

Efficiency versus Equality in Bargaining

Fabio Galeotti* Maria Montero[†] Anders Poulsen[‡]

March 2, 2018

Abstract

We consider how the outcome of bargaining varies with changes in the trade-off between equality, efficiency, and total-earnings maximization. We observe that subjects avoid an equal-earnings outcome if it is Pareto inefficient; a large proportion of bargaining pairs avoid an equal and Pareto efficient outcome in favor of one giving unequal and total-earnings maximizing payoffs, and this proportion increases when unequal outcomes imply larger earnings to one of the players, even though this also implies higher inequality; finally, we document a compromise effect that violates the independence of irrelevant alternatives condition.

Keywords: bargaining, efficiency, equality, communication, experiment, independence of irrelevant alternatives.

JEL Classification: C70, C72, C92.

*Univ Lyon, CNRS, GATE L-SE UMR 5824, F-69131 Écully, France.

[†]School of Economics, University of Nottingham, Nottingham NG7 2RD, UK.

[‡]School of Economics and CBESS, University of East Anglia, Norwich NR4 7TJ, UK.

We thank the Editor and four reviewers, whose comments and suggestions significantly improved the paper. We also thank Nejat Anbarci, Nick Feltovich, Werner Güth, Emin Karagözoğlu, Martin Sefton, Daniel Seidmann, Robert Sugden, and James M. Watson, as well as seminar audiences at UEA, the Annual NIBS meeting, the META workshop in Nottingham, GATE, UPV/EHU and the Tinbergen Institute, for helpful comments. All errors are ours. Support from the Network for Integrated Behavioural Science (NIBS) and the British Academy is gratefully acknowledged. This research was performed within the framework of the LABEX CORTEX (ANR-11-LABX-0042) of Université de Lyon, within the program Investissements d’Avenir (ANR-11-IDEX-007) operated by the French National Research Agency (ANR).

The welfare properties of *equity* and *efficiency* may determine the focal equilibrium in any game, whether there is an arbitrator or not. (Myerson, 1991, p. 373, italics in original)

1 Introduction

Bargaining is ubiquitous in economic and social life. An employer negotiates with a union about wages and working conditions. A buyer and seller negotiate over price, product specifications, delivery, and warranty terms. A couple negotiate over which house to buy. Creditors negotiate over the division of the assets of a bankrupt company.

Which properties of an agreement make it attractive, or *focal*? In this paper we present the findings from an empirical investigation related to Myerson's conjecture quoted above, that a focal agreement possesses some combination of desirable welfare properties, such as efficiency and equity. We consider three potentially salient welfare properties: equality of money earnings, Pareto efficiency, and total-earnings maximization. For simplicity, we refer to the latter as 'total-earnings efficiency' and to Pareto efficiency as 'efficiency'.¹

Suppose no agreement offers both equal and maximal total earnings. Then there is a *trade-off* between these properties: if bargainers agree on equal earnings they must accept a reduction in total earnings (perhaps even abandon efficiency); conversely, if they agree to maximize total earnings, they must accept some earnings inequality. Our central research questions are: how do bargainers tend to resolve this trade-off? And, how does this resolution vary with changes in the terms of the trade-off? Suppose, for example, that the feasible unequal agreements become more unequal but also offer larger total earnings. Will this make bargainers more or less likely to agree on an equal-earnings outcome? Although these questions seem both foundational and relevant to real world bargaining situations, no systematic and general investigation has, as far as we know, taken place.

¹We define these properties in terms of money amounts. Of course, Pareto efficiency in terms of preferences does not necessarily coincide with Pareto efficiency in terms of money amounts, since subjects may care about other subjects' money earnings.

To answer these questions, we design an experiment where pairs of subjects negotiate over a given finite set of *contracts*. A contract specifies an amount of money to each person within the pair. If the subjects can agree on a contract, each gets the implied money; otherwise neither person gets any money. By collecting data for a number of games that vary in the contracts and hence the efficiency-equality trade-off, we can in a *ceteris paribus* manner assess how the focality of contracts depends on and varies with properties such as earnings equality, efficiency, and total earnings.² Our approach differs from the typical bargaining experiment, where subjects negotiate over a fixed sum of money, and where there is consequently no trade-off between equality and total-earnings maximization.

In our experiment, subjects play several games and are rematched from round to round. Even if we cannot entirely rule out repeated interaction effects, our matching procedure makes them unlikely, and our results suggest they are, if anything, minimal and not systematic.

We embed our contract setup in an unstructured bargaining protocol that allows bargainers to make as many proposals as they wish within a certain period of time.³ They can communicate via chat, and any agreement is binding. The advantage of using a free-form unstructured bargaining protocol, apart from its inherent realism, is that it makes bargainers strategically equal (see the discussion in Gächter and Riedl, 2005; Camerer et al., 2017), so our data on how the efficiency-equality trade-off influences focality are not affected or confounded by some bargainers being in a strategically inferior position.

²If transfers are allowed, players could achieve equality, efficiency and total-earnings efficiency by agreeing on actions that maximize the size of the ‘cake’ and on transfers to equate earnings. In many real-world situations, there are however constraints on the transfers that can be made. Suppose two siblings inherit two indivisible objects, A and B. Both siblings prefer object A to object B. They are liquidity constrained, so the sibling that gets A cannot compensate the sibling that got B. They may agree to get one object each, or they may sell the items and divide the proceeds equally. Depending on how marketable the items are, the proceeds from the sale may be quite low, and so the equal-earnings outcome may not be total-earnings efficient, and may even fail to be efficient.

³See e.g. Roth and Malouf, 1979; Roth and Murnighan, 1982; Anbarci and Feltovich, 2013; Embrey et al., 2016.

While our primary goal is empirical, it is useful to consider the predictions that economic bargaining theory makes for how changes in the efficiency-equality trade-off affect the bargaining outcome. We consider two well-known cooperative bargaining theories (Thomson, 1994), the Nash Bargaining Solution (Nash, 1950) and the Kalai-Smorodinsky Solution (Kalai and Smorodinsky, 1975).⁴ As we show below, these bargaining solutions sometimes make different predictions about how changes in the efficiency-equality trade-off affect the bargaining outcome.⁵

Our main findings are as follows. First, in the benchmark case where there is an equal-earnings contract that is also total-earnings efficient, almost all bargainers settle on it, as the findings from the existing bargaining literature would lead one to expect. When the equal-earnings contract is efficient, but not total-earnings efficient, the focality of the equal contract falls gradually as we lower its total earnings. In some of our bargaining situations, more than half of bargaining pairs settle on an unequal-earnings contract.

Second, we observe a strong tendency for bargainers to avoid the equal-earnings contract when it is inefficient, regardless of what the unequal contracts offer.⁶

The third main result is based on a comparative statics exercise where we exogenously increase the conflict of interest over the unequal contracts. This is done by fixing the equal contract payoffs and raising the payoff that each unequal contract offers to the player who gets his or her largest money earnings from it. The data show that, when the equal-earnings contract is efficient, this results in a decrease (increase) in the proportion of agreements

⁴In deriving predictions we assume as a natural benchmark that both bargainers are self-interested, but we also allow for preference heterogeneity.

⁵We emphasize that our primary purpose is to shed empirical light on the effects of changes in the equality-efficiency trade-off rather than test these bargaining theories. Testing a theory about bargaining would require measuring bargainers' preferences and ensuring that these are common knowledge (see Roth and Murnighan, 1982; Roth, 1995).

⁶Recall that we measure Pareto efficiency in money terms, and our subjects may have additional concerns other than the maximization of money earnings. Still, our finding that most subjects disregard an equal and inefficient contract appears even stronger when allowing for this possibility.

on the equal (unequal) contract(s). This suggests that bargainers are more occupied with maximizing their own, and possibly to some extent the other person's payoff, than with ensuring equality of earnings.

The fourth main finding concerns the observed cheap-talk communication. In bargaining games with an equal but not total-earnings efficient allocation and *two* unequal and total-earnings efficient allocations, many bargaining pairs create and agree to use a randomization device, typically based on the rock-paper-scissors game, to decide which unequal allocation they should agree on. Although informal and non-binding, the 'recommendation' by this device is almost always followed by the bargainers. Interestingly, in bargaining situations where there is a *single* unequal contract alternative to an equal and efficient contract, rock-paper-scissors is almost never used by subjects. Bargainers are thus willing to rely on randomization as a way to break a 'tie' between two symmetric contracts, but not to overcome conflict in general.

The last main finding is that in most of the bargaining situations, and on average across all situations, more bargaining pairs settle on an equal contract when there are two unequal contracts rather than one. We interpret this as a violation of the Independence of Irrelevant Alternatives (IIA) axiom (Nash, 1950). This axiom, which plays an important role in cooperative bargaining and social choice theory, states that if a contract is agreed on when there is a large set of available contracts, then the same contract (assuming it is still feasible) is selected when the set of alternative contracts is reduced.⁷ As we explain below, the violation happens because a contract can be focal not only due to its absolute payoff properties, but also from being a *compromise* between other feasible agreements.⁸ IIA fails to capture the latter, contextual, source of focality, and in our data this effect is sufficiently strong to violate IIA.

⁷A more formal statement of IIA is: suppose a contract x is feasible both when the set of feasible contracts is S , and when it is T , where $S \subseteq T$. Then, if x is agreed on when the set is T , x is also agreed on when the set is S .

⁸See de Clippel and Eliaz (2012) for a theoretical analysis of an ordinal bargaining solution that generates such a compromise effect.

As far as we know, we are the first to empirically demonstrate a violation of the IIA axiom in a bargaining environment. Nydegger and Owen (1975) experimentally test IIA by comparing a basic situation in which two subjects must agree on how to divide a dollar with a constrained situation in which player 1 cannot receive more than 60 cents whereas player 2 can still potentially receive the whole dollar. They observe that subjects divide the dollar equally in both cases, consistent with IIA. However, their experiment is a relatively weak test since there is a contract that is both equal and total-earnings maximizing, which makes it strongly focal. Our experiment tests IIA in a more demanding setting where no contract is both equal and maximizes total earnings.⁹

As we describe below, neither the Nash nor the Kalai-Smorodinsky solution can capture all these empirical findings. Allowing for preference heterogeneity improves the theories' predictions, but the Nash Bargaining Solution (NBS) remains unable to capture the compromise-based source of focality of an equal earnings agreement.

The rest of the paper is organized as follows. Section 2 discusses the related literature. In Section 3 we describe the experimental design and procedures. Section 4 outlines our hypotheses. The data are presented and analyzed in Section 5. Section 6 concludes. The Online Appendix contains experimental instructions and additional analysis.

2 Related Literature

Existing empirical bargaining research (for surveys, see Roth, 1995; Camerer, 2003; specific studies include Nydegger and Owen, 1975; Roth et al., 1988; Herreiner and Puppe, 2010; Feltovich and Swierzbinski, 2011; Karagözoğlu and Riedl, 2015; Camerer et al., 2017) has done very little work on investigating how the bargaining outcome varies with changes in the trade-offs between equality and efficiency. The reason is simply that in these studies bargainers typically negotiate over a fixed sum of money. An equal-

⁹Bone et al. (2014) also observe a tendency to compromise in their data. However, their experiment is not a direct test of IIA since they do not compare larger choice sets with smaller ones.

earnings outcome is then also total-earnings efficient, which means that there is no trade-off between equality and total-earnings maximization.¹⁰ The only studies that we are aware of that vary the trade-off are Roth et al. (1988), Herreiner and Puppe (2010), and Isoni et al. (2014).

Isoni et al. (2014) consider unstructured bargaining situations where subjects claim valuable assets placed on a 'bargaining table'.¹¹ Their design only allows for two possibilities for the efficiency–equality trade-off: *either* an equal earnings contract is inefficient *or* it is total-earnings efficient and hence efficient. The equality-efficiency trade-off in their bargaining environment is thus either very severe, or there is no trade-off at all. A key contribution of our paper is to also consider the arguably important intermediate case, where an equal-earnings division is efficient but not total-earnings efficient.

Roth et al. (1988) consider binary lottery bargaining games where people bargain over the division of one hundred lottery tickets, and where the number of lottery tickets a player gets determines his or her probability of winning a personal money prize. In terms of expected monetary payoffs, these games generate a linear Pareto frontier with two potentially focal payoff pairs on it, namely an equal and efficient pair (obtained by an unequal division of the lottery tickets) and an unequal pair that offers larger total earnings (generated by an equal division of lottery tickets). Roth et al. (1988) vary the slope of the Pareto frontier, and hence the efficiency-equality trade-off, by simultaneously reducing the low prize and increasing the high prize. This is observed to make the equal allocation less focal. However, this response can be due to i) the equal allocation being less focal due to

¹⁰Note that if a contract is total-earnings efficient, it must also be efficient, but the converse is not true. As a simple example, consider an equal contract offering (Player 1, Player 2) earnings (40, 40). If the only other feasible contract is (30, 80), then (40, 40) is total-earnings inefficient but still efficient. If the alternative contract is (50, 80), then the equal contract is both inefficient and total-earnings inefficient.

¹¹In their environment not only payoff divisions but also spatial cues generated by the locations of the assets on the table can be focal. These spatial cues are payoff-irrelevant, but can still be a source of focality – see Schelling, 1960; Mehta et al., 1994; Crawford et al., 2008; Isoni et al., 2013.

its lower payoffs, and/or ii) unequal allocations becoming more focal due to higher total payouts. We think it is important to disentangle these factors, and so we in a *ceteris paribus* fashion either i) keep the equal allocation fixed and vary the unequal alternative allocation(s), or ii) keep the unequal allocation(s) fixed and vary the equal allocation.

In the experiment by Herreiner and Puppe (2010) subjects negotiate for 10 minutes over how to divide four indivisible objects. This gives rise to a non-linear Pareto frontier, and these differ across their games in a way that makes it difficult to attribute any observed behavioral difference to the presence or absence of some specific property of the feasible payoffs or frontier. Our experiment varies the set of available contracts more systematically, often by changing just one contract at a time. Another important feature of Herreiner and Puppe's design is that an equal and efficient contract is very close to being total-earnings maximizing; thus their design does not allow one to assess the focality of equal-earnings contracts that are efficient but far from being total-earnings maximizing.

3 Experimental Design and Bargaining Games

3.1 Design and Procedures

The experiment was conducted at the experimental lab of the Centre for Behavioural and Experimental Social Science, at the University of East Anglia (Norwich, United Kingdom). We ran 7 sessions with 16 participants each, making a total of 112 subjects.¹² The experiment was programmed and conducted with the software z-Tree (Fischbacher, 2007); recruitment was done using ORSEE (Greiner, 2015). Average earnings (including a £4 show-up fee) were £16.15. No session lasted more than one hour.

Subjects arrived to the lab and were allocated a desk. Instructions (see the Online Appendix) were circulated and read aloud by the experimenter. Subjects were informed that they would make decisions in 22 bargaining situations, referred to in the instructions as 'scenarios'. In each scenario they would be randomly matched with another participant (stranger matching),

¹²Background information on the participants is in the Online Appendix.

and presented with a set of (two or three) feasible contracts, displayed on subjects' computer screen in a random order. Each contract specified a number of points to each paired subject, such as (50,50) or (40,240).¹³

The two matched subjects were referred to as Persons 1 and 2.¹⁴ A subject was informed that he or she would in some scenarios be referred to as Person 1, and in others as Person 2.¹⁵

In each scenario each pair of subjects had 120 seconds to negotiate over which contract to agree on. During this time they could make contract proposals, and write messages to each other. A subject made a contract proposal by clicking with their mouse on one of the feasible contracts (see Figure 1 for a screenshot). As long as an agreement had not been reached a subject was free to change his or her contract proposal, or to retract it without replacing it with a new one, in real time and as frequently as desired. Subjects could also decide not to make any proposals at all. A binding agreement was reached if and only if the two players proposed (that is, clicked on) the same contract.¹⁶ If the subjects did not reach an agreement before the 120 seconds expired, neither earned any points from the scenario.

The subjects could also write cheap-talk messages to each other while making proposals. There were no constraints on the number and content of messages, except that subjects were told to avoid writing messages that revealed their identity, that physically threatened the other subject or that discussed what might or might not happen outside the lab. Subjects were informed that, if it was detected that a participant wrote any such messages,

¹³Using a small number of contracts allows us to vary the trade-off between efficiency and equality in a simple and transparent manner that minimizes other confounding factors.

¹⁴The advantage of using labels for the two subjects is that it is easy to describe and refer to a contract, and it is clear who gets how much. Moreover, each matched pair of subjects see the same representation of contracts on the screens (and this is common information). The potential disadvantage is that labels may have an effect on behavior. We did not find any differences between the two players in the data (see the Online Appendix).

¹⁵The alternative approach, that a subject was either Person 1 or 2 in all scenarios, has the disadvantage of reducing the number of possible matchings dramatically (a subject in the role of Person 1 could only be matched with those in the role of Person 2), and hence a subject would more frequently be matched with the same other participant.

¹⁶The same agreement technology is used in other papers, such as Roth and Murnighan (1982) and Feltovich and Swierzbinski (2011).

the subject would not receive any money earnings. Subjects could make proposals without sending messages, and vice versa.

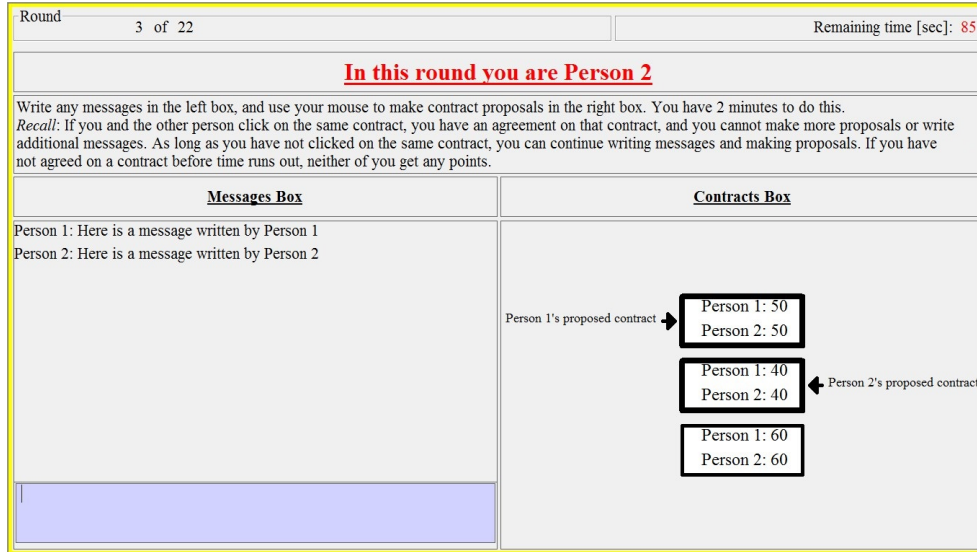


Figure 1: Screenshot.

The 22 bargaining scenarios were not known to subjects in advance. Subjects were informed that they would not be matched with the same participant in all the rounds, and that different subjects would encounter the scenarios in a different order.¹⁷

Since different subjects encounter the scenarios in a different order, learning effects in the data affect all scenarios to a similar extent and hence should not lead to systematic aggregate effects that would bias comparisons across scenarios. Having subjects play different scenarios against different co-participants minimizes repeated-game effects, since it is difficult to reward or punish a specific individual for his or her behavior in past scenarios. One

¹⁷There were some unavoidable constraints on the matching protocol, due to the real-time nature of the bargaining: if a participant plays a given scenario as, say, the seventh in his game sequence, then another subject in the room must also be playing that scenario as her seventh scenario. The matching protocol maximized the dispersion of matchings subject to these constraints (on average, subjects encountered the same person 1.46 times). Details can be found in the Online Appendix, where we also find that the effects of game sequence and past outcomes on current bargaining behavior are insignificant. The Online Appendix also contains an analysis of individual behavior across the twenty two games.

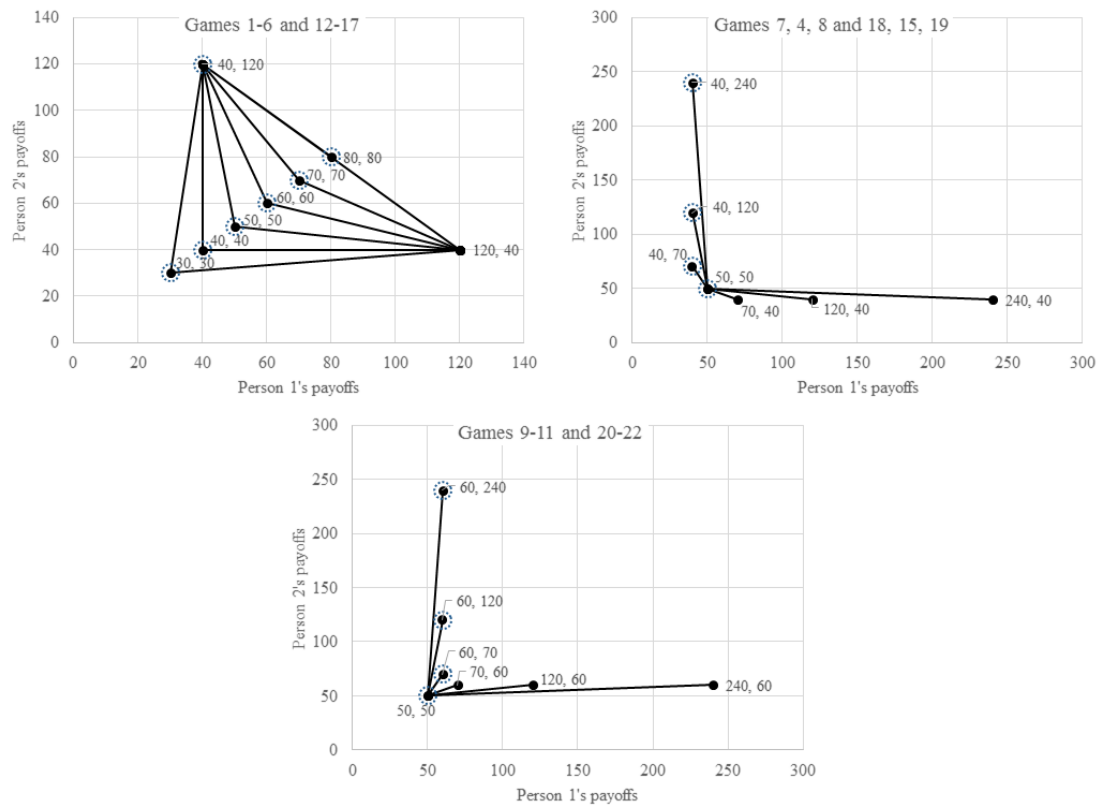
concern is that subjects could coordinate on a rule or heuristic for the overall game, composed of the 22 scenarios. For example, subjects could use a rule whereby in each scenario they select a total-earnings efficient outcome, to ensure that the overall outcome will maximize total earnings. Alternatively, they could select an equal-earnings outcome in all scenarios, to make sure that the overall outcome is also equal. Our design does not allow us to rule out this possibility, and subjects did occasionally mention such rules during bargaining, as we shall see from the coding data. However, if such a rule was used consistently on a large scale, we would expect approximately the same rate of agreement on an equal contract in all our 22 scenarios. In fact, behavior varies in a systematic way across scenarios, suggesting that observed behavior is mostly specific to the current scenario rather than reflecting some rule or heuristic that is being applied to all of them.

When everyone had completed the 22 scenarios the computer randomly selected three rounds for payment; the same three rounds were selected for all subjects in a given session. The conversion rate from points to pounds was 20 points = £1.

3.2 The Bargaining Games

We refer to an unstructured bargaining situation with a given set of feasible contracts as a *bargaining game* (called a scenario in the experiment). Each subject played 22 bargaining games, shown in Table 1. A visual depiction is given in Figure 2. The numbers are measured in experimental points. As we explain below, these games differ in the trade-off between equality and efficiency, and were selected to allow us to test a number of hypotheses about this trade-off that are presented in Section 4.

Games 1–11 are of the form $\{(z, z), (w, y), (y, w)\}$, where $0 < z < y$ and $0 < w < y$. Since there are two alternatives to the equal contract, each of which offers the largest money payoff to a different player, we call them *two-sided* games. Games 12–22 are of the form $\{(z, z), (w, y)\}$, and are *one-sided* games. Each two-sided game has a corresponding one-sided game. We collected data for both types of games since each seemed relevant for



Note: One-sided games are represented by dotted circles.

Figure 2: Visual depiction of all games.

real world bargaining situations.¹⁸

When $z \leq w$, the equal-earnings contract (z, z) is inefficient. These are Games 5-6, 9-11, and their corresponding one-sided versions 16-17, 20-22. When $z > w$, the equal-earnings outcome is efficient. This applies to Games 1-4, 7-8, 12-15, and 18-19. If $2z \geq w + y$, the equal contract is total-earnings efficient, and thus also efficient. There are two such games, 1 and 12.¹⁹

¹⁸Suppose two siblings inherit a house. They can either sell the house and divide the proceeds equally, or one of them can keep the house and pay the other a small (due to liquidity constraints) rent. If both siblings live in the same city, the situation is naturally two sided. However, if one sibling lives abroad, the situation is one sided: either they sell

Game	Feasible contracts	Equal contract efficient?	Equal contract total-earnings efficient?
1	(80,80),(40,120),(120,40)	Y	Y
2	(70,70),(40,120),(120,40)	Y	N
3	(60,60),(40,120),(120,40)	Y	N
4	(50,50),(40,120),(120,40)	Y	N
5	(40,40),(40,120),(120,40)	N	N
6	(30,30),(40,120),(120,40)	N	N
7	(50,50),(40,70),(70,40)	Y	N
8	(50,50),(40,240),(240,40)	Y	N
9	(50,50),(60,70),(70,60)	N	N
10	(50,50),(60,120),(120,60)	N	N
11	(50,50),(60,240),(240,60)	N	N
12	(80,80),(40,120)	Y	Y
13	(70,70),(40,120)	Y	N
14	(60,60),(40,120)	Y	N
15	(50,50),(40,120)	Y	N
16	(40,40),(40,120)	N	N
17	(30,30),(40,120)	N	N
18	(50,50),(40,70)	Y	N
19	(50,50),(40,240)	Y	N
20	(50,50),(60,70)	N	N
21	(50,50),(60,120)	N	N
22	(50,50),(60,240)	N	N

Notes: Y = Yes; N = No.

Table 1: The 22 bargaining games.

4 Theory and Hypotheses

In this section we formulate hypotheses for how changes in the efficiency-equality trade-off influence which contracts are agreed upon. Our bargaining environment allows players to freely decide when to make a contract proposal, to communicate via free chat, and to sign a binding agreement.

the house and split the proceeds equally, or the sibling that lives locally keeps the house.

¹⁹We only collected data for two games with an equal and total-earnings efficient contract, since we anticipated that this contract would be strongly focal. The data confirmed this.

Hence, we derive our hypotheses from the predictions of cooperative bargaining solution concepts (see Thomson, 1994). Besides the Nash bargaining solution (NBS) which is the main solution concept for two-person bargaining problems, we derive predictions also for the Kalai and Smorodinsky (1975)'s solution (KSS).²⁰ We start from the benchmark assumption of self-interested bargainers with utilities equal to own material payoffs (henceforth 'selfish' bargainers), but we also state the predictions of these solution concepts for a heterogeneous population in which many different utility functions are represented. For expositional reasons, we first derive hypotheses for two-sided games and then compare these with one-sided games.

4.1 Two-sided Games: Changes in the Total Earnings of the Equal and Efficient Contract

In Games 1-4 the payouts of the equal-earnings contract are gradually reduced, while keeping the unequal contracts unchanged. The NBS predicts (one of) the contract(s) with the highest Nash product. If utilities coincide with material payoffs, the Nash product of a contract $x = (x_1, x_2)$ is $x_1 x_2$. Under the assumption of selfish preferences, the NBS values efficiency (since increasing one player's material payoff always leads to a higher Nash product) and equality (since, for a fixed $x_1 + x_2$, allocations that are more equal have a higher Nash product). Hence, an equal and total-earnings maximizing contract is always predicted if present (as in Game 1). The equal contract is still chosen in Game 2 since $70 \times 70 > 40 \times 120$, but not in Games 3 (since $60 \times 60 < 40 \times 120$) and 4. While being equal and total-earnings maximizing guarantees that a contract will be predicted by the NBS, being equal and Pareto efficient is not sufficient.

Hypothesis 1. *(Games 1-4, NBS with selfish preferences) The equal (unequal) allocation is chosen in Games 1 and 2 (3 and 4).*

In contrast, the KSS selects the equal allocation in all four games. Intu-

²⁰These solution concepts and their predictions are described in detail in the Online Appendix. The NBS and the KSS have been defined for finite sets of contracts by Mariotti (1998) and Nagahisa and Tanaka (2002) respectively. Another important variant of the NBS, due to Conley and Wilkie (1996), is not defined for finite sets of contracts.

itively, the KSS attempts to bring both players as close as possible to their ideal allocation; none of the two players should be making too large a concession. With selfish preferences, this means that in all contract sets of the form $\{(z, z), (w, y), (y, w)\}$, where $0 < w < z < y$, the equal allocation is predicted.

Hypothesis 2. (*Games 1-4, KSS with selfish preferences*) *The equal allocation is chosen in all four games.*

An extensive experimental literature has found that agents care not only about their own material payoff but also about the material payoffs of others. In particular, agents may care about equality (see Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000) or may care positively about the other player's earnings (see Charness and Rabin, 2002; Engelmann and Strobel, 2004). Suppose we have a heterogeneous population where some players are selfish, others are inequity averse, and others have social welfare preferences.²¹ For any given pair of bargainers, we can calculate the theoretical prediction of each solution concept (NBS and KSS) under the assumption of complete information about preferences.²² Then both theories predict a gradual decrease (rather than a sudden drop from 100 to 0 percent, or no drop at all as was the case with selfish preferences) in the proportion of pairs that agree to the equal contract as its payoffs decrease. For the NBS, some pairs would switch to an unequal contract after a small drop in the earnings of the equal contract (this is the case if at least one of the players has a sufficiently high σ_i), while others would stick to the equal contract even if its total payoffs are low (this is the case if both players are sufficiently

²¹ More formally, $u_i(x) = x_i$ (selfish), $u_i(x) = x_i - \alpha_i(\max(x_j - x_i, 0)) - \beta_i(\max(x_i - x_j, 0))$ (inequity averse) and $u_i(x) = (1 - \sigma_i)x_i + \sigma_i x_j$ for $x_i \leq x_j$; $u_i(x) = (1 - \rho_i)x_i + \rho_i x_j$ for $x_i > x_j$ (social welfare) all have strictly positive probability. Moreover, within the inequity averse types, any (α_i, β_i) with $0 \leq \beta_i \leq \alpha_i$, $0 < \alpha_i$ and $\beta_i < 1$ has strictly positive density. Within the social welfare types, any (σ_i, ρ_i) such that $0 < \sigma_i \leq \rho_i < 1$ and $\sigma_i < 0.5$ has strictly positive density. Even though we assume a continuum of possible types, a finite number of types would lead to the same comparative statics predictions provided that the type space is sufficiently rich.

²²We view complete information as an "as if" simplifying assumption; our subjects can communicate but they are unlikely to learn each other's exact preferences.

inequity averse). For the KSS, an unequal contract can only be agreed on if at least one bargainer has social welfare preferences.²³ The drop in the predicted frequency of the equal contract is also gradual for the KSS, with pairs where one or both players have sufficiently strong social welfare preferences switching earlier, and pairs with weaker social welfare preferences switching only when the payoffs of the equal allocation are very low.

Hypothesis 3. (*Two-sided Games 1-4: NBS and KSS with a heterogeneous population*) Suppose the total payoffs offered by the equal contract decrease, while remaining efficient. Then i). Agreements on the equal contract become less likely. ii). Agreements on an unequal contract become more likely.

Note that i) and ii) do not follow logically from each other, since there could also be a change in the frequency of disagreement (even though this is not predicted by the theories).

4.2 Two-Sided Games: Changes in the Payoff to the Favored Player in the Unequal Contracts

Consider next the effect of increasing the payoff to the favored player in the unequal contracts (the payoff denoted y) while keeping an efficient equal-earnings contract fixed (Games 7-4-8). With selfish preferences, the NBS predicts an unequal allocation in the three games we consider since $70 \times 40 > 50 \times 50$, and increasing the payoff to the favored player from 70 to 120 to 240 would only increase the Nash product of the unequal allocations. On the other hand, it follows from our earlier discussion that the KSS predicts the equal allocation in all three games.

Hypothesis 4. (*NBS, selfish preferences*). An unequal allocation is chosen in Games 7-4-8.

Hypothesis 5. (*KSS, selfish preferences*). The equal allocation is chosen in Games 7-4-8.

²³Players who care positively about the other's payoff would not be making as much of a concession in agreeing to a disadvantageously unequal allocation, and may even rank the equal allocation last.

While the two theories make opposite predictions under selfish preferences, once we allow for a heterogeneous population there is considerable consensus between the theories as to the direction of the comparative statics predictions. In any heterogeneous population (of the type described in footnote 21), the KSS predicts an increasing frequency of agreements on an unequal allocation. This increase is entirely driven by the presence of players with social welfare preferences.

The predictions of the NBS depend on the exact distribution of preferences (see the Online Appendix). For example, in a heterogeneous population where all players have either selfish or social welfare preferences, the frequency of the unequal allocation increases gradually as we go from Game 7 to 4 and then to 8. In contrast, in the presence of inequity averse players the direction of the prediction is ambiguous. Sufficiently inequity averse players always agree to the equal allocation when they meet each other, and become more likely to agree to the (advantageously) unequal allocation as y increases when they meet a player with selfish or social welfare preferences. The ambiguity arises in pairs where both players are inequity averse, but not so much that they would necessarily agree to the equal allocation. For these pairs, when y increases, a contract that offers y becomes more attractive to the favored player, but less attractive to the unfavored player, and the Nash product may increase or decrease.

Since the KSS unambiguously predicts an increase in the frequency of the unequal allocation in a heterogeneous population, while the NBS also predicts it in many cases (including, but not limited to, the case where the proportion of inequity averse players is sufficiently small),²⁴ we take this prediction as our hypothesis.

²⁴Indeed, an increasing frequency in Games 7-4-8 can occur even if all players are inequity averse. As a simple example, suppose there are three types of players in the population, ranked by how inequity averse they are. These are type H ($\alpha = 1, \beta = 0.75$), type M ($\alpha = 0.3, \beta = 0.25$) and type L ($\alpha = 0.05, \beta = 0$). For $y = 70$, the only case in which an unequal allocation is chosen is if both players are of type L . For $y = 120$, an unequal allocation is chosen if one of the players is type L and the other is either L or M . For $y = 240$, an unequal allocation is chosen if at least one of the players is of type L , irrespective of the type of the other player.

Hypothesis 6. (*KSS with a heterogeneous population; NBS with a heterogeneous population where most players have selfish or social welfare preferences*) Suppose the payoff to the favored player in the unequal contracts increases, while the equal and efficient contract remains fixed. The proportion of agreements on an unequal (equal) contract increases (falls) in Games 7-4-8.

4.3 Two-sided Games with Monetarily Inefficient Contracts

Next, we consider the games with an equal but monetarily inefficient contract. If players are primarily concerned with maximizing their own earnings (or with maximizing some weighted average of their own and the other subject's earnings), both bargaining solutions predict that the equal contract would be avoided, since both players would be better-off by agreeing to another contract.²⁵

Hypothesis 7. (*Games 6 and 9-11; both theories with selfish or social welfare preferences*) No bargainers agree to a monetarily inefficient contract, regardless of the properties of the alternative contracts.

The theories may still predict the equal contract if bargainers are sufficiently inequity averse. By raising the payoff to the favored player in Games 9-11, we are making it increasingly challenging for Hypothesis 7 to hold.

4.4 Two- Versus One-Sided Games

From a theoretical point of view, the value of one-sided games is in the comparison between one-sided games and two-sided games.²⁶ Consider

²⁵We also have one game (Game 5, with contracts (40, 120), (40, 40), (120, 40)) in which the equal allocation is only weakly inefficient. The NBS still rules out the equal contract in this case, but the KSS allows it if both players are selfish.

²⁶In fact, under very general assumptions on preferences, the NBS and KSS coincide when there are only two contracts. Denote the two contracts by $x = (x_1, x_2)$ and $x' = (x'_1, x'_2)$. Suppose each player prefers a different contract (say, player 1 prefers x and player 2 prefers x') and both contracts are preferred to disagreement. Contract x is predicted by the KSS if player 2 would be making a smaller concession by agreeing to x than player 1 would be making by agreeing to x' , that is if (normalizing the utilities of disagreement to 0) $\frac{u_2(x)}{u_2(x')} \geq \frac{u_1(x')}{u_1(x)}$. The NBS predicts contract x if $u_1(x)u_2(x) \geq u_1(x')u_2(x')$, which is equivalent. The two solution concepts also agree in other cases, such as when both players prefer the same contract (that contract is of course selected) or when one of the players prefers disagreement to one of the contracts (the other contract is selected).

Nash's Independence of Irrelevant Alternatives (IIA) axiom (Nash, 1950), which is one of the axioms satisfied by the NBS (an extension of this axiom to multivalued solutions can be found in Mariotti, 1998). Suppose that a contract x is feasible in a game with contract set T and also in a game with contract set S , where S is a subset of T . An implication of the axiom is that, if x is the unique solution when the contract set is T , then x must also be the unique solution in the smaller set S . Let T be a two-sided game with equal allocation x , and S the corresponding one-sided game where x is still feasible. Then IIA predicts the following:²⁷

Hypothesis 8. (*Independence of Irrelevant Alternatives*) *In the two-sided game, the proportion of bargainers agreeing on the equal contract never exceeds the one for the corresponding one-sided game.*

The intuition is quite clear: if the equal contract was chosen over two other alternative contracts, it should still be chosen over only one of them.²⁸

An alternative to Hypothesis 8 is that the equal contract can be attractive not just because of its inherent property of offering equal earnings, but also because it serves as a *compromise* between two other rival contracts, as in the two-sided game. If this compromise source of focality is sufficiently strong, the equal contract may be *more* focal in the two- than in the one-sided game. This alternative hypothesis is consistent with the KSS, since agreeing to the unequal allocation that favors the other player is less of a concession if there is no favorable unequal allocation.²⁹

²⁷We ignore the knife-edge case in which the equal contract has the same Nash product as both or one of the unequal contracts in T and the solution is multivalued. We expect these cases to be rare in our data, unlike ties between the two unequal contracts.

²⁸Since subjects are rematched from round to round, we do not observe the same pairs bargaining over the two sets of contracts, so we expect IIA to hold only in a probabilistic sense. Suppose IIA holds for each possible pair in the population; then any pair that agrees on the equal allocation in a two-sided game would still agree on it in the corresponding one-sided game. Hence, equal allocations must be at least as common in the one-sided game at the population level. The actual frequency of agreements on equality may occasionally be lower in the one-sided game because we are looking at a sample rather than at all pairs in the population, but it should not be *systematically* lower. If it is, IIA is violated.

²⁹Despite this intuition, the KSS does not necessarily predict a compromise effect in a heterogeneous population. If player 1 has social welfare preferences and player 2 is inequity

5 Experimental Findings: Bargaining Outcomes

In analysing the data, we focus on the bargaining outcomes. Additional analysis (such as of individual behavior) is reported in the Online Appendix. Table 2 shows descriptive statistics for each of the 22 games (the *Game* column).³⁰

The number of bargaining pairs that played each game is given in the *Obs* column.³¹ The table reports for each game the proportions of disagreements, agreements on the equal contract, and agreements on an unequal contract (columns *Rate of disagreement*, *Rate of agreement on equality*, and *Rate of agreement on inequality*). These proportions sum to 100, and the rate of agreement equals $100 - \text{Rate of disagreement}$. See also Figure 3 for a visual representation of the agreement proportions in the games (Games 4 and 15 have deliberately been included twice, in order to ease comparisons).

Table 2 also shows the average time to agree for the pairs that did reach an agreement (column *Average time to agree*),³² the average total earnings of a bargaining pair (*Average total-earnings efficiency*, both in points and as a percentage of the maximum achievable value), the number of proposals made per subject (*Average number of proposals*), and the average number of messages sent per subject (*Average number of messages*).³³ Whenever we make pairwise comparisons of games, we use session averages as the units of observation, in order to control for the non-independence of the observations at the individual level. All statistical tests are two-tailed, and, unless otherwise mentioned, significance refers to the 5 % level.

averse, they may agree on (40, 120) in the two-sided game (40, 120), (50, 50), (120, 40) but agree on the equal contract in the one-sided game (120, 40), (50, 50).

³⁰We do not find any statistically significant effects of subject labels (1 or 2) on behavior. Similarly, the way in which a given set of contracts was ordered on the screen has no effect. Details are in the Online Appendix. We therefore pool all data.

³¹A technical problem at the end of one of our sessions resulted in the loss of some data; as a result we have a different number of observations for some games.

³²The Online Appendix gives more details on agreement times.

³³We also looked at the number of characters in the text instead of the number of messages sent. The results are qualitatively identical and are available from the authors upon request. We thank Nick Feltovich for suggesting this analysis.

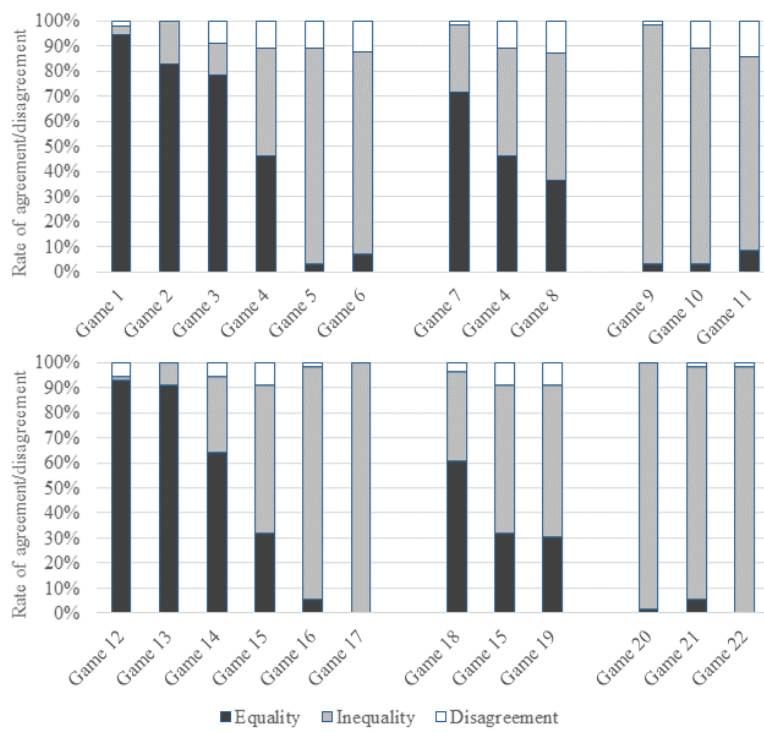


Figure 3: Proportion of different bargaining outcomes for each game

Game	Obs.	Feasible Contracts	Rate of disagreement	Rate of agreement on equality	Rate of agreement on inequality	Average time to agree (in sec.)	Average total-earnings efficiency (points)	Average total-earnings efficiency (%)	Average number of proposals	Average number of messages
1	55	(80,80), (40, 120), (120,40)	1.82%	94.55%	3.63%	39.3	157.09	98.18%	1.25	2.06
2	53	(70,70), (40, 120), (120,40)	0%	83.02%	16.98%	48.74	143.4	89.62%	1.24	3.05
3	56	(60,60), (40, 120), (120,40)	8.93%	78.57%	12.5%	50.02	114.29	71.43%	1.46	3.41
4	56	(50,50), (40, 120), (120,40)	10.71%	46.43%	42.86%	68.86	115	71.88%	1.45	4.61
5	55	(40,40), (40, 120), (120,40)	10.91%	3.64%	85.45%	77.67	139.64	87.27%	1.42	5.09
6	56	(30,30), (40, 120), (120,40)	12.50%	7.14%	80.36%	68.76	132.86	83.04%	1.48	4.87
7	56	(50,50), (40, 70), (70,40)	1.79%	71.43%	26.78%	57.31	100.89	91.72%	1.48	3.19
8	55	(50,50), (40, 240), (240,40)	12.73%	36.36%	50.91%	70.98	178.91	63.90%	1.43	4.98
9	56	(50,50), (60,70), (70,60)	1.79%	3.57%	94.64%	62.18	126.61	97.39%	1.38	3.94
10	55	(50,50), (60, 120), (120,60)	10.91%	3.64%	85.45%	75.57	157.45	87.47%	1.38	4.67
11	56	(50,50), (60, 240), (240,60)	14.29%	8.93%	76.78%	81.29	239.29	79.76%	1.42	5.15
All 2-sided	609	-	7.88%	39.57%	52.55%	63.20	145.88	83.76%	1.40	4.10
12	56	(80,80), (40, 120)	5.36%	92.86%	1.78%	35.91	151.43	94.64%	1.33	2.04
13	56	(70,70), (40, 120)	0%	91.07%	8.93%	47.89	141.79	88.62%	1.13	2.87
14	56	(60,60), (40, 120)	5.36%	64.29%	30.35%	55.49	125.71	78.57%	1.27	2.78
15	56	(50,50), (40, 120)	8.93%	32.14%	58.93%	61.41	126.43	79.02%	1.28	3.38
16	56	(40,40), (40, 120)	1.79%	5.36%	92.85%	33.98	152.86	95.54%	1.06	2.01
17	56	(30,30), (40, 120)	0%	0.00%	100%	30.32	160	100.00%	1.02	1.74
18	56	(50,50), (40, 70)	3.57%	60.71%	35.72%	47.8	100	90.91%	1.15	2.65
19	56	(50,50), (40, 240)	8.93%	30.36%	60.71%	52.96	200.36	71.56%	1.23	3.28
20	55	(50,50), (60, 70)	0%	1.82%	98.18%	28.35	129.45	99.58%	1.07	1.41
21	56	(50,50), (60, 120)	1.79%	5.36%	92.85%	30.47	172.5	95.83%	1.15	1.63
22	56	(50,50), (60, 240)	1.79%	0.00%	98.21%	43.02	294.64	98.21%	1.22	2.5
All 1-sided	615	-	3.41%	34.96%	61.63%	42.27	159.61	90.21%	1.17	2.39
All	1224	-	5.64%	37.25%	57.11%	52.44	152.78	87.00%	1.29	3.24

Note: Obs. = Observations (number of pairs).

Table 2: Average aggregate bargaining outcomes. Note: The average number of proposals and average number of messages are per subject.

We start by considering efficiency across the games. The rate of disagreement is in general quite low (the average is below 6%). In several games there are no disagreements at all, and the rate of disagreement never exceeds 15%.

In spite of the generally high rate of agreement, there are some noteworthy differences between the games. In two-sided games the rate of agreement tends to decrease as the equal contract becomes less efficient (Games 1-6) (Page's trend test, $p = 0.014$). The agreement rate also decreases when, keeping the equal contract fixed, we make the unequal contract(s) more unequal. When the equal contract is efficient (Games 7, 4, and 8), increasing the amount y offered in the contracts $(40, y)$ and $(y, 40)$ results in an increasing rate of disagreement (Page's trend test, $p = 0.054$). Something similar happens in Games 9-11, where the equal contract is inefficient (Page's trend test, $p = 0.041$). Interestingly, there is no significant pattern in one-sided games.

We summarize these findings in:

Finding 1. *In two-sided games, the rate of agreement tends to fall as the equal contract offers lower total earnings and eventually ceases to be efficient. The rate of agreement also tends to fall when the unequal contracts become more unequal. No significant pattern is found in one-sided games.*

Since disagreement is ruled out by both the NBS and the KSS, these findings are not predicted by the theory. An intuitive explanation is that the equal contract becomes less attractive when its total payouts drop (in the language of the NBS, its Nash product decreases; in the language of the KSS, bargainers would be making a greater concession by agreeing to it). Since in two-sided games there are two unequal contracts, each of which favors a different player, subjects may bargain harder for the unequal allocation that favors them and this makes it more difficult to agree.

In one-sided games there is only one unequal allocation, so that reducing the payoffs of the equal contract does not necessarily make it harder to agree (in fact, conflict disappears if payoffs drop so much that the equal contract becomes inefficient). Similarly, making the unequal contracts more

unequal increases the conflict of interest in two-sided games. In one-sided games, the conflict of interest is less pronounced since there is only one unequal allocation; indeed, if players have social welfare preferences, increasing the payoff to the favored player would reduce the conflict of interest in one-sided games but not in two-sided games. Thus, the increase in disagreements for two-sided games appears to be the result of chicken-type bargaining involving the two unequal contracts.

5.1 Changes in the Total Payoff Efficiency of the Equal Contract

In Games 1–4 more pairs settle on an equal and efficient contract than on an unequal and total-earnings maximizing contract. As we lower the total payouts of the equal contract while keeping the payoffs of the unequal contracts fixed, there is a monotonic and statistically significant drop in the rate of agreement on the equal contract (Page’s trend test, $p < 0.001$).³⁴

Finding 2. (Hypotheses 1, 2, and 3) *In games 1–4, almost all bargaining pairs agree on the equal contract when it is total-earnings efficient. The proportion of agreements on the equal contract falls gradually as its total payouts decrease, but a majority of agreements remain on the equal contract as long as it remains efficient.*

Finding 2 shows that the focality of an equal and efficient contract decreases only gradually as its total earnings are lowered; if an equal and efficient contract ceases being total-earnings efficient, there is no sharp decline in its focality. Hence, total-earnings efficiency is a sufficient, but not a necessary condition for subjects to accept an equal agreement. On the other hand, the equal allocation being efficient does not prevent a substantial proportion of bargainers from settling on an unequal allocation.³⁵

The findings on Games 1–4 are not consistent with the KSS for selfish preferences, since it predicts the equal contract in all games (cf Hypothesis 2). They are qualitatively more in line with the NBS for selfish prefer-

³⁴In pairwise comparisons based on a Wilcoxon signed-rank test, at the 5% level of significance, the relationship between the rates of agreement on equality for two-sided games is: Game 1 = Game 2 > Game 3 > Game 4.

³⁵Our data thus fail to lend support to the ‘principle’ stated in Herreiner and Puppe (2010), p. 230: “First, determine the most equal distribution of rewards. If this contract is Pareto optimal, then choose it.”

ences (Hypothesis 1), which predicts a fall in the proportion of agreements on the equal allocation as its total payoffs decrease, although the data do not support the NBS's prediction of a sharp fall (from 100% to zero) in the proportion of equal agreements. Note also that overall most agreements remain on the equal contract in Games 1–4, hence quantitatively the KSS explains a greater proportion of the data than the NBS with selfish preferences. The data are however consistent with there being a distribution of different preferences in the population, regardless of which bargaining solution is assumed (Hypothesis 3).

5.2 Changes in the Earnings of the Favored Player in the Unequal Contract(s)

Comparing Games 7, 4, 8 shows that when the unequal contracts offer more to one player and the same to the other (that is, we increase y in the contracts $(40, y)$ and $(y, 40)$), a significantly smaller proportion of bargaining pairs agree on an equal contract, *and* a significantly larger proportion agree on an unequal contract (Page's trend test, $p < 0.05$).³⁶

Finding 3. (Hypothesis 4, 5, and 6) *Consider the two-sided games with an equal and efficient contract $(50, 50)$ and unequal contracts of the form $(40, y)$ and $(y, 40)$, where $y = 70, 120, 240$ (Games 7, 4, 8). When y increases, significantly fewer subjects agree on an equal and efficient contract, and significantly more agree on an unequal contract.*

These findings are inconsistent with both bargaining theories under the assumption that bargainers are selfish (cf Hypotheses 4 and 5). They are, however, consistent with there being a distribution of bargainers with different preferences in the population.

We can interpret Finding 3 as follows. When we go from Game 7 to 4, and from 4 to 8, three things happen: first, each unequal contract offers a higher reward to one of the players; second, the total earnings in the unequal contracts increase; third, the unequal contracts become more unequal.

³⁶Game 7 > Game 4 = Game 8 for agreement on equality (pairwise comparisons based on a Wilcoxon signed-rank test). For agreements on inequality, the inequalities go in the opposite direction.

We intuitively expect the first factor to lead a self-interested subject to bargain harder in favor of his or her preferred unequal contract; the second factor makes unequal contracts more attractive to subjects who, in addition to their own, care about total earnings; the last factor, to the contrary, makes the equal contract more attractive to subjects who dislike inequality. The data suggest that in the population as a whole the first and second effects dominate the third.

We may also ask what sort of preference distribution fits our data best for each of the theories. Since the KSS predicts the equal allocation for any combination of selfish and inequity averse preferences in Games 1-4 and 7-8, any agreements on an unequal allocation would require the presence of at least one bargainer with social welfare preferences. We observe agreement rates on inequality as high as 50% (Game 8). A requirement that 50% of pairs contain one player with social welfare preferences translates into a proportion of social welfare preferences in the population of approximately 29%. This is just a lower bound, since some of the players who agreed to the equal contract may also have social welfare preferences. Conversely, for the NBS, selfish players would always agree to an unequal allocation, thus the 71% frequency of agreements on the equal allocation in Game 7 would be due to at least one bargainer being inequity averse (or, more implausibly, to both bargainers having social welfare preferences with high values of ρ_i and low values of σ_i). A requirement that 71% of pairs contain at least one inequity averse player would translate into at least 46% of bargainers being inequity averse.³⁷

5.3 Monetarily Inefficient Contracts

Consider next the two-sided games with a monetarily inefficient equal contract; these are Games 5-6 and 9-11. From Table 2 we see that in none of these games does the percentage of agreements on the equal allocation exceed 10%. As we saw, in Games 1-4, the frequency of the equal contract

³⁷Relatively mild levels of inequity aversion would be sufficient to explain the data. Both bargainers having $\alpha_i \geq \beta_i \geq 0.1$ would be sufficient for the equal allocation to be predicted by the NBS in Game 7.

drops only gradually as its total payouts decrease. However, as soon as the equal contract becomes inefficient (Games 5-6), the proportion of agreements on the equal contract drops dramatically.

Finding 4. (Hypothesis 7) *When the equal-earnings contract is monetarily inefficient, it is agreed on very rarely.*³⁸

These findings provide support for Hypothesis 7. Note also that when the equal contract is only weakly monetarily dominated (Game 5), it remains as infrequently agreed on as when it is strictly monetarily dominated (Games 6,9-11).

5.4 IIA versus Compromise

Recall that we interpret Independence of Irrelevant Alternatives (IIA) as predicting that the rate of agreement on an equal contract in a one-sided game is at least as large as in the corresponding two-sided game. Table 2 and Figure 3 show, however, that this is not what we typically observe. In almost all the two-sided Games 1–8 the proportion of agreements on equality is *larger* than in the corresponding one-sided Games 12–19, violating IIA.³⁹

Finding 5. (Hypothesis 8) *The average rate of agreement on the equal contract is significantly higher in two-sided than in one-sided games ($p = 0.05$), violating Independence of Irrelevant Alternatives.*

These findings reject Hypothesis 8 in favor of a compromise effect. An equal-earnings outcome not only gets some focality from its unique and absolute property of offering equality of earnings (a property which holds regardless of which other contracts are available), but also because it serves as a *compromise* between more extreme contracts. This context-dependent property is relevant in two-sided but, by design, not in one-sided games, and in our data it is strong enough to violate IIA.

³⁸Sign tests, $p > 0.1$ for all comparisons between the sample mean and 0.

³⁹One-to-one comparisons between a two-sided game and the corresponding one-sided game show that the average rate of agreement on the equal contract is significantly larger in Game 6 than 17 ($p = 0.085$), 7 than 18 ($p = 0.049$), and 11 than 22 ($p = 0.048$).

The systematic deviations from IIA we observe allow us to reject the NBS under very weak assumptions on preferences (not just any combination of selfish, inequity averse and social welfare preferences, but any preferences where the utility of a contract x depends only on x itself and not on what other contracts are available).

5.5 A Payoff Based Metric

In this section we consider if it is possible to construct a simple payoff-based measure that can capture the main observed differences between the bargaining games.⁴⁰

Consider a game with two contracts, one equal and another unequal. Denote these by $\{(y, w), (z, z)\}$ where $y > w$ and $y + w \geq 2z$ (that is, the unequal allocation always maximizes total payoffs). The equal allocation may or may not be Pareto efficient, depending on whether $z > w$.

Suppose first that the equal allocation is efficient, that is, $z > w$. We define the Gain Sacrifice Ratio, GSR, as the ratio between the gain in payoffs by one player and the loss in payoffs for the other player, if they agree to the unequal allocation instead of the equal one ($GSR = \frac{y-z}{z-w}$). Since the unequal allocation maximizes total payoffs, this ratio is always at least 1. The larger GSR is, the more one of the players gains, relative to the other's loss when they agree to the unequal allocation instead of the equal one – the trade-off between equality and efficiency becomes starker.

The GSR ratio can also be computed when the equal allocation is strongly inefficient, i.e. $z < w$ (we discard the knife-edge case $z = w$). In this case the ratio is negative, but a larger absolute value can be viewed as a more severe equality-efficiency trade-off. For example, in game $\{(50, 50), (240, 60)\}$, GSR would be $\frac{240-50}{50-60} = -19$: agreeing to the unequal allocation makes both players better-off (hence the 'sacrifice' is negative) but the payoff increase of one of the players is nineteen times greater than that of the other.⁴¹

⁴⁰We thank a referee for drawing our attention to this issue.

⁴¹Interestingly, there is a link between the GSR and the utility functions discussed earlier. A positive GSR is related to the value of σ that would make a player with social welfare preferences indifferent between the equal contract and the disadvantageously unequal contract. The utilities would be $u(w, y) = (1 - \sigma)w + \sigma y$ and $u(z, z) = z$ respectively.

For two-sided games of the form $\{(y, w), (z, z), (w, y)\}$, one can define the gain-sacrifice ratio in exactly the same way. It now measures how much one of the players can gain relative to the other's loss if they agree to one of the unequal contracts instead of the equal contract.⁴²

Our games vary in the GSR, and Figure 4 shows how the bargaining outcomes depend on this metric. For positive values of the GSR (equal contract is Pareto efficient) we see a clear (non-linear) relationship between the GSR and bargaining outcomes. When GSR increases, i) the proportion of agreements on an unequal allocation increases, ii) that on the equal allocation decreases, and iii) the proportion of disagreements increases.

The situation is very different for negative values of the GSR (that is, when the equal contract is Pareto inefficient). For two-sided games, we observe that an increase in the absolute value of the GSR results in a lower frequency of agreements on inequality and a greater frequency of disagreements. In contrast, GSR has very little explanatory power of disagreements in one-sided games. We provide a statistical analysis of all these patterns in the Online Appendix.

We also consider an alternative metric, the Total Payoff Efficiency Ratio, $TPER = \frac{y+w}{2z}$, the ratio between total earnings in the unequal and the equal contract, and find that overall the GSR perform better. See the Online Appendix for details.

If they are equal, we obtain $(1 - \sigma)w + \sigma y = z$, which can be rewritten as $\frac{1-\sigma}{\sigma} = \frac{y-z}{z-w}$, which is the GSR. The larger the GSR, the lower the threshold value of σ and consequently the less concern for the other's payoffs we need to assume in order for a player to prefer the disadvantageously unequal contract. On the other hand, if the GSR is negative, the relationship is with the value of α in the Fehr-Schmidt utility function. Suppose a player is indifferent between the equal allocation and the disadvantageously unequal allocation where both players get more. Then $u(w, y) = w - \alpha(y - w) = z = u(z, z)$. This can be rewritten as $\frac{1+\alpha}{\alpha} = -\frac{y-z}{z-w}$, where the right-hand side is the absolute value of the GSR. The larger the GSR in absolute value, the less concern for inequity one needs to assume in order for an inequity averse player to prefer the equal allocation.

⁴²We are implicitly assuming that the equal contract acts as a reference point from which gains and losses are calculated.

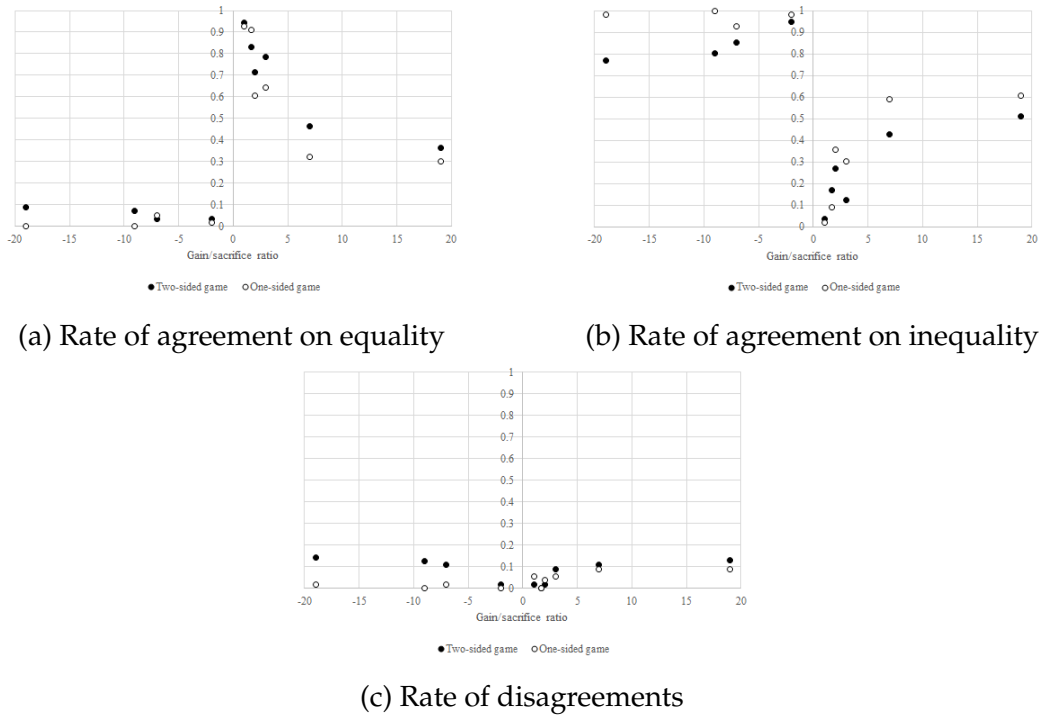


Figure 4: Gain-sacrifice ratio and bargaining outcomes

5.6 Analysis of Chat Conversations

We analyzed the content of conversations using the same method as Brandts and Cooper (2007) and Cooper and Kühn (2014). Two assistants separately coded the conversations by ticking categories out of a list. Coding was binary for each category, with a conversation coded as a 1 if it contained the category and zero otherwise. The categories are not mutually exclusive, and coders could tick as many or as few categories as they wished.

5.6.1 Chat Categories

Below is a list of the categories that were most common in our data; the entire list is in the Online Appendix.

Fairness/equality: one or both subjects refer to fairness or equality.

Rock-paper-scissors (RPS): the contract agreed upon is determined by the outcome of a game of RPS.⁴³

⁴³This was typically achieved by both subjects agreeing to type one of the three words or letters (R, S, P) when the timer on their screens reached a particular value (for example,

Standoff: each of the two subjects proposes, either by chat or by clicking, the *unequal* contract that favors them. By definition, this is only possible in two-sided games.

Standoff-equal: one of the two subjects proposes, either by chat or by clicking, the unequal contract that favors them, and the other subject proposes the equal contract.

Generous proposal: one subject *takes the initiative* in proposing, either by chat or clicking, the contract that favors the *other* subject. We asked the coders to tick this category only if the subject proposed this contract from the outset, and the other subject had not previously proposed the same contract, by chat or clicking.

Both better-off: one or both subjects refer to some contract(s) as being better for both subjects compared to other contracts.

Cost-benefit: one or both subjects refer to the costs and benefits of agreeing on one rather than another contract (e.g., “I gain more than you lose”, “this will not cost you much”, or similar).

Be generous: a subject appeals to the other’s generosity.

Off topic: one or both subjects write messages that are not directly related to the experiment, e.g. cracking jokes, talking about the weather, music, football, etc.

Supergame: this is a composite category where we pool several individual categories that might be related to supergame reasoning. This includes any references to past or to future scenarios, such as general rules of behavior for all scenarios, claims of having been unlucky or generous in past scenarios, and references to karma.

As in Brandts and Cooper (2007) and Cooper and Kühn (2014), we use the average of the coders’ output. For example, one of the coders found references to fairness/equality in 15.5% of all games and the other in 17.6% of all games; hence on average they found references to fairness in 16.5%

when there were 90 seconds left) in order to achieve simultaneity of moves. The subjects then agreed on the contract preferred by the winner of the RPS game.

of all games, and this is the value we use.⁴⁴ The Cohen’s kappa coefficient of agreement (Cohen, 1960) between the two coders is high, particularly for the categories that were more common in the data (see bottom of Table 3 and the Online Appendix).

5.6.2 Analysis

Table 3 summarizes the results of the coding, focusing on categories that were assigned to at least 5 percent of the bargaining situations by the coders (in particular, the table reports the percentage frequencies of the chat categories for every game and possible bargaining outcome).⁴⁵ The full data set can be found in the Online Appendix.

In the remaining of this section we focus on the five most prominent categories – *RPS*, *fairness/equality*, *standoff*, *standoff-equal* and *generous proposal* – and the hybrid category *supergame*. For each of these categories, we investigate (i) in which games it is observed, and (ii) what it accomplishes in terms of bargaining outcome. To answer the first question, we ran, for each category, an ordered logistic regression over all bargaining games. The dependent variable is whether a category is assigned to a given game by neither, only one or both coders. Independent variables include a dummy for each of the 22 games, and the round in which the game occurred. The results are reported in the Online Appendix. In what follows, when we compare the frequency of each category across games, the p values are based on pairwise comparisons between the coefficients of these regressions.

To answer the second question, we simply rely on descriptive statistics (this is because we cannot treat the different categories of chat conversation as being independent). In particular, we compare the proportion of disagreements, agreements on equality, and agreements on inequality respectively across the different chat categories. These proportions are reported at

⁴⁴These percentages are calculated over all games, not just over games where there was a chat. It is possible for three of our categories to be ticked even though there was no chat (*standoff*, *standoff-equal* and *generous proposal*).

⁴⁵There are only five categories that occur in more than 10% of the games: *RPS*, *fairness/equality*, *standoff*, *standoff-equal* and *generous proposal*. There are also four categories with frequencies between 5% and 10%: *both better off*, *cost-benefit*, *be generous* and *off topic*. The composite *supergame* category appears in 6.2% of the games.

the bottom of Table 3.

Rock-paper-scissors

RPS is much more common in two-sided than in one-sided games (25% of all two-sided games were coded as RPS, compared with only 1.9% of one-sided games; $p < 0.001$). This suggests that while bargainers find it salient to use RPS as a mechanism to decide between two unequal allocations (reconciling total-earnings efficiency with *ex ante* equality), they do not find it attractive to use as a general mechanism to resolve conflict over two contracts. For example, in the (one-sided) Game 14, 30% of games were coded as a standoff between the equal and the unequal allocation. None of the pairs used RPS to resolve the standoff situation.

If we focus on how the frequency of RPS varies across two-sided games, we see that its frequency increases in Games 1–6 as the equal allocation becomes less attractive (Game 1 = Game 2 = Game 3 < Game 4 = Game 5 = Game 6, $p < 0.01$) – this is consistent with subjects looking for alternative agreements and using RPS as a means to choose between the two unequal allocations. Similarly, the frequency of RPS increases when we compare Games 7, 4 and 8 (Game 7 < Game 4 = Game 8, $p < 0.05$),⁴⁶ and when we compare 9, 10 and 11, although for the last three games the differences are not significant. This is consistent with our previous finding that subjects agreed on inequality more often when the unequal allocations offer greater total payoffs, even though the allocations are more unequal.

It is also worth noting that, even if RPS is often associated with agreements on inequality (22.5% of these agreements are linked to RPS), not all agreements on inequality are the result of using RPS. As we have seen, agreements on inequality are more common in one-sided games, where RPS is almost never used. Even in two-sided games, only about half (46.72%) of unequal agreements are the result of using RPS.

⁴⁶As we mentioned earlier, statistical tests are computed based on pairwise comparisons between the coefficients of an ordered logistic regression.

Game	RPS	Fairness	Standoff	Standoff-Equal	Generous Proposal	Both Better Off	Cost Benefit	Be Generous	Off topic	Supergame ^d
1	3.6%	21.8%	3.6%	21.8%	1.8%	1.8%	1.8%	1.8%	14.5%	1.8%
2	13.2%	17.0%	7.5%	22.6%	1.9%	0.0%	0.0%	5.7%	15.1%	5.7%
3	10.7%	35.7%	17.9%	14.3%	1.8%	0.0%	0.0%	1.8%	12.5%	0%
4	28.6%	21.4%	21.4%	17.0%	3.6%	0.9%	10.7%	8.9%	8.9%	21.4%
5	40.0%	15.5%	40.0%	5.5%	10.0%	0.0%	4.5%	17.3%	14.5%	11.8%
6	40.2%	13.4%	39.3%	8.0%	15.2%	14.3%	6.3%	8.9%	12.5%	8.9%
7	13.4%	23.2%	16.1%	21.4%	3.6%	0.9%	8.9%	6.3%	8.0%	3.6%
8	29.1%	18.2%	36.4%	7.3%	5.5%	0.0%	10.9%	17.3%	0.9%	11.8%
9	28.6%	6.3%	25.9%	8.0%	12.5%	8.0%	13.4%	5.4%	4.5%	2.7%
10	32.7%	15.5%	31.8%	8.2%	16.4%	11.8%	4.5%	17.3%	3.6%	3.6%
11	34.8%	25.9%	31.3%	10.7%	9.8%	15.2%	0.9%	11.6%	2.7%	8.9%
All 2-sided	25.0%	19.5%	24.7%	13.1%	7.5%	4.8%	5.7%	9.3%	8.9%	7.3%
12	0.0%	20.5%	-	15.2%	0.0%	0.0%	6.3%	2.7%	4.5%	0.9%
13	5.4%	25.0%	-	19.6%	0.0%	0.0%	5.4%	7.1%	5.4%	4.5%
14	0.0%	21.4%	-	30.4%	8.9%	0.9%	7.1%	8.0%	7.1%	5.4%
15	5.4%	17.0%	-	28.6%	8.9%	0.0%	39.3%	17.0%	3.6%	24.1%
16	0.0%	5.4%	-	7.1%	33.0%	0.0%	1.8%	4.5%	1.8%	3.6%
17	0.0%	3.6%	-	2.7%	23.2%	11.6%	0.0%	3.6%	4.5%	1.8%
18	3.6%	15.2%	-	17.9%	8.9%	0.0%	19.6%	5.4%	5.4%	4.5%
19	5.4%	19.6%	-	25.9%	14.3%	0.0%	27.7%	16.1%	6.3%	7.1%
20	0.9%	9.1%	-	4.5%	35.5%	11.8%	0.0%	0.9%	4.5%	0%
21	0.0%	5.4%	-	8.0%	41.1%	17.0%	0.9%	0.9%	3.6%	1.8%
22	0.0%	8.0%	-	8.0%	33.9%	17.9%	1.8%	1.8%	6.3%	1.8%
All 1-sided	1.9%	13.7%	-	15.3%	18.9%	5.4%	10.0%	6.2%	4.8%	5.0%
All	13.4%	16.5%	12.3%	14.2%	13.2%	5.1%	7.8%	7.7%	6.8%	6.1%
<i>Disagreement</i>	2.2% ^b	19.6%	44.9%	27.5%	2.9%	5.8%	20.3%	16.7%	14.5%	18.8%
<i>Equality</i>	1.2%	25.2%	7.1%	19.0%	0.2%	0.8%	3.9%	4.2%	7.0%	3.2%
<i>Inequality</i>	22.5%	10.6%	12.4%	9.8%	22.7%	7.9%	9.2%	9.2%	5.9%	6.9%
κ	0.979	0.879	0.951	0.866	0.811	0.639	0.651	0.799	0.537	0.665

Notes: The table includes only categories that were assigned to at least 5 percent of the bargaining games (averaging across coders). ^a Supergame is a composite category made of several individual categories related to supergame reasoning. κ = Cohen's kappa coefficient of agreement (Cohen, 1960) between the two coders. At the bottom of the table (before the κ) we also report the percentage frequency of the chat categories for agreements reached on equality, agreements reached on inequality, and disagreements. For example, in games that resulted in an equal agreement the fairness chat category was observed 25.2% of the time, but only 10.6% of the time in games that resulted in agreement on an unequal contract. ^b There are two bargaining situations where subjects used RPS but they agreed on a contract after the deadline (i.e. they clicked on the same contract few instants after the time had elapsed and while the computer was loading the next screen). The coders have recorded these situations as RPS even if no agreement was officially reached before the deadline.

Table 3: Percentage frequency of the most common chat categories

Finding 6. *The frequency of RPS differs, often significantly, across games. Within two-sided games, RPS is more often used in games where the unequal contracts offer greater total payoffs relative to the equal contract. In contrast, RPS is almost never used in one-sided games.*

Fairness

A comparison of two-sided with one-sided games shows that fairness/equality references tend to be more common in two-sided games ($p = 0.004$); this is consistent with the greater frequency of equal agreements in two-sided games.

Focusing on one-sided games, fairness is mentioned between 15 and 25% of the time when the equal allocation is efficient (Games 12–15, 18 and 19), but the frequency is much lower (between 4 and 9%) in Games 16, 17, 20–22, where the equal allocation is inefficient. This suggests that most of our subjects are not strongly inequity averse, since they are not preoccupied with equality of earnings when there is only one unequal allocation available and this unequal allocation offers greater earnings to both subjects. References to fairness also tend to be more often associated with agreement on equality (fairness is discussed in 25.2% of the interactions that ended in an agreement on equality).

Finding 7. *References to fairness are more common in two-sided games. Within one-sided games, references to fairness are more common when the equal allocation is efficient. No such pattern exists for two-sided games.*

Standoffs

Standoffs tend to become more common as the degree of conflict of interest between the two bargainers increases. The comparison across games of the frequency of standoffs follows the same pattern as the comparison of disagreement rates.

It is worth noting that standoff-equal is more frequent than standoff in some two-sided games where the total earnings of the equal contract are similar to those of the unequal contract (Game 1, where the equal contract is total-earnings efficient, and Games 2 and 7). This is consistent with

the equal contract losing its focality only gradually as its total earnings decrease.

Finally, (both symmetric and asymmetric) standoffs in general tends to be linked to disagreement (44% of the chat conversations that ended in disagreement were classified as *Standoff*, and 27.5% as *Standoff-Equal*).

Generous Proposals

Generous proposals are more common in one-sided games ($p < 0.001$). Of course, a generous proposal in Games 17 and 20–22 is also a self-interested proposal, since the unequal allocation offers more to both subjects. Most of the generous proposals occurred in those games or in Game 16, where the equal and the unequal allocation give the same payoff to the proposer.

Generous proposals are important for agreements on inequality, particularly in one-sided games, where 30.34% of agreements on inequality follow this category of conversation.

Even though generous proposals are less frequent in two-sided games, they still account for 13.6% of agreements on inequality. If we add the frequency of agreement on inequality following a generous proposal to those following RPS (there is only one situation which has been coded as both by one coder) we find that around 60% of agreements on inequality follow a generous proposal or a game of RPS. Hence, “tough” protracted bargaining is relatively uncommon in our data.

Supergame

This category appears more frequently in games with a lot of conflict of interest between the players (Games 4, 5, 8, 11 and 15). It is disproportionately common in games that ended in disagreement (in 18.8% of the cases), suggesting that communication related to supergame reasoning was not a good strategy on average.

6 Conclusion

People often need to bargain in order to reach a joint decision, but it may not be possible to find an agreement that is both equal and total-earnings maximizing, or even efficient. In such situations, how do bargainers trade-off these properties? Do they tend to settle on an equal and efficient (or per-

haps even inefficient) contract, or rather on an unequal and total earnings maximizing contract?

We report on the findings from experiments where subjects are free to make proposals, can communicate, and sign binding agreements, and we vary the severity of the trade-off between efficiency and equality, in order to understand how the typical focal agreement varies with the efficiency-equality trade-off.

The data show that an inefficient contract is almost never agreed on. Second, equality and efficiency together ensure strong focality only if the total earnings are also sufficiently high; otherwise many bargainers settle on an unequal and total-earnings maximizing contract. Moreover, the data indicate that the bargainers are more occupied with maximizing their own monetary payoff, and possibly to some extent that of their counterpart, and less with ensuring equality of money earnings.⁴⁷ Finally, the data reveal that equality of earnings gets its focality from two sources, namely its absolute property of offering equal earnings, and from being a compromise between unequal contracts over which there is a conflict of interest. The second of these properties results in a systematic failure of the Independence of Irrelevant Alternatives (IIA) axiom.

Our results also have implications for bargaining theory. We reject the assumption of selfish preferences, regardless of the bargaining solution (Nash or Kalai-Smorodinsky). Once we allow for preference heterogeneity, there is considerable agreement between the theories as to the direction of the comparative static predictions, and our data are broadly in line with those. Nevertheless, even if we allow for heterogeneity, the violation of IIA indicates that there are systematic deviations from the predictions of the Nash Bargaining Solution even under very weak assumptions about preferences.

Our laboratory setting allowed us to study, in a controlled and system-

⁴⁷For example, our data for the one-sided game (50,50), (40,240) reveal that around 60% of pairs settled on (40,240). Since any subject can unilaterally achieve (0,0) by refusing to agree, this implies that 60% of subjects prefer the disadvantageous allocation (40,240) to (0,0) and thus are only mildly inequity averse if at all (in the Fehr-Schmidt specification, $\alpha_i < 1/5$).

atic way, how bargainers trade off equality, efficiency, and total-earnings maximization. A potential drawback (shared with many other lab experiments) is a lack of external validity (for a recent debate see Fréchette and Schotter, 2015). Future studies can test whether our results hold with non-student populations – especially those more familiar with negotiations (e.g. business men, salespeople, union negotiators) –, they can increase the stakes of the experiment or consider bargaining environments with more contracts, where the assessment of the efficiency-equality trade-off might be more demanding.

References

- Anbarci, N. and N. Feltovich, 2013, How sensitive are bargaining outcomes to changes in disagreement payoffs?, *Experimental Economics* 16, 560–596.
- Bolton, G. E. and A. Ockenfels, 2000, Erc: A theory of equity, reciprocity and competition, *The American Economic Review* 90, 166–193.
- Bone, J., J. D. Hey, and J. Suckling, 2014, What price compromise?, *Theory and Decision* 77, 359–376.
- Brandts, J. and D. J. Cooper, 2007, It's what you say, not what you pay: An experimental study of manager-employee relationships in overcoming coordination failure, *Journal of the European Economic Association* 5, 1223–1268.
- Camerer, C., 2003, *Behavioral game theory: experiments in strategic interaction*, The Roundtable series in behavioral economics (Russell Sage Foundation).
- Camerer, C. F., G. Nave, and A. Smith, 2017, Dynamic unstructured bargaining with private information: theory, experiment, and outcome prediction via machine learning, *Management Science* .
- Charness, G. and M. Rabin, 2002, Understanding social preferences with simple tests, *Quarterly Journal of Economics* 117, 817–869.
- de Clippel, G. and K. Eliaz, 2012, Reason-based choice: a bargaining rationale for the attraction and compromise effects, *Theoretical Economics* 7, 125–162.

- Cohen, J., 1960, A coefficient of agreement for nominal scales, *Educational and Psychological Measurement* 20, 37–46.
- Conley, J. P. and S. Wilkie, 1996, An extension of the nash bargaining solution to nonconvex problems, *Games and Economic Behavior* 13, 26–38.
- Cooper, D. J. and K.-U. Kühn, 2014, Communication, renegotiation, and the scope for collusion, *American Economic Journal: Microeconomics* 6, 247–78.
- Crawford, V. P., U. Gneezy, and Y. Rottenstreich, 2008, The power of focal points is limited: Even minute payoff asymmetry may yield large coordination failures, *American Economic Review* 98, 1443–1558.
- Embrey, M., K. Hyndman, and A. Riedl, 2016, Bargaining with a residual claimant: An experimental study, working paper.
- Engelmann, D. and M. Strobel, 2004, Inequality aversion, efficiency, and maximin preferences in simple distribution experiments, *The American Economic Review* 94, 857–869.
- Fehr, E. and K. M. Schmidt, 1999, A theory of fairness, competition, and cooperation, *Quarterly Journal of Economics* 114, 817–868.
- Feltovich, N. and J. Swierzbinski, 2011, The role of strategic uncertainty in games: An experimental study of cheap talk and contracts in the nash demand game, *European Economic Review* 55, 554–574.
- Fischbacher, U., 2007, z-tree: Zurich toolbox for ready-made economic experiments, *Experimental Economics* 10, 171–178.
- Fréchette, G. and A. Schotter, 2015, *Handbook of Experimental Economic Methodology*, The Handbooks in Economic Methodologies Series (Oxford University Press).
- Gächter, S. and A. Riedl, 2005, Moral property rights in bargaining with infeasible claims, *Management Science* 51, 249–263.
- Greiner, B., 2015, Subject pool recruitment procedures: organizing experiments with orsee, *Journal of the Economic Science Association* 1, 114–125.
- Herreiner, D. K. and C. Puppe, 2010, Inequality aversion and efficiency with ordinal and cardinal social preferences – an experimental study, *Journal of Economic Behavior and Organization* 76, 238–253.

- Isoni, A., A. Poulsen, R. Sugden, and K. Tsutsui, 2013, Focal points in tacit bargaining situations: Experimental evidence, *The European Economic Review* 59, 167–188.
- Isoni, A., A. Poulsen, R. Sugden, and K. Tsutsui, 2014, Efficiency, equality, and labels: Experimental evidence on focal points in explicit bargaining situations, *The American Economic Review* 104, 3256–3287.
- Kalai, E. and M. Smorodinsky, 1975, Other solutions to Nash’s bargaining problem, *Econometrica* 43, 513–518.
- Karagözoğlu, E. and A. Riedl, 2015, Performance information, production uncertainty, and subjective entitlements in bargaining, *Management Science* 61, 2611–2626.
- Mariotti, M., 1998, Nash bargaining theory when the number of alternatives can be finite, *Social Choice and Welfare* 15, 413–421.
- Mehta, J., C. Starmer, and R. Sugden, 1994, Focal points in pure coordination games, *Theory and Decision* 36, 163–185.
- Myerson, R. B., 1991, *Game Theory – Analysis of Conflict* (Harvard University Press).
- Nagahisa, R. and M. Tanaka, 2002, An axiomatization of the kalai-smorodinsky solution when the feasible sets can be finite, *Social Choice and Welfare* 19, 751–761.
- Nash, J., 1950, The bargaining problem, *Econometrica* 18, 155–162.
- Nydegger, R. and G. Owen, 1975, Two person bargaining: An experimental test of the nash axioms, *International Journal of Game Theory* 3, 239–249.
- Roth, A. E., 1995, Bargaining experiments, in: J. H. Kagel and A. E. Roth, eds., *Handbook of Experimental Economics* (Princeton University Press), 253–348.
- Roth, A. E. and M. W. K. Malouf, 1979, Game-theoretic models and the role of information in bargaining, *Psychological Review* 86, 574–594.
- Roth, A. E. and J. K. Murnighan, 1982, The role of information in bargaining: An experimental study, *Econometrica* 50, 1123–1142.
- Roth, A. E., J. K. Murnighan, and F. Schoumaker, 1988, The deadline effect in bargaining: Some experimental evidence, *The American Economic Review* 78, 806–823.

Schelling, T. C., 1960, *The Strategy of Conflict* (Harvard University Press).

Thomson, W., 1994, Cooperative models of bargaining, in: R. Aumann and S. Hart, eds., *Handbook of Game Theory with Economic Applications*, vol. 2 of *Handbook of Game Theory with Economic Applications*, chap. 35 (Elsevier), 1237–1284.