Journal of Economic Behavior & Organization 131 (2016) 21–36

Contents lists available at ScienceDirect

Journal of Economic Behavior & Organization

journal homepage: www.elsevier.com/locate/jebo

Fairness and efficiency in a subjective claims problem☆

Anita Gantner, Rudolf Kerschbamer*

Department of Economics, University of Innsbruck, Austria

ARTICLE INFO

Article history:
Received 22 November 2014
Received in revised form 28 June 2016
Accepted 25 July 2016
Available online 1 August 2016

JEL classification:
D63
C91
D61

Keywords:
Fair division
Subjective claims
Experiment
Mechanisms

ABSTRACT

In a subjective claims problem agents have conflicting perceptions on what constitutes a fair division of a jointly produced cake. In a large-scale experimental study involving a three-agent subjective claims problem, we compare the performance of four mechanisms which use agents' reports on fair shares as input and yield a division of the cake (or less) as output. The mechanisms differ with respect to the desirable properties they possess and they are compared in terms of efficiency and perceived allocative and procedural fairness. Successful in terms of both fairness and efficiency are two mechanisms that explicitly ask for an assessment of the partners' fair shares and that do not induce agents to exaggerate their assessment of the own fair share. One of the two successful mechanisms does not ask for an assessment of the own fair share while the other punishes overly selfish own claims.

1. Introduction

In the oil and gas industry, a single oil pool often underlies the land parcels of several owners. If owners act non-cooperatively, each drilling his own well, oil reserves are depleted at an excessive rate, resulting in large efficiency losses (see Libecap and Wiggins, 1984). An obvious solution to this 'tragedy of the commons' problem is unitization, that is, treating the oil pool as a production unit controlled by a single firm. Unitization, however, brings in the new problem of how the proceeds of the oil field should be divided among the multiple owners. Due to physical characteristics of oil and gas deposits, individual contributions to the joint profit are difficult to compare.[1] As a result, parties typically have conflicting subjective perceptions about the relative contribution of their own property to the joint profit and thus also about what constitutes a fair allocation.[2] An important question then is how to aggregate these conflicting subjective perceptions into a fair division of the joint profit.

☆ We would like to thank the co-editor Marie-Claire Villeval and two anonymous referees for their constructive comments and suggestions. We also thank seminar participants at Theem Kreuzlingen, Bilkent University, ESA and GfEW Meetings and, in particular, Shyam Sunder for useful suggestions and comments. This project was supported by funds of the *Austrian National Bank* (OeNB Anniversary Fund, grant number: 13595) and the *Austrian Science Fund* (FWF, grant number: P 22669-G11).

* Corresponding author at: Department of Economics, University of Innsbruck, Universitätsstr. 15, 6020 Innsbruck, Austria.
E-mail address: rudolf.kerschbamer@uibk.ac.at (R. Kerschbamer).

[1] For example, a simple division according to the size of contributed land might be considered unfair, since some parcels are located over the centre of the oil field while others are located in the periphery (see Cramer et al., 2009).

[2] Indeed, conflicting subjective perceptions about what constitutes a fair allocation of the joint profit and the resulting failure of negotiations about unitization has led major oil producing states in the US to provide legal mechanisms by which an oil pool can be unitized with less than unanimous consent of the right owners. See, for example, Rule 530 in Colodrado's Oil and Gas Conservation Act.

http://dx.doi.org/10.1016/j.jebo.2016.07.019
0167-2681/

The motivation for the present study came from another example where parties' contributions to a joint surplus are difficult to compare: In Tyrol, Austria, a new water power plant will be built in a location that has common borders with three municipalities. The company who owns and operates the power plant pays a certain amount of money to the three municipalities in exchange for the right to use land (i.e. the valley that would need to be flooded in order to operate the power plant) and water. The amount to be paid is no longer an issue as the decision to build the plant has already been made.[3] Now the municipalities have to decide how to distribute the money amongst themselves, and the question arises what would be the fair share for each of them. The question is non-trivial because the municipalities give up different amounts of land and water rights, and it is unclear how land and water should be evaluated on a common scale.

In order to evaluate fair shares, one could consider estimating the contribution of a given partner by removing him from the partnership. However, the market values of the individual rights are fairly small compared to the combined value of the resources. Hence, this procedure of estimating the contribution of a partner is of limited help. Another way of resolving the division problem would be to ask an outside observer as impartial arbitrator to assign shares to the prospective partners. However, such an external agent is typically in a worse position to discern the partners' relative contributions and their fairness perceptions than the partners are themselves, as his information might be incomplete, or biased if provided from the partners' side. The partners may thus prefer to rely on their own subjective perceptions on how the profit should be divided fairly and they may find it useful to use a rule or procedure to reach a consensus. The obvious question then is which rule this should be. The present paper investigates this question by comparing in a large-scale experimental study with more than 600 participants the performance of different cake division mechanisms for a subjective claims problem involving three agents. The three-agent case is actually relevant in practice and it has some special features that disappear when there are four or more agents (see the discussion below).

Our two motivating examples incorporate two different sources for the inefficiency of the non-cooperative solution: In the petroleum-industry example the inefficiency is due to the over-exploration of the joint resource, as captured in the familiar 'tragedy of the commons' problem: Multiple owners are each endowed with the right to use the common resource, and no one has the right to exclude anybody. By contrast, in the power plant example the source of the inefficiency is the under-exploration of the common resource due to what Heller (1998) has termed the 'tragedy of the anticommons' problem: Multiple owners of inputs are each endowed with the right to block their coordinated use, and inputs are almost useless individually.

Our aim is to take a unifying approach for the study of such problems by focusing on their similarities, which are the subjective perceptions of involved parties' fair shares of the common resource and the inefficiency of the non-cooperative solution. To this aim, a stylized version of the problem, which shall be referred to as the *subjective claims problem*, is formulated as follows: Several agents – the partners – have contributed inputs with a stand-alone value of zero to a joint project whose final value is a fixed sum of money denoted by S.[4] Once S is produced, the partners have to divide it amongst themselves. Since inputs are difficult to compare, partners typically have conflicting subjective perceptions about what constitutes a fair division of the joint profit. In the following, we shall refer to an allocation of S that is considered fair by partner i as her *subjective evaluation of claims*. In the case of n partners such a subjective evaluation of claims is a vector with n entries summing up to S.[5] This formulation implicitly assumes that the partners are interested not only in their own material payoff or share of S, but also in the fairness of the allocation. The subjective claims problem is then to find an allocation $s=(s_1, \ldots s_n)$, where $\sum_i s_i \leq S$, or alternatively, a procedure that implements such an allocation, which is considered fair by the partners and which is also efficient.

The present paper compares in a large-scale lab experiment the performance of different cake division mechanisms in a subjective claims problem involving three agents. The mechanisms use agents' reports on fair shares as input and yield a division of the cake (or less) as output.[6] The mechanisms differ with respect to the desirable properties they possess.

Economic theorists have identified several desiderata a cake division mechanism might or might not possess. An obvious minimum requirement is consensuality: If there is a way to divide the cake that agrees with all individual reports then this should be the outcome. A further desirable property is that the rule should not induce agents to exaggerate the own fair share. One way to achieve this is not to ask the agents for an assessment of the own fair share. This property is called impartiality.[7]

[3] The fact that the municipalities agreed to give up their resources at a point in time where they did not know what each of them gets in exchange might seem strange to an economist. But this is exactly what happened in this case.

[4] While in the examples above the combined value of the inputs is worth much more than the sum of the stand-alone values, each input still has some value individually. We abstract from positive stand-alone values for the sake of simplicity.

[5] In contrast to the n privately known vectors of subjective claims described here, the more familiar *objective claims* (or bankruptcy) *problem* refers to one publicly known vector of objective claims, which is infeasible because the sum of the objective claims exceeds the available amount. See Moulin (2002) and Thomson (2003) for surveys, and Gächter and Riedl (2005, 2006) for experiments on the objective claims problem.

[6] Bargaining may be another way to obtain a solution to the subjective claims problem – see Karagözoglu and Riedl (2015) for an application in a related context. In an additional part of our experiment, we also investigated the performance of various unanimity bargaining procedures (one procedure per session). Due to the large amount of data we obtained, we report here only the comparison of the static mechanisms. The performance of the unanimity bargaining procedures is discussed in our companion paper (Gantner et al., 2013).

[7] Impartiality combines objectivity, which requires that the implemented allocation does not depend on any partner's statement about what he deserves relative to the others, and strategy-proofness, which requires that no partner is able to affect his own share by the input that he provides.

The **Impartial Division Rule** proposed by DeClippel et al. (2008) is a rule that satisfies both, consensuality and impartiality. It does so by asking each partner only for an assessment of the relative claims of the other two partners and by implementing an allocation that is consistent with these relative shares if such an allocation exists. DeClippel, Moulin and Tideman show that for the case of three agents the Impartial Division Rule is the only direct mechanism that satisfies consensuality and impartiality. Thus, by investigating the performance of this mechanism we can address the question whether consensuality and impartiality are actually desirable properties. A downside of the Impartial Division Rule for the three-agent case is that if there is no division that is consistent with the reported opinions about relative fair shares then the rule divides less than the total amount available. This results in an inefficiency that increases in the degree of inconsistency of the reported relative shares.[8] This allows us to determine the efficiency cost of requiring impartiality and consensuality and to see who pays this cost. This is done by comparing the Impartial Division Rule to other mechanisms that may reduce inefficiency but are not impartial and consensual. For our main comparisons, we selected two other mechanisms that have in common that they replace impartiality by a punishment of overly greedy assessments of the own fair share.

The **Modified Nash-Demand Rule** (Mumy, 1981) asks each partner only for an assessment of the own fair share. If the sum of reported 'own claims' does not exceed *S*, each partner receives the reported own claim. Otherwise each partner receives the own claim minus a penalty that is a multiple, greater than 1, of the difference between the sum of reported own claims and the amount available. Thus, this mechanism potentially has two sources for inefficiency: One arises when the sum of reported own claims exceeds *S*, while the other arises when this sum falls short of *S*. In comparison to the Impartial Division Rule, this mechanism imposes a larger punishment in case of disagreement of the reports, thus giving a higher incentive for coordination of reports. Thus, the comparison to the Impartial Division Rule potentially gives information about the consequences of forcing agents to coordinate (by increasing the punishment when they do not do it) instead of requiring impartiality.

Both the Impartial Division Rule and the Modified Nash-Demand Rule distribute less than the value of *S* if reports are inconsistent. The third mechanism in our comparison has only an efficiency cost if reports are overly 'prudent': The **Extended Divide-the-Dollar Rule** (Brams and Taylor, 1994) asks each partner to report an evaluation of the own claim as well as the partners' claims. If the three own claims sum up to *S* or less, each partner receives the reported evaluation of the own claim. Otherwise, own claims are paid out sequentially up to the point where *S* is depleted, giving priority in the order of 'greed levels', starting with the lowest one. To calculate a partner's greed level, the mechanism subtracts the average of the claims assigned to this partner by the other two members of the partnership from the reported own claim. If this difference is positive, the partner is greedy and the difference is his greed level. The only source for inefficiencies under this mechanism are own claims that sum up to less than the value of the cake. Since reported fair shares are likely to be biased by material self-interest, this mechanism is expected to be highly efficient. Thus, comparing its performance to that of the Impartial Division Rule potentially gives information about the efficiency cost of requiring impartiality and about who pays this cost.

Finally, we test – in some additional sessions – a fourth mechanism that also gives up impartiality but does not replace it by a punishment of overly selfish own claims: The **Average Division Rule** asks each partner to report an evaluation of the own claim as well as the partners' claims. The share of a partner implemented by this mechanism is then calculated as the average of the claims assigned to this partner, implying that the mechanism is fully efficient by design. By having a mechanism in our comparison that is fully efficient by design but does neither respect impartiality nor impose a punishment for overly selfish own claims we intend to assess the importance of imposing one of those two requirements (in addition to consensuality, which is also respected by this mechanism). If agents report their subjective evaluations of claims truthfully even without such corrective devices, then this mechanism should perform well because it implements a compromise allocation in this case.

The details of our experiment are as follows: First, we generate a subjective claims problem involving three partners in a lab experiment. This is done by having subjects perform real effort tasks within different cohorts and earn points depending on their relative performance within their cohort.[9] The points a subject earns are his contribution to the partnership, and all three subjects in a partnership earn their points from exerting real effort in a different cohort. This makes it difficult for subjects to compare their performances. Furthermore, the production function translating individual contributions to the joint profit is non-linear.[10] The joint profit is then distributed among partners with the help of the above-mentioned mechanisms.

None of these mechanisms has been tested in lab experiments before. In the present study, we compare them in terms of efficiency and fairness. The efficiency comparison refers to the fraction of *S* that is finally paid out to the partners.[11] As

[8] The authors also show that for more than three agents, there is a family of division rules satisfying these two properties, which are also efficient. Thus, by investigating the special case of three agents, we test this mechanism under adverse conditions regarding efficiency.

[9] We let subjects acquire their contributions to the cake in real effort tasks to induce strong entitlements. Numerous empirical studies confirm that earned rights induce stronger entitlements than those acquired by luck (see e.g. Hoffman and Spitzer, 1985; Burrows and Loomes, 1994; Schokkaert and Lagrou, 1983).

[10] We let partners acquire their points in different cohorts and apply a non-linear production function in order to increase the likelihood of inducing *conflicting* subjective evaluations of claims.

[11] While one may reject a mechanism for the sole reason that it performs badly in terms of efficiency, one cannot recommend a mechanism in this context only because it is highly efficient. For instance, a mechanism that allocates the entire cake to one partner is maximally efficient, but it is probably considered rather unfair by most of the partners.

to the fairness comparison, we use three different yardsticks: First, before introducing the mechanisms, subjects are asked in a hypothetical fairness question what they consider a fair way to divide the cake. The answers to this fairness question shall serve as a first yardstick for the evaluation of the allocations produced by the three mechanisms. A second yardstick is subjects' evaluation of the mechanisms' procedural fairness, which is elicited after subjects have been exposed to the three division procedures but before they receive feedback about their payoffs. The third yardstick we use in our comparison is subjects' evaluation of the mechanisms' allocative fairness, which requires them to compare the actual outcomes of the three mechanisms.

Our results clearly show that the two mechanisms which explicitly force agents to assess the partners' claims and which do not provide obvious incentives to inflate the own claim – that is, the Impartial Division Rule and the Extended Divide-the-Dollar Rule – yield outcomes that are closer to subjects' fairness evaluations. The Modified Nash-Demand Rule that uses only agents' own claims as input is not only worse in terms of fairness, but due to its rule of imposing a fine when the sum of claims exceeds the available cake size, it is also far less efficient. The Average Division Rule is fully efficient by design, but performs rather badly in terms of procedural and allocative fairness, since it provides strong incentives to inflate the own claim and systematically disadvantages partners who try to implement a fair allocation.

2. Mechanisms to resolve the subjective claims problem

The subjective claims problem we consider features the three partners A, B, and C, who have jointly produced the cake S, which now has to be divided amongst them. We denote the subjective evaluation of claims by partner i by $c^i = (c^i_A, c^i_B, c^i_C)$, where c^i_j stands for the amount agent j should receive from partner i's perspective.[12] Throughout we assume that $c^i_A + c^i_B + c^i_C = S$ for $i = A, B, C$. Each of the mechanisms we consider yields an allocation $s = (s_A, s_B, s_C)$, where $\sum_i s_i \leq S$, which we will evaluate in terms of efficiency (in the sense that $S - \sum_i s_i$ is minimized) and in terms of allocative and procedural fairness. The mechanisms we compare differ in the amount and kind of information they process. Suppose the partners are asked to report their subjective evaluations of claims and denote the report of agent i by $m^i = (m^i_A, m^i_B, m^i_C)$, where m is mnemonic for message. (Note that if agent i reports truthfully then $m^i = c^i$). Then, from partner i's report, only i's own claim, m^i_i, is used as input in the first mechanism considered below; only others' claims, m^i_j and m^i_k, with $\{i, j, k\} = \{A, B, C\}$, are used in the second mechanism; and the whole vector of reported claims, m^i, is used in the third and fourth mechanism.

2.1. The Modified Nash-Demand Rule

Mumy (1981) proposed a mechanism where each agent i is asked to report only the own claim m^i_i and where the amount agent i receives is given by

$$s_i = m^i_i - \max\left\{ a\left(\sum_j m^j_j - S\right), 0\right\},$$

where $a > 1$. In words, if the sum of reported own claims does not exceed S, then each agent receives exactly the reported own claim. Otherwise each agent receives the reported own claim minus a 'fine' that is proportional to the difference between the sum of reported own claims and the available amount. Note that there are two sources of inefficiencies under this mechanism. One arises when the sum of reported own claims exceeds S, while the other arises when this sum falls short of S. Both kinds of inefficiencies increase linearly in the difference between the sum of reported own claims and the sum of money available. As for the theoretical prediction, it is far from trivial to see what it would be for the case where agents are not only interested in their own material payoff but also in the fairness of the final allocation as we assume in this paper, since it depends on the exact shape of agents' preferences and their belief about the partners' preferences in that case. In order to keep things tractable, we refer here and in the following only to the theoretical benchmark for the case where it is common knowledge that all agents are exclusively interested in their own monetary payoff.[13] Under this assumption, any combination of reports such that $\sum_j m^j_j = S$ is an equilibrium under this mechanism.

2.2. The Impartial Division Rule

In the mechanism proposed by DeClippel et al. (2008), each agent is asked to report an evaluation of the relative shares that the other two agents deserve. That is, for $\{i, j, k\} = \{A, B, C\}$ each agent i is asked only how much partner j should get

[12] A fully formulated theoretical model would start from the assumption that agents have complete and transitive preferences over all possible allocations, with the subjective evaluation of claims c^i being agent i's most preferred allocation. Although our approach implicitly assumes that agents have complete and transitive preferences over all possible allocations (otherwise the notion of a most preferred allocation would not be well defined), we introduce a notation only for the most preferred allocation because this is the object we elicit experimentally.

[13] The selfish benchmark is mentioned for completeness only. Our focus is not on whether subjects' behaviour in the experiment is consistent with this benchmark or not, but rather on which outcomes occur and how far they are from subjects' views of a fair division.

compared to partner k. Denoting partner i's report for the ratio of j's share to k's share by r^i_{jk} (so $r^i_{jk} = m^i_j/m^i_k$ in the notation used before), the Impartial Division Rule yields the following payoffs for the three partners:

$$s_A = \frac{S}{1 + r^B_{CA} + r^C_{BA}}; \quad s_B = \frac{S}{1 + r^A_{CB} + r^C_{AB}}; \quad s_C = \frac{S}{1 + r^A_{BC} + r^B_{AC}}.$$

As is easily seen, under this mechanism an agent's report has no impact on the own share – the latter is rather determined exclusively by the reports of the two partners. This is the impartiality property discussed earlier. Consensuality manifests itself in the property that if there is a division of S that is consistent with the reported opinions about relative shares, then the mechanism assigns these shares. In this case the whole cake is distributed. Otherwise the rule distributes strictly less than S, and the size of the inefficiency increases in the degree of inconsistency of the reported relative shares – see Tidemann and Plassmann (2008) for a discussion.[14] As for the theoretical prediction, the assumption of common knowledge about agents' exclusive interest in own monetary payoffs implies that any combination of reports constitutes an equilibrium for this rule, as no agent can affect his own share by his report.

2.3. The Extended Divide-the-Dollar Rule

The third mechanism we include in our comparison was proposed by Brams and Taylor (1994). The authors start with a setting with publicly known, player-specific entitlements $e = (e_A, e_B, e_C)$ that sum to the available amount S. The mechanism proposed for this 'objective entitlements' setup asks each agent i to report only the own claim m^i_i. If the sum of reported own claims does not exceed S, then each agent receives exactly the reported own claim; otherwise, the partners are given priority in order of their 'greed level' defined by $m^i_i - e_i$, starting with the lowest greed level.[15] Relevant for our subjective claims environment is an extended version of this mechanism, where players' entitlements are endogenously determined. For this purpose, each player i is now asked to report not only m^i_i, but rather a complete vector $m^i = (m^i_A, m^i_B, m^i_C)$. If own claims are feasible – in the sense that they sum up to S or less – they are again implemented. If own claims sum up to more than S, then this mechanism assigns to each player i a greed level defined as the difference between the reported own claim, m^i_i, and i's entitlement, which is the average of the claims assigned to i by her two partners, $(m^j_i + m^k_i)/2$. Own claims are then paid out sequentially in ascending order of greed levels up to the point where S is depleted. The only source for inefficiencies under this mechanism are own claims that sum up to less than the cake size S. While Brams and Taylor derive formal results for several other variants of the Divide-the-Dollar game, the discussion of this variant remains informal. Own investigations (under the assumption that agents care only about their own monetary payoff) reveal that the game induced by this rule has – in general – no pure-strategy equilibrium and that the existence and properties of mixed strategy equilibria heavily depend on the size of S and on the grid on S.

2.4. The Average Division Rule

One might argue that it is easy to find other mechanisms that might outperform the considered ones. A particularly interesting candidate (suggested by an anonymous referee) asks each agent i to report a complete vector $m^i = (m^i_A, m^i_B, m^i_C)$ and then implements for each partner the average reported share. That is, this mechanism yields

$$s_i = \frac{m^i_i + m^j_i + m^k_i}{3}$$

for partner i, where $\{i, j, k\} = \{A, B, C\}$. We refer to this mechanism as the 'Average Division Rule'. If agents are selfish, they claim the entire cake for themselves and the implemented allocation is an equal division of the dollar. Alternatively, they may express their fairness views in their proposed divisions. Note that while this mechanism is fully efficient by design, it lacks the properties of impartiality (respected by the Impartial Division Rule) and sanctioning greedy demands (the Extended Divide-the-Dollar Rule and the Modified Nash-Demand Rule have this property). This makes it well-suited for a comparison with the other three mechanisms, as we obtain information regarding the efficiency cost of imposing the requirement that a rule should not induce agents to exaggerate the own fair share. If agents report their subjective evaluations of claims truthfully, the mechanism yields an efficient outcome as a compromise of agents' reports.

3. Experimental design

The experiment was programmed and conducted with the software z-Tree Fischbacher (2007). The experiment consists of four parts.

Part I: Real effort task and emerging cake size. The real effort task consists of a general knowledge quiz. Prior to the quiz, subjects are randomly assigned to one of three same-size cohorts and they are informed that (i) each subject in a

[14] For more than three agents, DeClippel et al. (2008) characterize a family of adjusted rules that always distribute exactly the available amount.

[15] Ties result in an allocation where the shares of those involved in the tie are proportional to their entitlements.

Table 1
Group compositions in experiment.

Cake size	Points			# Observations	
S	A	B	C	Groups	Subjects
$S = 24$	2	2	3	90	270
$S = 36$	2	3	4	90	270
$S = 60$	3	4	4	90	270

cohort will be exposed to the same set of questions; (ii) each subject in a cohort will receive points depending on her relative performance (in terms of correctly answered quiz questions within a given time period) within her cohort; (iii) after the quiz each subject will be assigned to a group of three partners, each coming from a different cohort; (iv) the points a subject acquires in the quiz will be her contribution to the joint profit of the group; and (v) the joint profit of the group will later be distributed amongst group members with the help of several procedures, which are all fully paid out. Each cohort consists of 6 subjects, and the points assigned to subjects are as follows: The two high performers within a cohort (i.e., ranks 1 and 2) are assigned 4 points, the two medium performers (ranks 3 and 4) receive 3 points, and the two low performers (ranks 5 and 6) get 2 points. After the real effort task, subjects are informed about their own rank within their cohort and the points they achieved. Then they are assigned to a group consisting of three partners (labelled *A*, *B*, *C*), one from each of the three cohorts. Upon assignment to a group, subjects are informed about their two partners' contributions in points, but not about their partners' precise performance or rank within their respective cohort. The points subjects bring into the group enter a non-linear production function, which determines the size of the cake *S* to be distributed[16]:

$$S = 12 + (pointsA) \cdot (pointsB) \cdot (pointsC)$$

By using a relative performance measure *within* cohorts, but selecting the three partners from *different* cohorts, we intend to induce conflicting subjective evaluations of claims. A division according to the number of contributed points might be considered unfair because contributions in points are only a noisy signal for the actual performances in the quiz and because subjects have no possibility to directly compare their own quiz performance to that of the other two group members. The non-linear production function further aggravates the difficulty of finding a fair division.

Groups in the experiment are composed such that we have groups with a small cake size of $S = 24$, groups with a medium cake size of $S = 36$ and groups with a large cake size of $S = 60$ – see Table 1 for details. Note that with this choice of group compositions we have groups where two partners contribute the same low number of points, groups in which all partners have different contributions in points and groups where two partners have the same high contribution.

Part II: The fairness question. After being informed about their partners' contributions in points and the resulting cake size, subjects are privately asked what they consider a fair division of the jointly produced cake. That is, each subject *i* is asked to report a vector of his subjective evaluation of claims, $m^i = (m^i_A, m^i_B, m^i_C)$, where the entries have to sum up to *S*, knowing that the answer to this question is irrelevant for the earnings in the experiment. The answers to the fairness question shall serve as one of the benchmarks for our comparison of the cake-division mechanisms in terms of allocative fairness.

Part III: Actual division of the cake. In each experimental session, each group is successively exposed to each of the three main mechanisms in random order.[17] For the Modified Nash-Demand Rule, the factor *a* determining the fine in case reported own claims exceed *S* is set to 1.1. Group composition and cake size remain constant across mechanisms. In each case, the mechanism is first described and subjects have to answer some control questions to make sure that they understand how the mechanism works. Then subjects are asked for the necessary input for the mechanism, after which they receive the description of the next mechanism. Feedback about the allocation implemented by a mechanism is given only at the end of the experiment (see Part IV below) in order to avoid that the outcome of a mechanism affects subjects' behaviour for other mechanisms. All mechanisms are paid off, and for each point earned in the experiment subjects were paid 25 cents.

Part IV: Procedural and allocative fairness. Before subjects receive feedback about their payoffs under the different cake-division mechanisms, they are asked to rank the mechanisms in terms of procedural fairness. After receiving information about their payoffs, subjects are asked to rank the allocations implemented by the mechanisms in terms of allocative fairness.

4. Experimental results

4.1. Fairness question

One might argue that in a subjective claims problem known fairness standards (apart from the egalitarian norm) should not play a role for fairness evaluations, since contributions are at least partly influences by luck, which is often considered

[16] All of this information was contained in the instructions, which are displayed in Appendix B (not intended for publication).

[17] As already mentioned, in some sessions subjects were successively exposed to four mechanisms, the three main ones and the Average Division Rule. Note also that in addition to the three or four mechanisms subjects in all sessions also played a unanimity bargaining procedure. Different bargaining procedures were tested in different sessions, and the bargaining procedure was always presented after the static mechanisms.

Table 2
Fairness standards and observed assignments in fairness question.

Cake size		S = 24		S = 36			S = 60	
Partner		A/B	C	A	B	C	A	B/C
Contribution		2	3	2	3	4	3	4
Fairness standard								
Egalitarian	Payoffs	8	8	12	12	12	20	20
	obs. in %	42.7	18.8	28.8	15.5	8.8	41.1	25.0
Proportional	Payoffs	6.86	10.28	8	12	16	16.36	21.82
	obs. in %	38.3	62.2	26.6	26.6	31.1	15.5	27.2
Liberal	Payoffs	7.43	9.14	9.33	12	14.67	17.10	21.45
	obs. in %	5.5	2.2	4.4	5.5	8.8	27.2	9.4
Observed assignments in fairness question								
Own fair share		8.04	9.90	10.22	12.64	16.36	18.91	22.65
Fair share from others		7.14	8.72	8.61	11.48	14.30	16.00	20.95
Average fair share		7.44	9.10	9.15	11.87	14.99	16.97	21.51

a factor that is irrelevant for fairness perceptions because it is outside the agent's control. Yet, due to the absence of other (better) benchmarks, it is likely that some known fairness standards serve as anchors for subjective evaluations of claims.[18] The upper part of Table 2 displays the allocations implied by three well-known division standards and the relative frequencies of answers to the fairness question that are consistent with each standard.[19] We refer to the *egalitarian standard* when S is distributed equally among the partners, to the *proportional standard* when shares of S are assigned proportionally to the points each partner has contributed, and to the *liberal standard* when each partner receives an equal share of the fixed part of the production function and the remainder of S is divided proportionally to the points contributed. Table 2 shows that the overall consistency of subjects' answers to the fairness question with common fairness standards is rather low. In particular, for the cake size of $S = 36$ – where all contributions are different – only about half of the observed assignments of fair shares are consistent with one of the standards considered. This confirms our presumption that evaluations of claims are subjective in our scenario, i.e. it is not clear which standard is of relevance here.

The lower part of Table 2 displays for each contribution type and cake size how much subjects, on average, report as being fair for themselves (*'own fair share'*), what their partners report as being fair for them (*'fair share from others'*), and what the average fair share for each position is when all three partners' reported fairness evaluations are included in the calculation of the average (*'average fair share'*). In calculating these figures we pool the data for the two partners who contribute the same number of points within a given cake size, since there is no sound basis to differentiate between them (except for the label), and since the two-sample Kolmogorov–Smirnov test for equality of distributions shows no significant differences. A first lesson we learn from Table 2 is that there is a significant self-serving bias in the answers to the fairness question: The own fair share is significantly larger than the fair share from others in all relevant comparisons (paired t-test: $p < 0.01$ for all contribution types and cake sizes). Own fair shares also exceed the shares the proportional standard would predict for almost all types and cake sizes for common significance levels. There are two exceptions: Those subjects who bring the most points to the partnership report own fair shares below the proportional share in $S = 24$ (t-test: $p < 0.02$), and own fair shares that are not significantly different from what the proportional standard would predict in $S = 60$ (t-test: $p = 0.19$). This observation is consistent with a taste for a more 'compressed' distribution than the one implied by proportionality, i.e., one that implies smaller payoff differences across subjects. This taste for a more compressed distribution can be seen more clearly in the assignments from others, where the self-serving bias does not obliterate part of the effect: Fair shares from others are significantly different from the proportional standard in most comparisons (t-test: $p < 0.01$ for all types and cake sizes except for the medium contributor in $S = 60$, where $p < 0.09$), and while high and medium contributors are consistently assigned less than what proportionality would predict, low contributors are assigned more.

Looking at individual data, the self-serving bias manifests itself in the choice of fairness standards – while low contributors tend to report own fair shares that are (roughly) consistent with the egalitarian standard, high contributors tend to report fair shares consistent with the proportional standard.[20],[21]

[18] The survey by Karagözoglu (2012) shows how norms of equity and desert influence bargaining behaviour when agents jointly produce the cake that is to be divided. Regarding the self-serving bias in fairness judgements, which is shown in numerous studies when stakes are involved (for field experiments see Babcock et al., 1996, for lab experiments see Kagel et al., 1996; Konow, 2000), Konow (2003) notes that "although biases sometimes widen the range of predicted outcomes, behaviour still is constrained by fairness".

[19] When counting observations as consistent with a given standard, we allow for intervals that typically round numbers to the next half unit in case the standard does not yield integers. See Appendix A for more details.

[20] As Dana et al. (2007) show, the existence of "moral wiggle room" allows subjects to behave self-interestedly while maintaining the illusion of fairness in the presence of uncertainty between actions and their resulting outcomes. According to Cappelen et al. (2007), an application of this idea in the context of multiple fairness norms is that subjects tend to appeal to the one that benefits them most.

[21] Table A.1 in Appendix A displays a related observation: Low contributors frequently report fair shares that assign more to low contributors than what proportionality predicts, while higher contributors do so to a lesser extent. It is important to note, however, that a considerable fraction of *all* contribution types in all cake sizes assign a larger share than the one predicted by the proportional standard to the lowest contributor within their group. This confirms our earlier finding of a taste for a more compressed distribution than the one implied by proportionality.

Table 3
Modified Nash-Demand Rule: claims, fines and payoffs.

Cake size		$S=24$		$S=36$			$S=60$	
Partner		A/B	C	A	B	C	A	B/C
Contribution		2	3	2	3	4	3	4
Claim	Mean m_i^i	9.06	9.50	11.02	14.21	16.00	19.18	22.54
	Median m_i^i	8	9	10	12	15	18.75	21
Payoff	Mean s_i	5.47	6.12	6.40	8.75	10.75	14.24	16.98
	Mean s_i/m_i^i	0.70	0.63	0.66	0.70	0.68	0.78	0.77
Fine		4.43			6.53		5.8	
Efficiency (in %)		0.71			0.72		0.80	
When $\sum m_i^i > S$		0.53			0.60		0.72	
When $\sum m_i^i < S$		0.93			0.92		0.93	

Result 1 (Fairness Assessments). *Overall, the consistency of subjects' answers to the fairness question with common fairness standards is rather low. A self-serving bias is present for all contribution types and cake sizes, and subjects tend to assign more to the low contributor and less to the high contributor compared to the distribution implied by proportionality.*

4.2. Results for the mechanisms

Since the randomized order in which the procedures were presented to subjects showed no effect on the results for common significance levels, we will report pooled data of all sessions for a given procedure.

4.2.1. Modified Nash-Demand Rule

The results for the mechanism proposed by Mumy (1981) are displayed in Table 3. Again, we use pooled data for partners with the same contribution in points, as the Kolmogorov–Smirnov test shows no difference in the two distributions. Recall that only each partner i's own claim, m_i^i, is elicited by this mechanism.

Mechanism input: Reported own claims respect the ordering of contributions in points (median test: $p < 0.001$ for all comparisons of claims with different contributions), but they do not correspond to the proportional standard: While low contributors ask systematically more than proportionality would predict (WSR and t-test: $p < 0.001$ for all cake sizes), high contributors ask less (WSR: $p < 0.03$ for all cake sizes, only the t-test for $S = 36$ gives $p = 0.5$). This corresponds to the preference for a compressed distribution observed in the fairness question. For high contributors there is no difference between reported own claim and stated own fair share independent of the cake size. For low contributors, however, reported own claims tend to be higher than own fair shares in the fairness question (WSR: $p < 0.03$ and paired t-test: $p < 0.001$ for $S = 24$; t-test: $p < 0.08$ while WSR: $p = 0.28$ for $S = 36$), while for medium contributors the sign of this difference depends on the relative position they take within a cake size. Since stated own fair shares were already found to be self-servingly biased, we can expect from these mechanism inputs that the Modified Nash-Demand Rule will suffer from excess claims.

Mechanism output: For all cake sizes, mean payoffs are significantly lower than mean claims, as can be seen in Table 3, and this is true also for pairwise comparison on the individual level (MWU: $p < 0.01$ for all types and cake sizes). Payoffs are thus also lower than what subjects considered their own fair share in the fairness question, and furthermore, payoffs from this mechanism are also lower than fair shares assigned from others in the fairness question (WSR: $p < 0.001$ for all types and cake sizes). The efficiency of this mechanism, i.e. the share of S that is paid out, is rather low – it reaches values between 71% for the small cake size and 80% for the large cake size. While there are two sources of inefficiency, claiming too little or too much, the latter turns out to be the main culprit for the low payoffs: The sum of claims is lower than the available amount in less than 25% of cases for all cake sizes, while it is higher in more than 50%. In addition, efficiency is considerably lower with overclaiming compared to underclaiming (see lower part of Table 3).[22]

Overall, our results suggest that the self-serving bias in fairness evaluations already present in the fairness question is exacerbated by this mechanism: Subjects tend to report own claims that are higher than own fair shares in the fairness question, and this results in realized payoffs that are not only lower than reported own claims, but also lower than the stated own fair shares and even the fair shares from others in the fairness question.

Result 2 (Modified Nash-Demand Rule). *The self-serving bias found in the fairness question is exacerbated by the Modified Nash-Demand Rule: Subjects tend to report own claims that are at least as high as stated own fair shares. This results in an excess sum of own claims, which, in turn, leads to considerable inefficiencies. The second source for inefficiencies – claims that sum to less than the cake size – is empirically less relevant for all cake sizes.*

[22] Note that to avoid negative payoffs in the experiment, we set a subject's payoff to zero if $m_i^i - a(\sum_j m_j^j - S) < 0$, thus efficiency from overclaiming can be larger than what the difference between mean sum of claims and total fines may suggest.

Table 4
Impartial Division Rule: reports and payoffs compared to fairness standards.

Cake size		$S = 24$		$S = 36$			$S = 60$	
Partner		A/B	C	A	B	C	A	B/C
Contribution		2	3	2	3	4	3	4
Observed r^i_{jk}	Mean	0.86	0.99	0.80	0.48	0.66	1.09	0.75
	Median	0.67	1	0.75	0.5	0.67	1	0.75
Observed fair ratio		0.84	1.00	0.83	0.60	0.71	1.0	0.76
Proportional		0.67	1	0.75	0.5	0.67	1	0.75
Realized s_{jk}		0.84	1.00	0.80	0.51	0.62	1.06	0.75
Efficiency		0.98			0.99		0.99	

Agent *i*'s report r^i_{jk} is displayed in column *i*, and s_{jk} is the ratio of *j*'s to *k*'s realized payoff share, displayed in column *i*.

4.2.2. Impartial Division Rule

Recall that this mechanism asks each agent *i* only for an evaluation of the partners' relative shares, $r^i_{jk} = m^i_j / m^i_k$ for $\{i, j, k\} = \{A, B, C\}$, which makes it 'impartial'. While there exist impartial and efficient rules for dividing a cake for any number $n > 3$ of agents, in the special case of three agents the rule distributes the entire cake only if the reported relative shares are consistent; otherwise it distributes strictly less.[23] Thus, besides the main question how well this mechanism performs in terms of procedural and allocative fairness, it is also interesting to see whether efficiency is an issue in the case of three agents.

Mechanism input: Table 4 displays the elicited relative shares r^i_{jk} as well as the respective ratios implied by subjects' answers to the fairness question (*observed fair ratio*) and the ratios implied by the proportional standard. Subjects' reports for r^i_{jk} reveal a consensus on assigning equal shares for equal contributions (87% in $S = 60$ and 94% in $S = 24$ assign ratio 1 : 1 when partners have the same contribution). With different contributions, the median of r^i_{jk} corresponds precisely to the ratio implied by the proportional standard for each role in each cake size, and for medium and high contributors the mean of r^i_{jk} is also not significantly different from the ratio derived from the proportional standard. Only low contributors assign less to higher contributors than what their relative payoff according to the proportional standard would be, both in the medium and in the large cake size (WSR: $p < 0.01$).[24] Overall, in the small and large cake size, the observed r^i_{jk} does not differ significantly from the observed fair ratio for any contribution type. This means that subjects' reports in the Impartial Division Mechanism do not deviate substantially from their evaluations of claims stated in the fairness question, which we interpret as their true subjective evaluations of claims. Only for the cake size of $S = 36$, the medium and the high contributors assign more to the high and less to the low contributors compared to what they stated in the observed fair ratio (both WSR and *t*-test are significant at $p < 0.05$), and are thus more consistent with the proportional standard. Low contributors here also tend to reduce payoff differences, which again reflects their *observed fair ratio*.

Mechanism output: As Table 4 shows, realized payoff ratios under this mechanism are rather close to reported input ratios (the ratio of *j*'s to *k*'s payoff share, s_j/s_k, is displayed in partner *i*'s cell of the respective cake size): The WSR cannot reject that the mechanism implements results that are not significantly different from subjects' reports for all cake sizes. As a result, the comparisons above between mechanism inputs and proportional standard also hold for realized payoffs. Regarding the comparison between realized payoff ratios and observed fair ratios in the fairness question, the mechanism yields good results: There are no significant differences for the large and small cake size; only for the medium cake size, low contributors are disadvantaged compared to the observed fair ratios in the fairness question.

Comparing the absolute size of realized payoffs to the respective assignments from the fairness question (see Fig. A.1 in Appendix A), we find that the Impartial Division Rule implements absolute payoffs that are between the own fair share and the fair share from others in the fairness question. This seems justified, considering the self-serving bias in the answers to the fairness question. A mechanism that performs well in a subjective claims problem needs to take such a bias into account. Furthermore, the Impartial Division Rule is highly efficient, paying out over 98% of all cake sizes. Asking whether deviations from the *observed fair ratio* would bring about changes in the mechanism's efficiency, we find that this is not the case for the medium and large cake size. In the small cake size, there is some efficiency loss due to the deviation of actual reports from those elicited in the fairness question (WSR: $p < 0.05$), the difference in means is only 0.006%, though. Overall, this mechanism gives good incentives to report true evaluations of others – assuming that those elicited in the fairness question are true. This and the property that it avoids self-evaluations seem to be the main strengths of the mechanism. A possible weakness is the fairly intransparent payoff rule which does not easily reveal the consequences of changing the own report on assigned shares.

---

[23] The reports of the three partners are consistent iff $r^C_{AB} = r^B_{AC} / r^A_{BC}$.

[24] Two outliers of $r^B_{AC} = 50$ are excluded from reported descriptive statistics for $S = 24$ in Table 4. Including them would yield a mean reported ratio of 1.4 for the low contributors and give a distorted picture of the results.

Table 5
Extended Divide-the-Dollar Rule: reports and payoffs compared to fairness question results.

Cake size	$S=24$		$S=36$			$S=60$	
Partner	A/B	C	A	B	C	A	B/C
Contribution	2	3	2	3	4	3	4
Own claim	8.49	9.8	11.06	12.77	15.22	19.30	22.15
Entitlement	7.08	8.4	9.09	11.83	13.55	17.21	20.49
Greed level	1.42	1.41	1.98	0.95	1.67	2.08	1.66
Own fair share	8.04	9.90	10.22	12.64	16.36	18.91	22.65
Fair share from others	7.14	8.72	8.61	11.48	14.30	16.00	20.95
Payoff	7.31	8.95	9.25	12.07	13.84	16.93	21.05
Efficiency	0.98			0.97		0.98	

Result 3 (Impartial Division Rule). *By eliciting relative shares of the partners, the Impartial Division Rule avoids the self-serving bias, and reported shares are mostly in line with the proportional standard and also with the ratios of subjects' stated fair shares for others. Realized absolute payoffs are between the own fair share and the fair share from others, thus representing elicited subjective claims from the fairness question. Furthermore, the mechanism is highly efficient implying that there is hardly an efficiency cost of requiring impartiality.*

4.2.3. Extended Divide-the-Dollar Rule

In the mechanism proposed by Brams and Taylor, each partner i is asked to report an $m^i = (m^i_A, m^i_B, m^i_C)$. This makes it particularly well-suited for a comparison with the results from the fairness question, as they both ask for precisely the same information. Table 5 displays own claims, entitlements, greed levels and payoffs for this mechanism, as well as the corresponding reports from the fairness question as comparison.[25]

Mechanism input: Subjects' own claims differ significantly from their entitlements (WSR: $p < 0.01$ for each comparison), which implies a greed level different from zero for all contribution types and cake sizes (see Appendix A for more details). Since partners are paid out their claims in ascending order of their greed level and the greed level is calculated as the difference between own claim and own entitlement, the mechanism incorporates an incentive to act strategically. On the one hand, an agent might report a low claim for a partner, which decreases his entitlement and thus increases his greed – holding everything else constant. This makes it more likely that one's claim is fully paid off, as one's own greed is comparatively lower when that of others increases, and partners with lower greed levels get paid off first. If, on the other hand, a subject reduces his own claim, this increases the chance of being fully paid off (since the own greed level decreases), however, at the cost of receiving a lower amount when the subject gets paid off (since payoffs correspond at most to own claims).

In Appendix A we report on observed behaviour that is consistent with any of the strategic incentives described. Overall, we find that for each cake size, the partner with the highest contribution in a given partnership gets systematically disadvantaged by a lower entitlement (WSR: $p < 0.01$ and t-test: $p < 0.01$ for type C in $S = 36$; WSR: $p < 0.07$ and t-test: $p < 0.03$ for type C in $S = 24$; and WSR: $p < 0.03$ and t-test: $p < 0.01$ for pooled types B and C in $S = 60$). Lower contributors, on the other hand, receive entitlements at least as high as the fair shares from others in the fairness question. These tendencies are fully in line with the strategic incentives described earlier: Shifting the assignment from high to low contributors increases the chance that the latter advance in the order partners are paid off, but they take away a smaller share of the cake than anyone else would. Therefore, reducing the entitlements of higher contributors and increasing those of lower contributors serves the same purpose.

Mechanism output: As can be seen in Table 5, final payoffs in this mechanism are always smaller than own claims (WSR: $p < 0.01$ for all types and cake sizes). This is due to the observed self-serving bias in own claims and its consequence that not all claims are paid off when the sum of own claims exceeds the total amount available. Only for medium contributors in $S = 36$ this difference is small, as 82 out of 90 subjects receive the full own claim. The high own claims also imply that most of the time the full cake size is allocated, which can be seen in the high efficiency of this mechanism: for all cake sizes, 97% or more of the available amount is paid off.

Regarding the fairness of the implemented allocation, we first note that the payoffs of partners with the same contribution level within a partnership do not differ (MWU: $p > 0.6$ for low contributors in $S = 24$, MWU: $p > 0.3$ for high contributors in $S = 60$). For all contributions in all cake sizes, realized payoffs are below own fair shares as stated in the fairness question (WSR: $p < 0.03$ for all comparisons). Comparing each type's payoff to the fair share from others in the fairness question, we find that all types in the small cake size get a slightly higher average payoff than what others think is fair, but the difference is not significant. For $S = 36$, the low and medium contributor get a significantly higher payoff than the fair share from others (t-test: $p < 0.01$); the high contributor's mean payoff is smaller, but the difference is not significant (WSR: $p = 0.17$ and t-test: $p = 0.11$). For $S = 60$, the medium contributor gets significantly more (t-test: $p < 0.01$), while the two high contributors receive about what others thought was fair for them.

[25] Remember that the 'entitlement' of partner i is the average of the two partners' evaluations of i's claim.

Table 6
Average Rule: reports and payoffs compared to fairness question.

Cake size	S = 24		S = 36			S = 60	
Partner	A/B	C	A	B	C	A	B/C
Contribution	2	3	2	3	4	3	4
Own claim	14.21	16.31	15.93	21.18	21.43	36.37	38.52
Assignment from others	4.24	5.12	5.43	9.13	10.12	9.04	12.12
Payoff	7.57	8.85	8.93	13.16	13.89	18.15	20.92
Own fair share	8.40	10.31	9.56	13.68	17.43	19.02	22.84
Average fair share	7.57	8.87	8.43	12.12	15.43	16.10	21.96

Overall, we see that high contributors' disadvantage due to lower entitlements and lower own claims carries over to their payoffs in this mechanism – they are the only ones who systematically receive an amount comparable to what others think is fair for them, while others receive more. Fig. A.2 in Appendix A displays by how much realized payoffs differ from own fair shares as well as from fair shares from others as reported in the fairness question. Comparing realized payoffs to the discussed fairness standards, we find that all types in all cake sizes receive a payoff that is significantly different from what the egalitarian and proportional standard would imply. Unique exception is the medium contributor in $S = 36$ whose payoff is not significantly different from 12 (which, as already discussed, also corresponds to the payoff implied by other distribution standards).

Result 4 (Extended Divide-the-Dollar Rule). *Subjects claim more for themselves than what their partners think they are entitled to under the Extended Divide-the-Dollar Rule. This makes the mechanism highly efficient. Entitlements assigned by the partners are systematically lower for high contributors and higher for low contributors compared to the fair shares from others. This distortion is carried over to payoffs, leading to a disadvantage of high contributors.*

4.2.4. Average Division Rule

Our three main mechanisms have in common that they either have a built-in punishment for overly greedy demands (the Modified Nash-Demand Rule and the Extended Divide-the Dollar Rule have this property) or are impartial by design (as the Impartial Division Rule). A possible downside of each of them is that they yield efficiency losses in case of disagreement. The Average Division Rule is located at the other end of the spectrum. It has neither a punishment for greedy demands nor is it impartial, but it is fully efficient by design. It is therefore interesting to see how the Average Division Rule compares to the other mechanisms. To address this question, we tested this rule in some additional sessions with 144 subjects in total.

Table 6 displays the results of the Average Division Rule and how they compare to those from the fairness question. As in the main experiment, the results from the fairness questions display a self-serving bias for all contribution types, as can be seen in the last two rows of the table. Turning to mechanism inputs, we observe that only about one quarter of our subjects make entirely selfish claims. Claims are nevertheless extremely biased towards material self-interest: Even if we exclude all subjects with entirely selfish claims, the remaining subjects' own demands are significantly larger than own fair shares as stated in the fairness questions (WSR test significant at common levels for all contribution types and cake sizes).

Compared to the reported ratios under the Impartial Division Rule, subjects' assignments to the two partners is dampened in the Average Division Rule in the sense that when partners contribute a different amount of points to the cake size, the lower contributor receives more and the higher contributor receives less than under the Impartial Division Rule. This may be due to the fact that subjects' focus under the Average Division Rule is largely on own demands, and less emphasis is put on a fair evaluation of the partners. Compared to the Extended Divide-the-Dollar Rule, the shares assigned by an agent to the two partners are significantly smaller under the Average Division Rule (WSR test significant at common levels for all comparisons). This might be due to the fact that the Extended Divide-the-Dollar Rule reduces the self-serving bias through greed-based sanctions, while the Average Division Rule does not sanction selfish claims. On the other hand, observed claims under the Average Division Rule are not sufficiently selfish so that the average of all three claims would induce an equal division. The result is a distorted division with large shares for subjects who act selfishly, and small shares for those who display some concern for others in their reported assignment. This also implies a considerably larger variance in the payoffs compared to the other mechanisms (see Fig. A.3 in Appendix A). If the average fair share is used as a yardstick for a fair outcome, we find that payoffs from the Average Division Rule display a larger absolute deviation from this benchmark than those from the Impartial Division Rule and the Extended Divide-the-Dollar Rule (WSR: $p < 0.01$ for each contribution type).

Result 5 (Average Division Rule). *A simple mechanism in which each partner reports a division and the average is implemented leads to a distorted division with large shares for subjects who make very selfish claims and small shares for subjects who include their fairness ideas into their reported division.*

4.3. Comparison of procedural and allocative fairness

In this subsection we first compare the three main mechanisms in terms of subjects' ex-post evaluation of the procedural and allocative fairness. Then we discuss the results from the sessions where all four mechanisms have been tested.

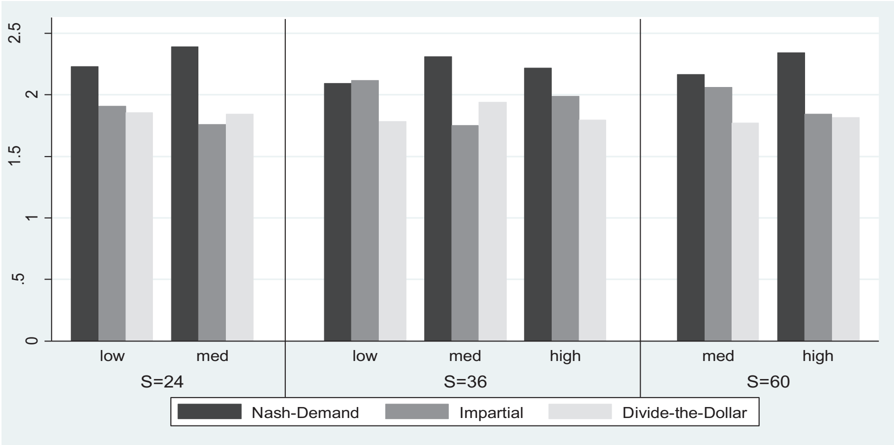

Fig. 1. Mean ranks of mechanisms' procedural fairness by cake size.

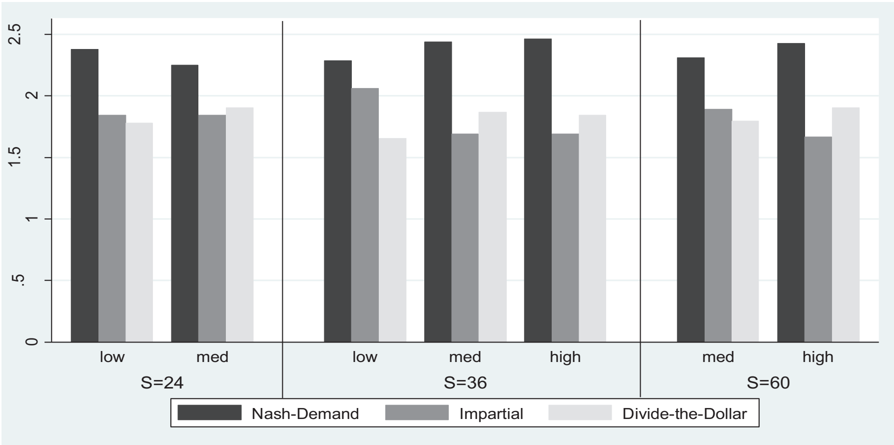

Fig. 2. Mean ranks of mechanisms' allocative fairness by cake size.

4.3.1. Comparison of main mechanisms

Before knowing their payoffs, subjects were asked to rank the mechanisms in terms of procedural fairness. Fig. 1 shows the mean rank that each contribution type in a given cake size assigns to each mechanism for its procedural fairness: The Modified Nash-Demand Rule is ranked worse than the others in most cases, while there is mostly no significant difference in the mean rank between the Impartial Division Rule and the Extended Divide-the-Dollar Rule. Indicating the significance level of 5% for a Wilcoxon signed-rank test by ** and the level of 1% by ***, the following relations hold: Extended Divide-the-Dollar Rule = Impartial Division Rule $\succ^{***}$ Modified Nash-Demand Rule for all contribution types and all cake sizes except for the medium contributor in the large cake size and the low contributor in the medium cake size, for whom we have Extended Divide-the-Dollar Rule $\succ^{**}$ Impartial Division Rule = Modified Nash-Demand Rule. Thus, when considering only the three main mechanisms, the Extended Divide-the-Dollar Rule is seen as best with respect to procedural fairness by subjects when they do not know which allocations the rules finally implement.

After being informed about their payoff from each mechanism, subjects were asked to rank the mechanisms according to the fairness of outcomes. Fig. 2 displays the results of subjects' view on the allocative fairness of the three main mechanisms. Here we see that the Modified Nash-Demand Rule again is ranked last for all comparisons while the preference in the comparison Impartial Division Rule versus Extended Divide-the-Dollar Rule depends on subjects' contribution type. The following relations hold: Extended Divide-the-Dollar Rule = Impartial Division Rule $\succ^{***}$ Modified Nash-Demand Rule for all types in $S = 24$, for the high and medium contributor in $S = 36$ and for the medium contributor in $S = 60$. For the high contributor in the large cake size we find Impartial Division Rule $\succ^{***}$ Extended Divide-the-Dollar Rule $\succ^{***}$ Modified Nash-Demand Rule, while for the low contributor in $S = 36$ we find Extended Divide-the-Dollar Rule $\succ^{***}$ Impartial Division Rule = Modified Nash-Demand Rule. The Impartial Division Rule might have looked slightly less appealing when judged only in terms of procedural fairness as one cannot state own claims, and only relative evaluations of others are possible. However, after seeing the allocation this mechanism produces, it moves upwards in the ranking of agents in the large cake size, while for all other types its ranking does not change. The Extended Divide-the-Dollar Rule probably has more intuitive appeal, but its potential for strategic behaviour, which harms high contributors in our experiment, makes its allocations less fair for these partners in the end.

4.3.2. Results from sessions testing all mechanisms

Looking at the data of the sessions where subjects were exposed to four mechanisms – the three main mechanisms and the Average Division Rule – we observe that the relative ranking of the three main mechanisms is very similar to that in the main comparison above – see Fig. A.4 in the Appendix for details. Regarding the Average Division Rule we observe that it performs rather poorly in terms of procedural fairness for the low and the medium contributor. The mechanism's rating then decreases further for all contribution types after subjects are informed about the outcomes and evaluate the allocative fairness of the mechanisms: The Average Division Rule fares worse than both the Impartial Division Rule and the Extended Divide-the-Dollar Rule in terms of distributional fairness for all contribution types and all cakes sizes (WSR: $p < 0.01$). It is ranked higher than the Modified Nash-Demand Rule, which remains the worst performer in all comparisons.

Result 6 (Procedural and Allocative Fairness). *While the Extended Divide-the-Dollar Rule ranks slightly better in terms of procedural fairness, the Impartial Division Rule gains when considering allocative fairness. The Modified Nash-Demand Rule is ranked consistently worse than the others in both comparisons. In sessions where the Average Division Rule was included in the comparison, it ranks worse than both the Impartial Division Rule and the Extended Divide-the-Dollar Rule, but better than the Modified Nash-Demand Rule with respect to both fairness measures.*

5. Conclusion

Different cake-division mechanisms have been compared experimentally in terms of fairness and efficiency in a subjective claims problem involving three agents. All mechanisms use reported subjective evaluations of claims as input and yield unique shares summing up to no more than the value of the cake as output. The mechanisms differ in the amount of information they process: The Modified Nash-Demand Rule (Mumy, 1981) only requires each partner's own claim as input, the Impartial Division Rule (DeClippel et al., 2008) asks for each agent's assessment of her partners' relative claims, and the Extended Divide-the-Dollar Rule (Brams and Taylor, 1994) requires each partner's assessment of her own as well as the partners' claims. While under the Impartial Division Rule no partner can affect his own share regardless of what he reports, under the other two mechanisms the reported own claim potentially has an impact on the own share. To avoid excessively greedy reported own claims, these mechanisms punish agents when the sum of reported own claims exceeds the available cake size. The Average Division Rule tested in some additional sessions uses the same input as the Extended Divide-the-Dollar Rule but has no built-in punishment for overly selfish claims. On the positive side, this mechanism is efficient by design.

We found that the Modified Nash-Demand Rule performs rather badly, since the considerable self-serving bias – i.e. the difference between what subjects consider fair for themselves and what others think is their fair share – implies that not the entire amount available is paid out. This is mainly due to the mechanism's built-in fine that is intended to deter agents from reporting too high own claims. The Impartial Division Rule and the Extended Divide-the-Dollar Rule perform significantly better in terms of efficiency, as at most 2-3% of the cake is not allocated. They also rank high in terms of perceived procedural and allocative fairness. The relative ranking of these two mechanisms depends on the relative position of the evaluating partner in terms of contribution to the cake. While low contributors typically prefer the Extended Divide-the-Dollar Rule, high contributors tend to prefer the Impartial Division Rule. This is consistent with our finding that low contributors are systematically advantaged and high contributors disadvantaged by the Extended Divide-the-Dollar Rule. The Average Division Rule tested in some additional sessions is efficient by design. However, it ranks low in terms of procedural and allocative fairness, mainly because subjects who try to implement their fairness ideas are systematically disadvantaged.

Overall, we think that the Impartial Division Rule in particular deserves high attention in this type of division problem for three reasons: First, no partner is systematically disadvantaged by this mechanism; second, its property of impartiality (own shares are only determined by others' reports, and one can only report evaluations of others) and the relative (rather than absolute) evaluation it requires make it less susceptible to the self-serving bias; and third, the mechanism's high efficiency implies that the efficiency cost of impartiality is rather low.

The ubiquitous self-serving bias also explains the findings of our companion study (Gantner et al., 2013) on the performance of three unanimity bargaining procedures to resolve the same subjective claims problem: A procedure that only requires agents to state their own demand has a disadvantage for the last mover, as little is left for him after the other two players make their demands. Bargaining procedures which force agents to consider the entire vector of claims turn out to be closer to subjects' own fairness assessments, and also to assessments of impartial spectators expressed in a vignette.

While bargaining may seem a natural way to resolve the subjective claims problem, our companion paper shows that all bargaining protocols considered lead to inefficiencies. By contrast, as the present paper has shown, a carefully chosen static mechanism implements a fair allocation at almost no efficiency loss. The fact that a well designed static mechanism – such as the Impartial Division Rule – outperforms the best bargaining protocol in our comparison in terms of effciency is a remarkable result, especially in light of the fact that our experiments involve only three players. Indeed, we would expect the relative advantage of static over dynamic procedures to increase in the number of players.[26] Since the high-stake environments in the motivating examples typically involve at least three parties, and since in these examples each percentage point of difference in efficiency corresponds to a large dollar amount, investigating how the relative performance of static over dynamic procedures changes in the number of players seems a fruitful area for future research.

Appendix A.

A.1. Additional results

Fairness question. The table below presents the results of the fairness question in more detail. When counting observations as consistent with a given standard, we allow for intervals that typically round numbers to the next half unit in case the standard does not yield integers. [27] These intervals are given in the *'adjusted'* row. Furthermore, we report the observed relative frequencies of some intuitive criteria such as assigning equal shares to equal contributions, or shares that reflect the contribution ordering.[28]

Table A.1
More fairness standards and observed assignments.

Cake size	$S=24$		$S=36$			$S=60$	
Partner Contribution	A/B 2	C 3	A 2	B 3	C 4	A 3	B/C 4
Fairness standard							
Egalitarian	8	8	12	12	12	20	20
Proportional	6.86	10.28	8	12	16	16.36	21.82
Adjusted	[6.5, 7]	[10, 11]	8	12	16	[16, 17)	(21.5, 22]
Liberal	7.43	9.14	9.33	12	14.67	17.10	21.45
Adjusted	(7, 7.5]	[9, 9.5]	[9, 9.5]	12	[14.5, 15]	[17, 18]	[21, 21.5]
Observations consistent with fairness standard (in %)							
Equal shares to equal contributions	90.5	96.6		–		98.8	89.4
Reflects contribution ordering	47.2	38.8	62.2	75.5	80.0	53.3	64.4
Low contributor gets more than proportional	50.5	23.3	66.6	53.3	41.1	73.3	38.3

Impartial Division Rule. Fig. A.1 compares the absolute size of realized payoffs from the Impartial Division Rule to the

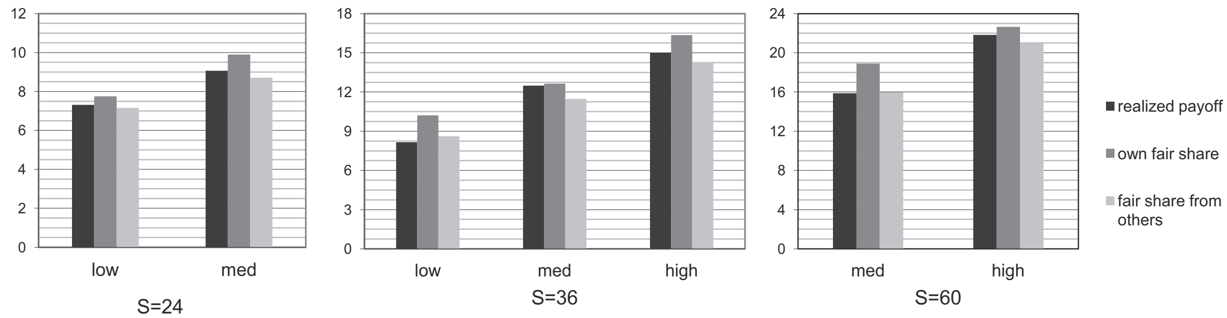

Fig. A.1. Fairness and impartial division mechanism.

respective assignments from the fairness question. This mechanism implements absolute payoffs that are between the own fair share and the fair share from others in the fairness question.

Extended Divide-the-Dollar Rule. Fig. A.2 compares the absolute size of realized payoffs from the Extended Divide-the-Dollar Rule to the respective assignments from the fairness question.

[26] See e.g. Chatterjee and Sabourian (2000), and Cai (2000) for theoretical work, or Cadigan et al. (2009) for an experimental study on the so-called holdout problem, a situation in which the required agreement by multiple parties creates an incentive for individual strategic delay in order to capture a larger share of the surplus.

[27] Sometimes the interval is larger. For instance, for $S=24$ the proportional standard predicts 6.86 for the two low contributors and 10.28 for the medium contributor. If we count an observation of 6.5 for the low contributors as consistent with this standard, this would leave 11 for the medium contributor, therefore we consider observations in the interval [10,11] as consistent with the proportional standard for this type. This is certainly an ad-hoc criterion, however, we tried various other criteria, such as allowing for a fixed deviation in both directions of a point prediction, and results are rather stable.

[28] Here we refer to the strict ordering, i.e. the egalitarian standard would not reflect contribution ordering in our definition, as higher contributions are not rewarded with higher shares.

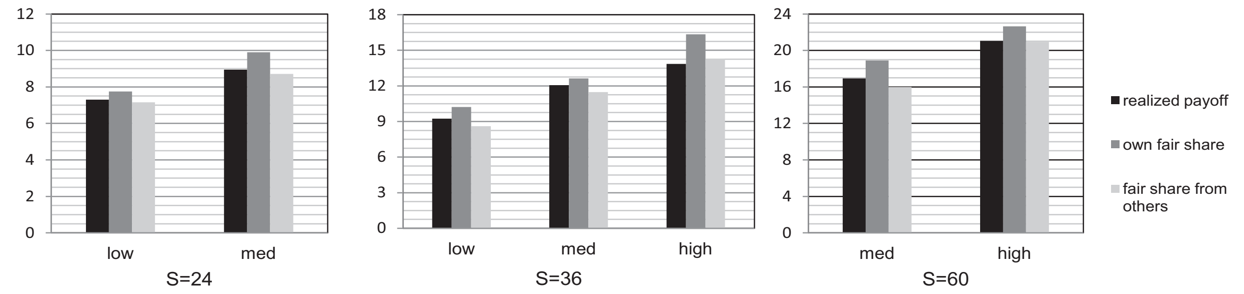

Fig. A.2. Fairness and Extended Divide-the-Dollar mechanism.

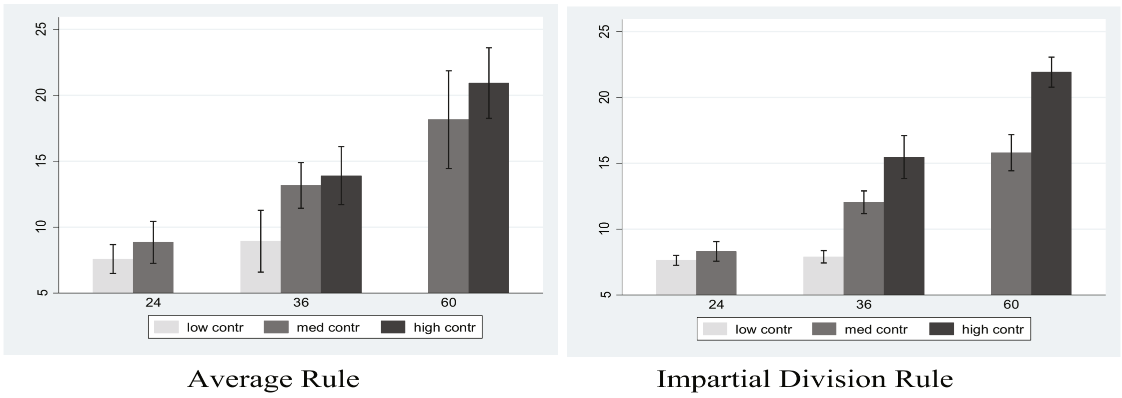

Fig. A.3. Standard deviations in payoffs for average rule and impartial division rule.

Greed levels: Overall, greed levels in the Extended Divide-the-Dollar Rule do not depend on subjects' contributions towards the cake (WSR: $p = 0.49$ for the comparison between low vs. medium contributor in $S = 24$; $p = 0.7$ for low vs. high contributor in $S = 36$; and $p = 0.8$ for medium vs. high contributor in $S = 60$). Only exception is the medium contributor in $S = 36$: his greed level is significantly lower than the low contributor's (WSR: $p < 0.02$ and *t*-test: $p < 0.03$) and also lower than the high contributor's (WSR: $p = 0.12$ and *t*-test: $p < 0.09$). This is probably due to the fact that only for this type all fairness standards discussed above point towards the same allocation of 12. In the experiment, 2/3 of all subjects in this role claimed precisely this amount, and also 2/3 of all subjects in this role were assigned an entitlement of 12 or more from their partners, which explains the low greed level.

Strategic behaviour: Claims and stated own fair shares do not differ significantly for low contributors in the small cake size and for medium contributors in all cake sizes, suggesting that these types tend to submit truthful reports of their evaluation of own claims. High contributors, on the other hand, systematically claim less than their stated own fair share (WSR: $p < 0.05$ and *t*-test: $p < 0.03$ for $S = 36$; WSR: $p < 0.01$ and *t*-test: $p < 0.01$ for $S = 60$), which can be interpreted as an attempt to lower their greed levels. But, we also find that low contributors in the medium cake size claim more than they stated as own fair share (WSR: $p < 0.05$ and *t*-test: $p < 0.03$), which would – ceteris paribus – imply higher greed levels. Knowing that there was a significant self-serving bias in the fairness question, one might expect that reporting even higher own claims would further increase these subjects' greed. However, payoffs finally depend on relative greed levels. We thus need to check whether low contributors systematically assign smaller entitlements to a certain type compared to what they assigned to them as fair share. This would mean that they try to compensate high reports of their own claim by actively increasing some other partner's greed level. In fact, low contributors (partner A) in the medium cake size do precisely that – they assign slightly more to partner B, and considerably less to C (*t*-test: $p < 0.001$) compared to what they stated as fair for each of these partners. Note that the disadvantaged partner C is the high contributor, who would take away a large share of the cake in case he gets paid off first. In comparison, the high contributor C has some extra amount to allocate between his partners compared to the fairness question, since his own claim under this mechanism is lower than his stated own fair share in the fairness question. C allocates this extra amount by reporting, on average, 0.4 more for the medium contributor compared to what he assigned to B in the fairness question, and 0.8 more to the low contributor. This might be interpreted as a strategic move, too, as C would prefer B to be the last one being paid off, since B takes away more than A from the cake. Consistent with this strategic reasoning is also what we find for $S = 60$: Both high contributors report, on average, 0.7 more for the low contributor than for their fellow high-contribution partner (*t*-test: $p < 0.03$ for B, and $p < 0.1$ for C).

Average Rule. Fig. A.3 shows that the Average Rule leads to much larger standard deviations in payoffs compared to the Impartial Division Rule.

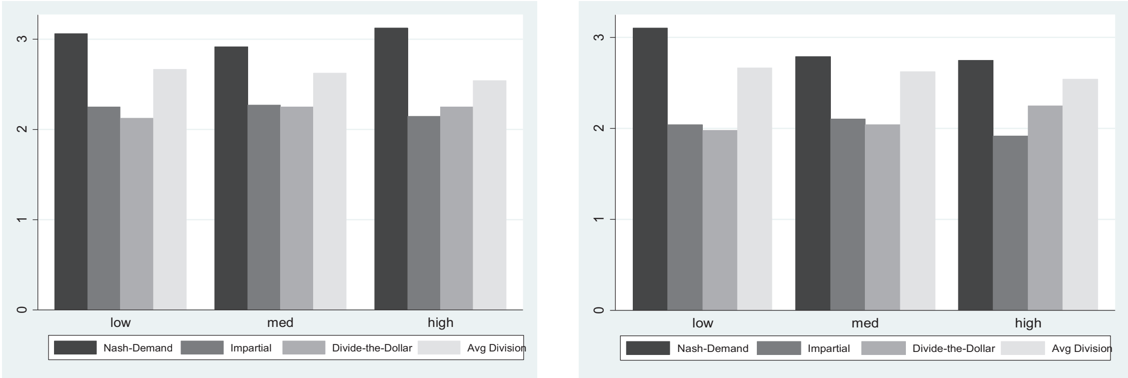

Fig. A.4. Mean ranks of all mechanisms' procedural (left) and allocative (right) fairness.

Fig. A.4 shows a comparison of the procedural and allocative fairness for the data of the additional treatment including all four mechanism.

Appendix B. Experimental Instructions

Experimental instructions associated with this article can be found, in the online version, at http://dx.doi.org/10.1016/j.jebo.2016.07.019.

References

Babcock, L., Wang, X., Loewenstein, G., 1996. Choosing the wrong pond: social comparisons in negotiations that reflect a self-serving bias. Q. J. Econ. 111 (1), 1–19.
Brams, S.J., Taylor, A.D., 1994. Divide the dollar: three solutions and extensions. Theory Decis. 37, 211–231.
Burrows, P., Loomes, G., 1994. The impact of fairness on bargaining behavior. Empir. Econ. 19, 201–221.
Cadigan, J., Schmitt, P., Shupp, R., Swope, K., 2009. An experimental study of the holdout problem in a multilateral bargaining game. South. Econ. J. 76 (2), 444–457.
Cai, H., 2000. Delay in multilateral bargaining under complete information. J. Econ. Theory 93 (2), 260–276.
Cappelen, A., Drange, A., Sorensen, E., Tungodden, B., 2007. The pluralism of fairness ideals: an experimental approach. Am. Econ. Rev. 97 (3), 818–827.
Chatterjee, K., Sabourian, H., 2000. Multiperson bargaining and strategic complexity. Econometrica 68 (6), 1491–1510.
Cramer, R., Schotanus, D., Ibrahim, K., Colbeck, N., 2009. Continuous well-flow estimates improve production allocation. Oil Gas J. 107 (47).
Dana, J., Weber, R.A., Kuang, J.X., 2007. Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. Econ. Theory 33, 67–80.
DeClippel, G., Moulin, H., Tideman, N., 2008. Impartial division of a dollar. J. Econ. Theory 139, 176–191.
Fischbacher, U., 2007. Z-Tree: Zurich toolbox for ready-made economic experiments. Exp. Econ. 10 (2), 171–178.
Gächter, S., Riedl, A., 2005. Moral property rights in bargaining with infeasible claims. Manag. Sci. 51 (2), 249–263.
Gächter, S., Riedl, A., 2006. Dividing justly in bargaining problems with claims. Soc. Choice Welf. 27 (3), 571–594.
Gantner, A., Horn, K., Kerschbamer, R., 2013. Fair division in unanimity bargaining with subjective claims. In: Working Papers in Economics and Statistics 2013–31, University of Innsbruck.
Heller, M.A., 1998. The tragedy of the anticommons: property in the transition from Marx to markets. Harv. Law Rev. 111, 621–688.
Hoffman, E., Spitzer, M.L., 1985. Entitlements, rights, and fairness: an experimental examination of subjects' concepts of distributive justice. J. Legal Stud. 14, 259–297.
Kagel, J., Kim, C., Moser, D., 1996. Fairness in ultimatum games with asymmetric information and asymmetric payoffs. Games Econ. Behav. 13 (1), 100–110.
Karagözoglu, E., 2012. Bargaining games with joint production. In: Bolton, G., Croson, R. (Eds.), Oxford Handbook of Economic Conflict Resolution. Oxford University Press, New York.
Karagözoglu, E., Riedl, A., 2015. Performance information, production uncertainty, and subjective entitlements in bargaining. Manag. Sci. 61 (11), 2611–2626.
Konow, J., 2000. Fair shares: accountability and cognitive dissonance in allocation decisions. Am. Econ. Rev. 90 (4), 1072–1092.
Konow, J., 2003. Which is the fairest one of all? A positive analysis of justice theories. J. Econ. Lit. 41, 1188–1239.
Libecap, G.D., Wiggins, S.N., 1984. Contractual responses to the common pool: Prorationing of crude oil production. Am. Econ. Rev. 74 (1), 87–98.
Moulin, H., 2002. Axiomatic cost and surplus sharing. In: Arrow, K.J., Sen, A.K., Suzumura, K. (Eds.), In: Handbook of Social Choice and Welfare, vol. I. Elsevier Science B.V.
Mumy, G.E., 1981. A superior solution to captain MacWhirr's problem: an illustration of information problems and entitlement structures. J. Polit. Econ. 89 (5), 1039–1043.
Schokkaert, E., Lagrou, L., 1983. An empirical approach to distributive justice. J. Public Econ. 21, 33–52.
Thomson, W., 2003. Axiomatic and game theoretic analysis of bankruptcy and taxation problems: a survey. Math. Soc. Sci. 45, 249–297.
Tidemann, N., Plassmann, F., 2008. Paying the partners. Public Choice 136 (1/2), 19–37.