

An Empirical Evaluation of Fair-Division Algorithms

Nicolas Dupuis-Roy (nicolas@dupuis.ca)

Département de Psychologie, C.P. 6128, succ. Centre-ville
H3C 3J7, Montréal, QC, Canada

Frédéric Gosselin (frederic.gosselin@umontreal.ca)

Département de Psychologie, C.P. 6128, succ. Centre-ville
H3C 3J7, Montréal, QC, Canada

Abstract

Fair-division problems are ubiquitous. They range from the day-to-day chore assignments to the Israeli-Palestinian conflict and include the division of an inheritance to the heirs (Brams & Taylor, 1999; Massoud, 2000). Many intuitive and self-implementable algorithms guaranteeing “fairness” have been devised in the past 50 years (Brams & Taylor, 1996). So far, very few empirical studies have put them to the test (Daniel & Parco, 2005; Schneider & Krämer, 2004). In fact, it is not even known to what extent the solutions derived from these algorithms are satisfactory to human players. Here, we present an experiment that investigated the satisfaction of two pairs of players who divided 10 indivisible goods between themselves. A genetic algorithm was used to search for the best division candidates. Results show that some of the best divisions found by the genetic algorithm were rated as more mutually satisfactory than the ones derived from six typical fair-division algorithms. Analyses on temporal fluctuation and non-additivity of preferences could partially explain this result. Ideas for the future implementation of a more flexible and unconstrained approach are discussed.

Keywords: Fair-division, cake-cutting algorithms, fairness, justice, ethics, genetic algorithms.

Introduction

Inspired by the age-old *I divide the cake and you choose the piece* procedure, Steinhaus, Banach, and Knaster developed, in the 1940's, the first mathematical algorithms aiming at fairly dividing a cake between more than two players. This started what has been recently phrased « the quest for a “magic formula” to resolve conflicts » (Schneider & Kramer, 2004, p.507). The goal is to design a division procedure guaranteeing one or several fairness criteria such as proportionality¹, envy-freeness², equitability³, or efficiency⁴—even against the greediest opponent. Additional desirable properties include simplicity, self-implementation, and applicability to any number of participants in real-life conflicts.

¹ A solution is proportional when each of the n players receives a share that is worth at least $1/n$ of their individual utility.

² A solution is envy-free when each of the n players receives a share that is worth more (or the same) for him than any other share, in terms of his individual utility.

³ A solution is equitable when all of the n players receive a share that is worth the same proportion of their individual utility.

⁴ A solution is efficient if no other division solution can increase a player's utility without decreasing the other player's utility.

Since the seminal work of Steinhaus and colleagues, dozens of fair-division algorithms have been designed (for a review, see Barbanel, 2004; Brams, 2008; Brams & Taylor, 1996; Moulin, 2003; Robertson & Webb, 1998; Young, 1994). One famous example is the *Adjusted Winner* (see Brams & Taylor, 1996 for complete details). Suppose m inherited goods are to be split between two heirs. Briefly stated, this procedure first requires that both heirs privately express their true material preferences for the disputed goods by distributing a total of 100 points over them (more points means higher preference). Then, each object is given to the player who allocated more points to it. If the players' total points are unequal after this initial distribution, then the procedure makes an additional adjustment: the most disputed object (i.e. the one for which the ratio of both players' points is closest to one) is split in such a proportion that both players end up with the same total amount of points. If the heirs express their true preferences—and Brams & Taylor (1996) demonstrated it is optimal for a rational and self-regarded agent to do so—they will be assured a proportional, envy-free, equitable, and efficient solution.

So far, very few empirical studies have put these fair-division algorithms to the test (Daniel & Parco, 2005; Pratt & Zeckhauser, 1990; Schneider & Krämer, 2004). In fact, it is not even known to what extent the solutions derived from these algorithms are satisfactory to human players. In the next paragraphs, we argue that it is not likely to be the case because most fair-division algorithms dealing with indivisible goods rest on invalid assumptions regarding human preferences. More specifically, they assume stable and additive material preferences.

Recall that in the *Adjusted Winner*, the players are instructed to make a single and overt appraisal of their material preferences. Likewise, most algorithms dealing with indivisible objects require that a fixed or unlimited amount of points or monies be distributed over the disputed goods. This is at odd with behavioral evidence that shows systematic fluctuation of individual utility over time (Rabin, 1998). A change occurring between preferences' appraisal and implementation of the solution derived from the fair-division algorithm could hamper its fairness and leave the group with a feeling of injustice. So, are preferences stable in time? And, if not, what is the impact on the fairness of the algorithms?

Most fair-division algorithms dealing with indivisible goods assume the additivity of material preferences. For

example, in the case of the *Adjusted Winner*, this means that the points given to two disjoint bundles of goods is equal to the sum of the points of its constituents. However, non-additive preferences are likely to be observed when goods have closely related functions: for example, a *bayan* and a *dayan* (the two drums of an Indian tabla set) might be worth more together than separated. Most fair-division algorithms do not tolerate non-additivity, which can prevent them from reaching the desired solution. Well aware of this potential problem, Brams and Taylor (1996, 1999) have suggested that goods be carefully organized, prior to the application of the algorithm, in separable bundles (e.g., *bayan* and *dayan* in the same bundle). This solution seems reasonable but non-additivity might be difficult to predict considering individual variability: for example, a non-musician might consider a *bayan* and a *dayan* as two exotic pieces that could be displayed either together or separately. Hence, can we really control for the additivity of preferences?

Taking into account these two issues, we hypothesize that the solutions derived from these algorithms are not likely to be satisfactory for human players. We tested this hypothesis and pinpoint two of its underlying causes by measuring the satisfaction of two pairs of players who divided 10 indivisible goods between themselves. More precisely, participants rated their satisfaction over a subset of all 1,024 possible divisions, and then estimated their material preferences by distributing a total of 100 points over the 10 objects. The division space was explored using a genetic algorithm (GA; Holland, 1975) that searched for best division candidates. More precisely, it was designed to converge toward divisions maximizing the satisfaction of both players (details are provided in the *Algorithm* section). The rationale behind the use of this GA was manifold. First, GAs can perform even when, as we hypothesized, non-additive (non-linear) and dynamic preferences are involved. Second, its natural tendency to duplicate solutions allows us to evaluate the temporal fluctuation of preferences. Finally, this GA is also a fair-division algorithm of its own that could be tested against standard fair-division algorithms.

Methods

Participants

Four co-workers (including the authors of this paper) that had known each other for at least three years were grouped in pairs. Players 2 and 3 were naïve to the purpose of this experiment. Each pair was informed that the goal of the experiment was to divide ten items between them, and that they would have to express their satisfaction individually about at least 300 candidate divisions. Prior to the experiment, subjects were given one minute to familiarize themselves with the items. Informed consent was obtained and no monetary compensation was provided.

Indivisible goods

Ten food items were designated as the indivisible disputed goods: asparagus, a six-pack of beer, a jar of almond butter,

a bag of chips, fresh dills, a can of concentrated orange juice, mushrooms, a pizza, a salmon filet, and half a pound of shrimps. They were chosen as possible ingredients of meals and we purposely selected potentially non-separable items (e.g., a salmon filet and fresh dills). Color pictures of these items were taken on the Internet and resized to span about 4 to 5 cm on the screen at full resolution. Viewing distance was about 60 cm.

Apparatus

Our Matlab (MathWorks) experimental programs used functions from the Psychophysics toolbox (Brainard, 1997; Pelli, 1997) and ran on two Macintosh G5 computers linked via the intranet. All stimuli were presented on two Sony Trinitron monitors.

Procedure

Paired participants were tested on two linked computers. They were each submitted to a minimum of 15 runs. During each run, they were shown a population of 20 candidate divisions (not necessarily the same ones). On the computer monitor, one row contained the goods given to one participant and another row contained the goods given to the other participant. Participants were told they had unlimited time to express, on a scroll bar, her/his satisfaction about the candidate division on a scale ranging from 0 (not satisfied at all) to 1 (fully satisfied). Each run lasted approximately 2.5 minutes. After the testing session, participants were asked to distribute 100 points over the ten items to reflect their preferences for these objects (material preferences). Finally, they were debriefed. The entire experiment lasted approximately one hour.

Genetic algorithm

The genetic algorithm (GA) started with two populations of 20 randomly generated candidate divisions. A vector of ten binary elements represented a candidate division; each binary element coded an object and its value specified the owner of this object. Each paired participant rated the candidate divisions from her/his population. The GA maximized each subject's rating independently. The fitness of a candidate division was computed independently for each population as follows: the satisfaction rating of a participant divided by the sum of all of her/his satisfaction ratings. This normalization per subject compensated for discrepancies between paired participants. Two new populations of candidate divisions were generated independently for each participant following the so-called roulette-wheel selection. Next, the two resulting populations were mixed through a single-point crossover (50%) and randomly split in two distinct populations. Finally, 1% of the candidate divisions were randomly mutated (for an overview of genetic algorithms, see Mitchell, 1996). This GA implementation ensured that the ratings of each participant equally contributed to the next populations of candidate allocations thus avoiding the adoption of any particular definition of fairness. The testing session was

stopped after a minimum of 15 runs and after having reached a minimum clone (duplicate candidate divisions) rate of 0.75 or a mean satisfaction level above 0.9.

Results and Discussion

The average satisfaction rating increased rapidly from 0.49 ($SD = 0.0003$) at the first run to 0.85 ($SD = 0.001$) after 13 runs. The clone rate also jumped from 0 to 0.78 over the first 13 runs, indicating a rapid convergence toward “fair” candidate divisions.

Do fair-division algorithms maximize satisfaction?

As pointed out in the introduction, most fair-division algorithms dealing with indivisible goods derive their solutions from a single appraisal of material preferences. In our case, participants were asked to express their material preferences by distributing a total of 100 points over the 10 objects. By contrasting the satisfaction ratings with the total number of points they would have obtained for each solution candidate (henceforth called material utility), we can test the extent to which solutions prescribed by fair-division algorithms on the basis of material preferences are satisfactory.

Figure 3 consists of two scatter plots, one for each pair of participants. Each plot summarizes the relationship between the participants’ satisfaction ratings and the material utility of the solutions. More specifically, material utility is the proportion of the 100 points a participant would have won for a given solution candidate. This measure could be seen as a variable that predicts the participant’s satisfaction rating. And, for fair-division algorithms to work correctly, satisfaction should closely map onto material utility.

The scatter plots depict four dimensions on a flat surface: the x-y axes represent each participant’s material utility (e.g. 1 means all objects are assigned to that participant) and the gray level of each joint squares (or rectangle formed of two squares) indicates the satisfaction rating (black means “not satisfied” and white “fully satisfied”). More precisely, each joint squares represents one candidate division rated by the two participants of a pair. The left ones show player’s 1 (or 3) satisfaction ratings and the right ones show player’s 2 (or 4) satisfaction ratings. This side-by-side arrangement was chosen to visually enhance the contrast between the participants’ satisfaction. Homogeneous gray rectangles show satisfaction agreement whereas highly contrasted joint squares mean satisfaction disagreement. All solution candidates located in the upper right portion (delimited by dotted lines) of the scatter plots are envy-free and proportional because both participants won at least 0.5 of their total material utility. The solutions located on the dotted diagonal are equitable. Black crosses indicate the divisions derived from six fair-division algorithms (based on material preferences): the *Sealed Bid Knaster* (Knaster, 1946), the *Adjusted Winner* (Brams & Taylor, 1996), the *Adjusted Knaster* (Raith, 2000), the *Division by Lottery* (Pratt, 2007), the *Descending Demand* (Herreiner & Puppe, 2002) and the *Balanced Alternation* algorithm (Brams &

Taylor, 1999). All of these algorithms except the *Descending Demand* and the *Balanced Alternation* proceed in a similar fashion: first, the players evaluate their material preferences by distributing a fixed (or unlimited) amount of points (or monies) over the disputed goods; second, each object is allocated to the highest “bidder”; and third, the procedure adjusts this first distribution by either splitting one of the objects, asking monetary compensation or randomly assigning one of the objects. Details about the algorithms are available in the cited papers.




Figure 1. Each rectangle (composed of two joint squares) represents one candidate division rated by the two participants of a pair. The x-y coordinates of each rectangle represent the material utility of each paired participants for the corresponding division. The right (players 2 and 4) and left (players 1 and 3) gray levels of a disc indicate the respective satisfaction of paired participants (black=“not satisfied” and white=“fully satisfied”). Red circles represent the divisions prescribed by the five fair-division algorithms.

In the upper scatter plot, material preferences explain 48% and 1% of the variance of satisfaction ratings of player 1 and 2 respectively. It can be appreciated at a glance: most light and relatively homogeneous rectangles are clustered on the left side of the scatter plot. This suggests that the most

satisfactory solutions have very low material utility. Also, this could mean that some players had difficulty estimating their material preferences. All of this is bad news for most fair-division algorithms, which mostly rely on these material preferences.

In the bottom scatter plot, material preferences explain 28.79% and 47.06% of the variance of the satisfaction ratings of player 3 and 4 respectively. Light and homogeneous rectangles are concentrated in the envy-free and proportionality region. There is also a clear overlap between the divisions prescribed by the fair-division algorithms (black crosses) and the candidate divisions selected by the GA (joint squares). Some of the divisions prescribed by the fair-division algorithms were highly satisfying (especially those of the *Descending Demand* and the *Balanced Alternation* algorithms) but other candidate divisions sampled by our GA were even more satisfying. In fact, some of the most satisfying candidate divisions were outside the envy-free and proportionality region. These results strongly advise against using only material preferences to derive fair divisions. Next, we consider two possible explanations for this result.

Satisfaction fluctuates

As mentioned in the introduction, one explanation for the small correlation between satisfaction and material utility is the fluctuation of the preference measure. Accordingly, a single and overt appraisal of material preferences might not be precise enough. To verify this possibility, we first analyzed the overall fluctuation of satisfaction ratings of the clones. All participants exhibited a fairly high level of fluctuation ($SD = 0.15, 0.24, 0.12$ and 0.06 respectively), especially for mid-rated clones (between 0.25 and 0.75 satisfaction units). Two questions logically follow this observation: Are these fluctuations caused by some random factors? And what could be the effect of these fluctuations on fair-division algorithms?

To address the first question, we looked at the slope of the line fitted on satisfaction ratings and time (or trial). We focused on variations occurring in the mid-range—between 25% and 75% —to stay at bay from floor and ceiling. Out of 14 sets of clones, five (57%) had a significant slope ($p < 0.05$, Bonferroni corrected). Linear variations of satisfaction ranged from $-0.036/\text{trial}$ to $0.0002/\text{trial}$. This suggests that at least part of the temporal fluctuations were systematic and linear. Verbal reports of the players offer one possible explanation: they all interpreted the selection process of the genetic algorithm (GA) as an indication of “who was winning” and, as a consequence, adapted their strategies. Other possible explanations concern all time-dependent behaviors (e.g. fatigue, habituation, etc) and order effects such that a given trial influences the next one. Unfortunately, pursuing these specific hypotheses goes beyond the scope of this paper.

To address the second question, we first needed to find an estimate of the group random fluctuation level that would be free of systematic (or strategic) fluctuations. We opted for

the standard deviation of the residuals of the line fitted on the clone ratings (see previous analysis). It turned out to be 0.15 in satisfaction units—this is the variation that cannot be predicted by the linear model. We did not look for any non-linear trend because we had no specific expectations about what type of non-linearity it should be. Nevertheless, this allowed us to control minimally for the possibility of a linear interaction between the GA and the players. Moreover, we think the standard deviation of the residuals is a conservative estimate of the random fluctuations because it takes into account all clones, including the highly- and lowly-rated ones, which do not fluctuate much (because of the floor and ceiling effects). To assess the possible impact of such fluctuations on fair-division algorithms, we ran a simulation on five of them: the *Sealed Bid Knaster*, the *Adjusted Winner*, the *Adjusted Knaster*, the *Division by Lottery* and the *Descending Demand*. We did not use the *Balanced Alternation* because it does not guarantee any criterion of justice. We compared the division obtained following the application of each algorithm with deterministic and stochastic (uniform noise with a $SD = 0.15$) agents. By repeating this process 100 times, we estimated the probability that an algorithm will meet its (promised) fairness criteria in a situation with a realistic noise level (Table 1).

Table 1: Percentage of the 100 simulated solutions that satisfied envy-free, equitability, maximin and efficiency. Added uniform noise had either a $SD = 0.15$ (the estimated noise level; bottom numbers) or a $SD = 0.075$ (a more conservative noise level; top numbers).

Procedures	S.B.K.	A.W.	A.K.	D.L.	D.D.
Efficiency	76%	70%	76%	66%	47%
	41%	39%	41%	36%	32%
Equitability	—	17%	14%	—	—
	—	8%	1%	—	—
Envy-freeness	100%	100%	100%	100%	—
	92%	95%	92%	94%	—
Maximin	—	—	—	—	57%
	—	—	—	—	33%

Results indicate that the observed random fluctuations have a decisive impact on the fairness of all procedures. Equitability seems to be the least robust criterion whereas envy-freeness (proportionality is equivalent to envy-freeness with only two players) seems to be the most robust. This was expected because equitability is a very restrictive criterion: both subjects need to win the same total utility. In comparison, envy-freeness is generally more inclusive: it fills up to 25% of the solution space. The least robust algorithm seems to be the *Descending Demand* (Herreiner & Puppe, 2002). In sum, we observed fluctuations of satisfaction that could not be predicted by the linear model, and showed that even a very conservative estimate of these random fluctuations was sufficient to affect the justice of fair-division algorithms.

Preferences are (non-predictably) non-additive

Another explanation for the small correlations observed between satisfaction and material utility is the presence of non-additivity in the material preferences. Additivity implies the independence of object preferences; we thus needed to search for dependencies in conjunctions of objects. This was tested by performing a multiple linear regression on the satisfaction ratings and the outer-products of the candidate division vectors. We estimated the statistical thresholds of regression coefficients using a Monte-Carlo simulation ($p < 0.05$; two-tailed). Significant (black cells) positive (white '+') and negative (white '-') regression coefficients are shown in figure 2. This figure displays two conjunction matrixes that indicate which combination of objects are strongly non-additive. Note that the upper and lower triangles in each matrix correspond to different paired participants.




Figure 2. Two conjunction matrices: Significant (black cells) positive (white '+') and negative (white '-') regression coefficients indicate strong non-additive preferences.

Results show a total of 26 significant preference conjunctions. This is not all that surprising because, as we explained in the Methods section, we did not choose the objects with the intent of separability—on the contrary.

More interestingly, inter-rater reliability, measured by the Fleiss' kappa index (0.04), indicates little agreement between raters. This means a consensus would be hard to obtain in real-life when allotting goods in separable bundles prior to the application of a fair-division algorithm—as proposed by Brams and Taylor (1996, 1999). Overall, the presence of non-additive preferences can partially explain why satisfaction ratings do not map onto material utility.

General discussion

We explored the satisfaction over a large subset of the 1,024 possible divisions of 10 indivisible goods between two pairs of participants using a genetic algorithm (GA). The data from two pairs of participants suggests that fair divisions derived from a single and overt appraisal of material preferences are not necessarily the most satisfactory solutions. Furthermore, we discovered that fair-division algorithms fared rather poorly in comparison to the most satisfactory division candidates selected by our GA. Subsequent statistical analyses suggested two possible explanations for this result: two assumptions of typical point-allocation fair-division algorithms, i.e. stability and additivity of preferences, were violated. More specifically, we observed fairly large fluctuations of the satisfaction ratings over time and further analysis showed that part of this fluctuation was linearly systematic. This means part of the fluctuations were predictable (probably due to order effects, strategic or time-dependent behaviors) and part were mere noise. Computer simulations of deterministic and stochastic players allowed us to evaluate the impact of the residual random fluctuations (the noise) on the fairness of five fair-division algorithms. All algorithms were affected. Unsurprisingly, the most inclusive criterion of justice, envy-freeness (and proportionality), was also the most noise-resistant. The conjunction analysis on pairs of objects revealed that the combinations of objects that were treated non-additively varied considerably from one participant to another. This suggests that the solution consisting in allotting these non-separable objects in bundles (see Brams and Taylor, 1996) would be hard to implement in real-life.

There is at least another possible explanation for the discrepancies between the solutions derived from fair-division algorithms and the one found by our GA: individual attitudes toward justice (e.g. inequity aversion, self-interest, reciprocity, etc) and knowledge about the other players' attitude could have injected additional non-linearity into the satisfaction ratings. For example, a participant averse to inequity would have probably been unsatisfied with a division assigning all objects to him and none to his partner—i.e. a division with maximum material utility. Such attitude could involve a non-linear and possibly non-monotonic relationship between material utility and satisfaction. Hence, non-additivity and temporal fluctuations were not necessarily the only sources of non-linearity in the data. This represents a possible confounding variable in both analyses. From a broader perspective, the diversity of attitudes toward justice (Fehr & Schmidt, 1999; Colman,

2003) could be a real threat to fair-division algorithms because they impose a limited set of justice criteria. Therefore, the choice of a fair-division algorithm could depend on which criteria are preferred. We thus wonder if a consensus can be reached on a set of justice criteria, considering possible discrepancies between individual attitudes toward justice?

There are caveats about our experiment: We did not control for possible non-linear interactions between the selection process of the GA and the participants' strategies. To prevent this from happening, we could have sampled a (much) larger subset of *randomly* selected candidate divisions. Also, this would have allowed a better evaluation of existing fair-division algorithms. Furthermore, additional naïve subjects would need to be tested to gain a better understanding of the interplay between satisfaction and material utility.

We used a GA to search the division space for the most satisfying divisions. Interestingly, our GA also constitutes a fair-division algorithm in itself: it converges toward maximum mutual satisfaction. Perhaps our GA-based fair-division converged toward more satisfactory divisions than typical fair-division algorithms because—unlike typical fair-division algorithms—GA-based fair-division does not assume a shared set of justice criteria and can promptly adapt to fluctuating and non-additive preferences. In any case, we believe that such a flexible and non-normative approach to justice is worth exploring, especially in the context of real-life fair-division problems. Many questions remain unanswered at this point: Can GA-based fair divisions elicit trust in players even though no guarantee of justice is promised in the face of greed? And will the inherent complexity of the GAs evoke suspicion in players? Or, rather, will it diminish the likelihood of manipulation?

Conclusion

As promising as fair-division algorithms might be, their implementation in realistic setting presents great challenges (e.g. Pratt & Zeckhauser, 1990), some of which can be addressed by experimental investigations. Future studies should consider the possible impact of different attitudes toward justice and focus on more detailed comparison between standard fair-division algorithms and alternative algorithms such as our GA.

Acknowledgments

Research supported by FQRNT scholarship and NSERC grants.

References

- Brams, S. J. (2008). *Mathematics and Democracy: Designing Better Voting and Fair-Division Procedures*. Princeton, NJ: Princeton University Press.
- Brams, S. J., & Taylor, A.D (1996). *Fair Division: From Cake-Cutting to Dispute Resolution*. New York: Cambridge University Press.
- Brams, S. J., & Taylor, A.D (1999). *The Win-Win Solution: Guaranteeing Fair Shares to Everybody*. New York: W.W. Norton.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10, 433–436.
- Camerer, C. F. (2003). *Behavioral Game Theory*. New York: 255 pages.
- Colman, A. M. (2003). Cooperation, psychological game theory, and limitations of rationality in social interaction. *Behavioral and Brain Sciences*, 26, 139-153.
- Daniel, T. E., & Parco, J. E. (2005). Fair, efficient and envy-free bargaining: An experimental test of the Brams-Taylor adjusted winner mechanism. *Group Decision and Negotiation*, 14(3), 241-264.
- Denoon, D., & Steven J. Brams. (1997). Fair division : A new approach to the Spratly Islands controversy. *International Negotiation*, 2(2), 303-329.
- Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics*, 114(3), 817-868.
- Haake, C.-J., Raith, M. G. & Su, F. E. (2000). Bidding for envy-freeness: A procedural approach to n-player fair-division problems. *Social Choice and Welfare*, 19, 723-749.
- Holland, J. H. (1975). *Adaptation in Natural and Artificial Systems*. University of Michigan Press: Ann Arbor, MI.
- Kramer, U. S., & Schneider, G. (2003). Fair formulas: Psychological and procedural determinants in the resolution of redistributive conflicts. *Kolner Zeitschrift Fur Soziologie Und Sozialpsychologie*, 55(1), 55.
- Massoud, T. G. (2000). Fair division, Adjusted Winner procedure (AW), and the Israeli-Palestinian conflict. *Journal of Conflict Resolution*, 44(3), 333-358.
- Mitchell, M. (1996). *An introduction to genetic algorithms*. Cambridge, MA: MIT Press.
- Moulin, H. J. (2003). *Fair Division and Collective Welfare*. Cambridge, MA: MIT Press.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision*, 10(4), 437–442.
- Pratt, J. W., & Zeckhauser, R. J. (1990). The fair and efficient division of the windsor family silver. *Management Science*, 36(11), 1293-1301.
- Pratt, J. (2007). Fair (and not so fair) division. *Journal of Risk and Uncertainty*, 35(3), 203-236.
- Raith, M. G. (2000). Fair-negotiation procedures. *Mathematical Social Sciences*, 39(3), 303-322.
- Robertson, J., & William Webb. (1998). *Cake-Cutting Algorithms: Be Fair If You Can*. MA: A K Peters.
- Schneider, G., & Kramer, U. S. (2004). The limitations of fair division - An experimental evaluation of three procedures. *Journal of Conflict Resolution*, 48(4), 506-524.
- Young, H. P. (1994). *Equity in Theory and Practice*. NJ: Princeton: Princeton.