution of conflict) within a group of participants and display specific desirable properties.

[19] This additional requirement defines a strong notion of a uniform mechanism. A weaker form that does not include this requirement is interesting from a theoretical point of view, but less useful in the context of fairness, since it is easy to see how it would include mechanisms producing outcomes readily perceived to be unfair.

procedural approach to fairness. They can be seen as a codified form of impartiality that seeks to avoid the unfairness arising from differential treatment of participants that should be considered equivalent - this, of course, is the notion of fairness that underpins our legal system, as well as most forms of competitions, including sports.[20]

Challenges in developing and applying uniform mechanisms arise (1) from difficulties in designing sets of rules that guarantee impartial treatment; (2) from assessing equivalence and the magnitude of deviations from equivalence between participants; (3) from preventing disproportionate differences in treatment (and outcome) when dealing with deviations between non-equivalent participants; and (4) from convincing the participants of the uniformity of the mechanism.

Challenge 1 is sometimes addressed by anonymisation, as used, for example, in the double-blind reviewing processes that govern the assessment of work submitted to top-tier archival publication venues in many areas of computing science. Challenge 2 follows from the difficulty of quantitatively assessing the positioning of participants (in the sense of qualifications, merit, fitness, *etc.*), and in particular, from separating aspects that should be considered irrelevant in the given context from those that are relevant, as well as from determining the relative importance of the latter. Rational and defensible solutions to the problem of quantitatively assessing participants play a key role towards meeting challenges 3 and 4. However, challenge 3 can be exacerbated by a generalised version of the notion of unfairness amplification discussed in the previous section.[21]

To illustrate this form of unfairness amplification, we return to our example of allocating training times on sports field to two teams. Let us assume that the teams play at different levels, and that an allocation mechanism is used that gives the better team additional time. This form of differential treatment appears transparent and defensible, and the mechanism *per se* may well be considered fair by both participating teams. Yet, it is easy to imagine how through the advantage of additional training time, the better team may further

---

[20] This notion is also discussed in Section IV of the chapter by Kingsford Smith and highly related to the connection between fairness and perspective-taking central to the chapter by Warren.

[21] This generalised notion also applies to inequalities in perspective taking resulting from differences in resource endowments or status, as discussed in Section III of the chapter by Warren.

increase its strength compared to the weaker team. Thus, over time, the advantages afforded to the stronger team can compound, leading ultimately to a strongly biased and self-perpetuating allocation of training time.

As can be seen from this example, compounding effects from repeated application of a mechanism can amplify and stabilise small initial differences in standing. Even for mechanisms initially regarded as uniform and fair, typically based on consideration of a limited number of applications, these differences can quickly reach magnitudes at which they will readily be perceived as markedly unfair. In this sense, one could even speak of an *emergence of unfairness* from what appeared to be a uniform mechanism. Unfortunately, most complex systems – whether they are natural or man-made – have the property that small effects, such as minor differential advantages, are, at least under certain circumstances, reinforced over time.

Many evolutionary and economic processes give rise to this phenomenon. For example, in the area of investment, access to more information tends to increase the probability of financial gains, and additional capital can be used to obtain better access to information. Even if market participants were guaranteed the same access to data (*e.g.,* real-time market data), the use of additional computational resources could easily lead to financial gains (*e.g.,* through better assessment of risk), which in turn could be used to further improve computational resources.[22] This last example is based on the observation that the usefulness of data generally increases with the ability to extract information from the data, which requires computation; it applies equally to situations where the computation is carried out by silicon-based computer hardware or by human brains.[23]

In light of the possibility of unfairness amplification or emergence, it is important to consider - and, if necessary, limit or counter-act - possible compounding effects in mechanisms that are used repeatedly. In our sports example, this could be achieved by guaranteeing a minimum amount of training time to every team, irrespectively of their

---

[22] A related observation can be found in Section III of the chapter by Kingsford Smith, who points out that equal access to securities information may provide equal inputs (in terms of data), but does not guarantee equal outputs (in terms of financial gains).

[23] This line of argument is based on the notion of bounded rationality, see, e.g.: Herbert Simon (1957): "A Behavioral Model of Rational Choice", in Models of Man, Social and Rational: Mathematical Essays on Rational Human Behavior in a Social Setting. Wiley, New York, USA.

standing. Progressive taxation and social security programs can at least to some extent be seen to fulfill a similar purpose. Where the mechanisms under consideration are too complex for rigorous mathematical analysis, computational simulation can and should be used to study their properties, as well as the efficacy of modifications aimed at addressing the problem of unfairness amplification.

Finally, it is important to recognise that fairness is often a matter of perception, and even a carefully constructed uniform mechanism can fail to produce outcomes perceived to be fair by the participants or independent observers. This recognition gives rise to challenge 4 identified earlier in this section: to convince the participants of the uniformity of a given mechanism. While an in-depth discussion would exceed the scope of this chapter, we note that there are several strategies that can help to address this challenge. Firstly, it is useful to separate the design of the mechanism from its application. Ideally, the design is carried out or at least coordinated by disinterested parties, *i.e.*, individuals or groups who do not stand to benefit or suffer from its application. In any case, the design process should be rational, transparent and inclusive – ideally, by ensuring that those likely to end up with differentially less desirable outcomes are involved equally to those who are likely to obtain advantages when the mechanism is applied. Simulations can play an important role in allowing all participants to gain intuitive understanding of the mechanism before it is accepted for real-world application. Secondly, we believe that by equipping the mechanism with affordances for eliciting feedback and performing corrective modifications based on experiences from its application, especially during early stages, it is possible to increase trust in the intention to ensure uniformity of the mechanism and fairness of outcomes, and that this trust aids the perception of fairness of the mechanism and of the results obtained from its applications.[24]

V. Conclusions

[24] These ideas were inspired in part by conversations between the author and Mark Warren that took place over the time that both of them spent as Distinguished Scholars in Residence at UBC's Peter Wall Institute for Advanced Studies; it is therefore perhaps not too surprising that these ideas complement, at least to some degree, the discussion of the important role of democratic institutions and processes in establishing and maintaining fairness found in Warren's chapter. Likewise, the notion of fairness as a matter of perception resonates with Warren's concept of fairness as perspective taking.

Of the three notions of fairness discussed in this chapter, equal share, equal chance and a strong notion of uniform mechanisms, the first two are easier to quantify, yet limited in their applicability, in the sense that we can easily find situations in which they would be violated without necessarily creating a perception of unfairness.[25] Equal share, while at least in principle easy to implement and test, in most cases, is overly simplistic. We have furthermore demonstrated that equal chance is much more difficult to ensure, and that minor deviations from equal chance can lead to large differences in outcome (*e.g.*, share) – through a compounding effect we refer to as *unfairness amplification*. Finally, we believe that the notion of a uniform mechanism, while overall still too weak to capture all relevant aspects of fairness, appears to provide a necessary condition in the sense that it is difficult to see how obviously non-uniform processes could be generally perceived to be fair. We further believe that care needs to be taken in the design of mechanisms, to ensure that their properties are desirable and well-understood – in particular, with respect to the potential risk of amplification and feedback effects. Ideally, uniform mechanisms should be protected against emergence and amplification of unfairness.

Overall, we posit that fairness is not an intrinsic property of a situation, procedure or mechanism, but to a large degree a matter of perception. Certain types of unfairness are easy to detect and broadly agreed upon, and quantitative measures of equality can be useful to recognise and ameliorate those. In particular, such measures can help us to sharpen our analysis and understanding of situations perceived to be unfair. Thus, while ultimately, fairness may be a notion that is sufficiently multi-faceted and subjective to resist our attempts to formally define and measure it, those attempts can still be useful in exploring salient aspects - facets of fairness.

---

[25] Still, even the simple notion of equal share has applications, for example in the context of financial markets, as discussed in Section III of the chapter by Kingsford Smith.