CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

Discussion Paper No. 2015-18

Fabio Galeotti, Maria Montero and Anders Poulsen

Efficiency versus Equality in Bargaining

November 2015

CeDEx Discussion Paper Series
ISSN 1749 - 3293

CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

The Centre for Decision Research and Experimental Economics was founded in 2000, and is based in the School of Economics at the University of Nottingham.

The focus for the Centre is research into individual and strategic decision-making using a combination of theoretical and experimental methods. On the theory side, members of the Centre investigate individual choice under uncertainty, cooperative and non-cooperative game theory, as well as theories of psychology, bounded rationality and evolutionary game theory. Members of the Centre have applied experimental methods in the fields of public economics, individual choice under risk and uncertainty, strategic interaction, and the performance of auctions, markets and other economic institutions. Much of the Centre's research involves collaborative projects with researchers from other departments in the UK and overseas.

Please visit http://www.nottingham.ac.uk/cedex for more information about the Centre or contact

Suzanne Robey
Centre for Decision Research and Experimental Economics
School of Economics
University of Nottingham
University Park
Nottingham
NG7 2RD
Tel: +44 (0)115 95 14763
Fax: +44 (0) 115 95 14159
suzanne.robey@nottingham.ac.uk

The full list of CeDEx Discussion Papers is available at

http://www.nottingham.ac.uk/cedex/publications/discussion-papers/index.aspx

Efficiency versus Equality in Bargaining

Fabio Galeotti* Maria Montero† Anders Poulsen‡

November 10, 2015

Abstract

We report experimental data from bargaining situations where bargainers can make proposals as often and whenever they want, and can communicate via written messages. We vary the set of feasible contracts, thereby allowing us to assess the focality of three properties of bargaining outcomes: equality, Pareto efficiency, and total earnings maximization. Our main findings are that subjects avoid an equal earnings contract if it is Pareto inefficient; a large proportion of bargaining pairs avoid equal and Pareto efficient contracts in favor of unequal and total earnings maximizing contracts, and this proportion increases when unequal contracts offer larger earnings to one of the players, even though this implies higher inequality. Finally, observed behavior violates the Independence of Irrelevant Alternatives axiom, a result we attribute to a 'compromise effect'.

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

JEL Classification: C70, C72, C92.

---

*University of Lyon, Lyon, F-69007, France; CNRS, GATE, Ecully, F-69130, 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 Werner Güth, Emin Karagözoğlu, Martin Sefton, and James M. Watson, as well as seminar audiences at UEA, the Annual NIBS meeting, the META workshop in Nottingham, GATE-LSE and UPV/EHU, for helpful comments. All errors are ours. Support from the Network for Integrated Behavioural Science (NIBS) and the British Academy is gratefully acknowledged.

1 Introduction

> 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)

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.

It is important for economists and other social scientists to predict which agreement, if any, the bargainers will reach. Will it be one that equates earnings? Will it be efficient, or maximize total earnings? 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, Pareto efficiency, and total-earnings maximization. For simplicity, we refer to the latter as 'total-earnings efficiency' and to Pareto efficiency as 'efficiency'.[1]

Our central research questions are: Which agreements are focal when there is a *tradeoff* between different welfare properties? And how does the focal agreement depend on and vary with *changes* in the terms of the tradeoff? Although these questions seem both foundational and relevant to real world bargaining situations, no systematic investigation has, as far as we know, taken place.

To answer these questions, we design an experiment where pairs of subjects negotiate over a given set of feasible *contracts*. A contract specifies an amount of money to each person within the pair. If the subjects agree on a contract, they get the implied money; otherwise neither person gets any money. We use an unstructured bargaining protocol that allows bargainers to make as many proposals as they wish within a certain period of time.[2] They can communicate via chat, and any agreement is binding.

---

[1]We define Pareto efficiency 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.

[2]See, for example, Anbarci and Feltovich, 2013; Gächter and Riedl, 2005; Isoni et al., 2014; Roth and Malouf, 1979; Roth and Murnighan, 1982, and the surveys in Camerer,

We systematically vary the set of feasible contracts, and this allows us to assess the focality of contract properties such as equality, efficiency, and total payoff maximization.[3]

The advantage of using a free form unstructured bargaining protocol, apart from its inherent realism, is that it makes bargainers strategically equal (see also the discussion in Gächter and Riedl, 2005). Our data on how the efficiency-equality tradeoff influences focality are thus not contaminated by effects due to asymmetries in the bargaining protocol. We see this is as a crucial advantage over using simpler and more restrictive bargaining protocols, such as the ultimatum game (Güth et al., 1982), where the Responder's rejection of a low offer can be driven both by a preference for an equal zero-zero payoff allocation over an unequal one, or by an adverse reaction to being in a strategically inferior position relative to the Proposer.

Our main findings are as follows. First, in the benchmark case where there is an equal 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. Second, we observe a strong tendency for bargainers to avoid the equal contract when it is inefficient.[4] Third, 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

2003; Roth, 1995.

[3]If transfers are allowed, players could, in principle, achieve equality, efficiency and total-earnings efficiency simultaneously by agreeing on actions that maximize the size of the 'cake' and on any transfers needed in order to equate earnings. In many real world situations, there are however financial, technological, informational, or legal constraints on the transfers that can be made, and, therefore, an equal-earnings contract may not be total-earnings efficient, and may even fail to be efficient. Examples are the division of a deceased person's estate among the heirs, or creditors dividing the assets from a bankrupt company. In the estate example, 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 contract can be efficient but not total-earnings efficient, or even inefficient. A similar example is where two partners dissolve a firm, or a couple divorce.

[4]We should bear in mind that we measure Pareto efficiency in money terms, and our subjects may have additional concerns other than the maximization of money earnings. Still, our findings that most subjects disregard an equal and inefficient contract appear even stronger when allowing for this possibility.

some of our bargaining situations, more (less) than half of bargaining pairs settle on an unequal (equal) contract. We are, to the best of our knowledge, the first to find that an equal and efficient but not total-earnings efficient contract can be less focal than an unequal and total-earnings efficient contract.

The fourth main result concerns the comparative statics effects on focality of increasing the conflict of interest over the unequal contracts. We keep the equal contract payoffs fixed and raise the payoff that each unequal contract offers to the player who gets his or her largest money earnings from it. Thus conflict of interest over which unequal contract to settle on increases, but total earnings also increase. The data show that, when the equal contract is efficient, this results in a decrease (increase) in the focality of the equal (unequal) contracts(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 fifth main finding concerns cheap talk communication. Within the class of bargaining situations with an equal but not total-earnings efficient allocation and two unequal and total-earnings efficient allocations, many bargaining pairs invent and agree to use a randomization device, typically based on rock-paper-scissors, to decide which unequal allocation they should agree on. Although informal and non-binding, the 'recommendation' by this rock-paper-scissors device is almost always followed by the bargainers. Interestingly, in bargaining situations where there is only a single unequal contract alternative to an equal and efficient allocation, rock-paper-scissors is almost never used.

The last main finding is that in most 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. This amounts to a violation of the Independence of Irrelevant Alternatives (IIA) axiom (Nash, 1950). This important axiom plays an important role in cooperative bargaining and social choice theory and 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.[5] As we explain below, the violation happens because

[5] A 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 $T \subseteq S$. Then, if x is agreed on when the set is S, x is also agreed on when the set is T.

an agreement can be focal not only due to its absolute payoff properties, but also from being a *compromise* between other feasible agreements. IIA fails to capture the latter, contextual, source of focality, and in our data this effect is sufficiently strong to violate IIA.

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. The data are presented and analyzed in Section 4. Sections 5 and 6 examine the dynamics of contract proposals and cheap talk communication. Section 7 concludes.

2 Related literature

Existing empirical bargaining research (for surveys, see Camerer, 2003; Roth, 1995; specific studies include Camerer et al., 2015; Feltovich and Swierzbinski, 2011; Fouraker and Siegel, 1963; Gächter and Riedl, 2005; Herreiner and Puppe, 2010; Isoni et al., 2013, 2014; Karagözoglu and Bolton, 2015; Karagözoglu and Riedl, 2015; Nydegger and Owen, 1975; Roth, 1995; Roth and Murnighan, 1982; Roth et al., 1988) has done very little work on investigating how focality varies with changes in the tradeoffs between equality and efficiency. In these studies the equal-earnings outcome is either total-earnings efficient (so there is no tradeoff)[6] or it is efficient but not total-earnings efficient, but the equality-efficiency tradeoff is not varied. The only studies that we are aware of that vary the tradeoff are Herreiner and Puppe (2010), Isoni et al. (2014), and Roth et al. (1988).

Isoni et al. (2014) consider unstructured bargaining situations where subjects claim valuable assets placed on a 'bargaining table'. In their environment not only payoff divisions but also the locations of the assets on the table can be focal (the latter is a purely contextual and payoff-irrelevant feature, but can still be a source of focality – see Crawford et al., 2008; Isoni et al., 2013; Mehta et al., 1994; Schelling, 1960). There are two important differences between our and their work. First, and crucially, in their set-up there are only two possibilities for the efficiency–equality tradeoff: *either*

[6]Note 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.

an equal earnings contract is inefficient, *or* it is total-earnings efficient and hence efficient. Indeed, in Isoni et al's set-up, efficiency is *equivalent* to total-earnings maximization. The equality-efficiency tradeoff in their bargaining environment is thus either very severe, or there is no tradeoff at all. The 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. Second, in Isoni et al's design it was not possible for subjects to send messages to each other.

Roth et al. (1988) (RMS) 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 in an agreement determines his or her probability of winning a personal money prize.[7]

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 one hundred lottery tickets) and an unequal pair that offers larger total earnings (generated by an equal division of lottery tickets).[8] RMS vary the slope of the Pareto frontier, and hence the efficiency-equality tradeoff, by simultaneously reducing the low prize and increasing the high prize.[9] This is observed to make the equal allocation less focal. However, this response can be due to a combination of i) the equal allocation being less focal due to its lower payoffs, and/or ii) unequal allocations becoming relatively more focal due to higher total payouts. Both would work in the direction of making the equal allocation less focal. 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 payouts of the equal

[7] If players are self-interested and the axioms of expected utility theory hold, this payment procedure ensures that players are risk neutral in the quantity of lottery tickets they get. This procedure was used by RMS and earlier papers, such as Roth and Malouf (1979) and Roth and Murnighan (1982), since it allows for a test of the Nash Bargaining Solution (Nash, 1950). See Roth (1995) for details.

[8] For example, in RMS's (5,20) game Player 1 (2) has a 5$ (20$) prize, so equating expected money earnings requires a 80-20 split of the lottery tickets; this gives expected money earnings (4,4). A 50-50 split of the tickets produces unequal but larger total earnings of (2.5,10).

[9] More precisely, they collect data for four games, with (Low, High) prizes equal to (10,15), (6,14), (5,20), and (4,36).

allocation.[10]

Also, RMS only study what we call 'one-sided games', namely bargaining situations where there is only a single potentially focal alternative to the equal allocation; we also consider 'two-sided' situations where there are two unequal alternatives to the equal allocation, each preferred by one of the players, and hence a conflict of interest over which one of these (if any) to agree on. We think the latter are as economically relevant as the former, and as we explain below a comparison across them allows us to observe an important behavioral 'compromise effect' that amounts to a violation of the independence of the IIA axiom. Last but not least RMS do not report data on any of the messages that players exchanged – in our analysis we analyze the content of messages, and this allows us to better understand the process through which an agreement, if any, was reached.

In the experiment by Herreiner and Puppe (2010) subjects negotiate for 10 minutes over how to divide four indivisible objects. Each bargainer has a monetary value associated with each bundle of goods, hence bargainers negotiate over which of sixteen payoff pairs to agree on (corresponding to the sixteen different divisions of the four 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 across games 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. We think this makes our experimental design more suited for an investigation of the effects of changes in the efficiency-equality tradeoff on bargaining behavior. Finally, another important feature of Herreiner and Puppe's design is that an equal and efficient contract is very close to also being total-earnings maximizing; thus their design does not allow one to clearly assess the focality of equal earnings contracts that are efficient but far from being total-earnings maximizing.

[10]We do not use the binary lottery method but instead let subjects bargain directly over deterministic money amounts. In our view, the simplicity of this direct procedure more than outweighs the theoretical advantages of the binary lottery method mentioned in Footnote 7. Indeed, Selten et al. (1999) observe that the binary lottery method makes subjects more rather than less sensitive to risk.

3 Experimental Design and Bargaining Games

3.1 Design and Procedures

The experiment was conducted using the experimental lab of the Centre for Behavioural and Experimental Social Science, at the University of East Anglia (Norwich, United Kingdom). We ran 8 sessions with 14 participants each, making a total of 112 subjects.[11] The experiment was programmed and conducted with the software z-Tree (Fischbacher, 2007); recruitment was done using ORSEE (Greiner, 2004). 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 Online Appendix A.1) 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), 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).

The two matched subjects were referred to as Persons 1 and 2.[12] A subject was informed that he or she would in some scenarios be referred to as Person 1, and in others as Person 2.[13]

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

[11] Background information on the participants is in Online Appendix D.2.

[12] The 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 Online Appendix B.1).

[13] The 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.

Figure 3.1 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.[14] 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, the subject would not receive any money earnings. Subjects could make proposals without sending messages, and vice versa.

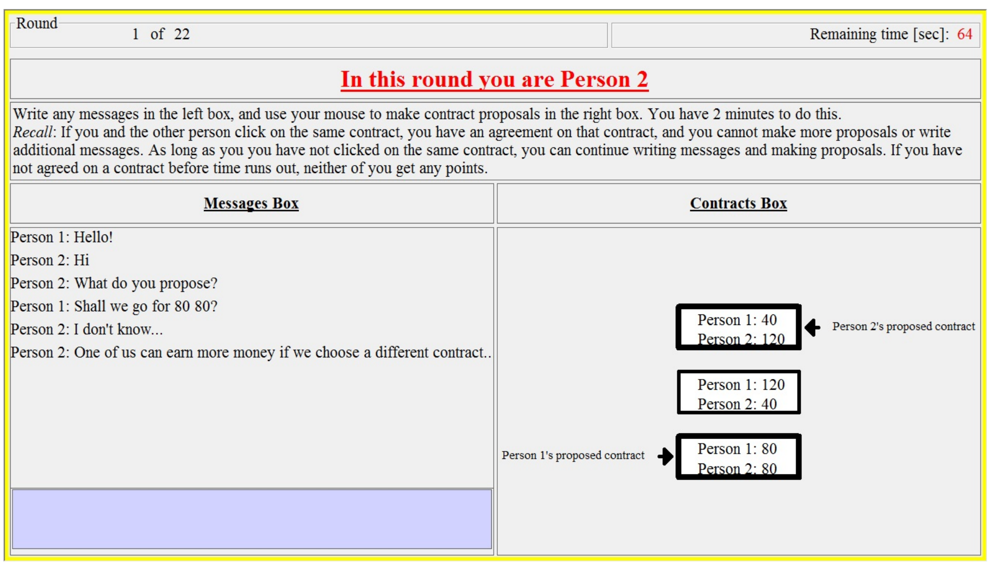

Figure 1: Screenshot.

The content of the 22 bargaining scenarios was not known to subjects

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

in advance. Subjects were informed that they would play the scenarios against different co-participants, and that different subjects would encounter the scenarios in a different order.[15]

When everyone had completed the 22 scenarios the computer randomly selected three rounds for payment. 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* (this is equivalent to a scenario in the experiment). We collected data for 22 different bargaining games, shown in Table 1 below. The numbers are measured in experimental points.

Games 1–11 are of the form $\{(z,z),(x,y),(y,x)\}$, where $z > 0$ and $0 < x < y$. 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 $\{(z,z),(x,y)\}$, and are *one-sided* games. Each two-sided game has a corresponding one-sided game.

When $z \leq x$, 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 $z > x$, the equal-earnings outcome is efficient. This applies to Games 1-4, 7-8, 12-15, and 18-19. If $2z \geq x + y$, the equal contract is total-earnings efficient, and thus also efficient. There are two such games, 1 and 12.[16]

3.3 Game Comparisons

Comparing behavior across bargaining games allows us to answer the following questions:

[15]There were some unavoidable constraints on the matching protocol, due to the real-time nature of the bargaining: if a participant encounters a given scenario as the, say seventh in his game sequence, then another subject in the room must also be playing that scenario as her seventh scenario. The matching protocol that we designed maximized the dispersion of matchings subject to these constraints. See Online Appendix C for details.

[16]We only collected data for two games with an equal and total-earnings efficient contract, since we anticipated that the equal-earnings outcome would be strongly salient. The data confirmed this.

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

Table 1: The 22 bargaining games.

Notes: Y = Yes; N = No.

Question 1: How does the focality of the equal contract depend on its efficiency and total-earnings properties?

We compare bargaining behavior in Games 1–6, and Games 12–17. Here we fix two unequal and payoff efficient contracts, $(x, y) = (40, 120)$ and $(y, x) = (120, 40)$ (only $(x, y) = (40, 120)$ for one-sided games), and vary the payoff z, where $z \in \{30, 40, 50, 60, 70, 80\}$.

Question 2: How does the focality of an equal and efficient (but not total-earnings efficient) contract vary with the earnings and inequality properties of the unequal and total-earnings efficient contracts?

We fix an equal contract $(z, z) = (50, 50)$, and consider different unequal contracts, $(40, y)$ and $(y, 40)$, with $y \in \{70, 120, 240\}$. These are Games 4, 7, and 8 (two-sided) and Games 15, 18, and 19 (one-sided).

When y increases, the contract that offers y becomes more attractive to a self-interested subject, and so he or she may bargain more intensely to get an agreement on that contract. We may then expect a lower (higher) proportion of agreements on the equal contract (an unequal contract). Of course, the frequency of disagreement can increase as well. Subjects may also care positively about the other subject's earnings (see Charness and Rabin, 2002; Engelmann and Strobel, 2004). An increase in y then makes a subject more likely to accept an agreement that gives him or her 40 and the other player y, than (50,50). Both self-interest and a positive concern for other people's earnings are therefore expected to make the unequal contracts more focal than the equal contract.[17] On the other hand, when y increases, the unequal contracts become more unequal, and subjects may dislike settling on a contract that gives them less, or more, than the other subject (see Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000). This can make the equal-earnings contract more focal. Thus it is *a priori* unclear what the overall effect of changing the y parameter is, and our experimental data will shed light on this.

[17] Suppose $y = 120$. If players agree on the unequal contract that gives 120 to one player and 40 to the other instead of the equal contract (50,50), one player is in effect sacrificing $50 - 40 = 10$ points in order to benefit the other player by $120 - 50 = 70$ points. Thus the 'benefit–sacrifice ratio' is 7. Equivalently, the price of transferring one point to the other player is $1/7$ points. When y is larger, 240, the benefit–sacrifice ratio increases to 19, and the price falls to $1/19$. Similarly, the opportunity cost of agreeing to the equal contract in terms of total earnings is $110 - 100 = 10$ when $y = 70$, $170 - 100 = 70$ when $y = 120$, and $280 - 100 = 180$ when $y = 240$.

Question 3: Does the focality of equality differ in one- and two-sided games?

We considered both one- and two-sided games because in real-world bargaining the set of feasible contracts may or may not be symmetric. Comparing one- and two-sided games allows for an empirical test of the Independence of Irrelevant Alternatives (IIA) axiom (Nash, 1950). To see this, note that IIA states that if subjects agree on an the equal allocation in, say, game 8:$\{(50,50),(40,240),(240,40)\}$, then they also agree to the equal allocation when playing game 19: $\{(50,50),(40,240)\}$. We therefore consider a finding that significantly more bargaining pairs agree on $(50,50)$ in Game 8 than in 19 to be violation of IIA.

3.4 Discussion of Experimental Design Features

Since we collect data using a within-subject design, there are several potential effects that can introduce dependencies between games.[18] In this section we describe how our design aims to minimize these effects, and argue that the behavioral patterns we would expect if these effects were significant are only marginally present or not found in the data.

3.4.1 Learning effects

Subjects can naturally be expected to learn as they move through the sequence of twenty-two scenarios, and to transfer their experience from early scenarios to later ones. However, since subjects do not encounter the scenarios in the same order and since the order is random for each subject, these effects should not lead to systematic aggregate effects that bias comparisons across scenarios.

3.4.2 Repeated-game effects

By repeated-game effects we mean any behavior in a current scenario that is intended to influence the behavior of co-players in future scenarios. For example, a subject might believe that he or she, by punishing a 'greedy' co-player in the current scenario, can induce future co-players to behave

[18] The alternative method would be to use a between-subject design, where each subject only makes a decision in a single scenario. However, this would require a very large number of subjects, sessions, and research budget.

more co-operatively. Or that he or she, by being 'generous' now can induce generosity from future co-players. We believe our design minimizes such reasoning and behavior. As already mentioned, subjects played different scenarios against different co-participants, making it difficult to reward or punish a specific individual for their behavior in past scenarios.

One way to look for repeated-game effects in the data is to consider the subjects' chat conversations, since it would appear natural for subjects to use the chat messages to reinforce and make any repeated-game reasoning more salient. We analyzed the chat messages and found very few messages where subjects express or suggest repeated-game reasoning.

3.4.3 Super game heuristics

It is conceivable that the subjects could coordinate on a normative 'super game 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 be also total-earnings efficient. Alternatively, they could select an equal-earnings outcome in all scenarios, to make sure that the overall outcome is also equal. If such a rule was used on a large scale by the subject population, we would expect to see approximately the same rate of agreement on an unequal allocation in all our 22 games. We do not observe this in the data – rather, as we shall describe below, behavior varies in a systematic way across games, suggesting that observed behavior is mostly specific to the game that is currently played rather than reflecting some rule or heuristic that is being applied to all the games.

4 Experimental Findings: Bargaining Outcomes

Table 2 shows descriptive statistics for each of the 22 games (the *Game* column).[19] The number of bargaining pairs that played each game is given in the *Obs* column.[20] The table reports for each game the proportions of each

[19]We 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 Online Appendix B.1. We therefore pool all data.

[20]A 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.

of the three possible outcomes of bargaining, namely a disagreement, an agreement on equality, or an agreement on inequality (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−*Rate of disagreement*. The table also shows the average time to agree for the pairs that did reach an agreement (column *Average time to agree*), 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*).

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.

4.1 Agreement Rates

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%.[21]

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). In particular, the agreement rates of Games 1 and 2 are significantly larger than those of Games 3-6 (Wilcoxon signed-rank tests, $p < 0.05$).[22] There is no significant pattern for the corresponding one-sided Games 12–17. We summarize these findings as:

Finding 1. *In two-sided games 1–6, the rate of agreement tends to fall as the equal contract offers lower and lower total earnings and eventually ceases to be efficient. In one-sided games (12–17) there is no significant pattern.*

[21]The rate of disagreement tends to be lower than in other unstructured bargaining experiments, such as Gächter and Riedl (2005) (where the average rate of disagreement is 16.7 %) and Roth and Murnighan (1982) (17%), but quite close to those found in other studies, such as Herreiner and Puppe (2010) (4.7%) and Isoni et al. (2014) (5.3 %).

[22]There is no statistically significant difference between Games 1 and 2, and between Games 3, 4, 5, and 6 respectively ($p > 0.1$).

A natural explanation is that the focality of the equal contract falls when its total payouts drop, and players are then more likely to consider alternative contracts. Since two-sided games present subjects with two such contracts, each of which is preferred by a different player, this increases the conflict of interest and makes it more difficult for subjects to agree. In one-sided games there is only one unequal allocation, so that the conflict of interest is less pronounced.

A comparison of two- and one-sided games shows that there tends to be more disagreement in the two- than in the corresponding one-sided game. Pooling the data of all two-sided and all one-sided games, the agreement rates are significantly higher for one-sided games ($p = 0.022$).

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**

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

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

Finding 2. *The average rate of disagreement in two-sided games is significantly higher than in one-sided games.*

We can also consider how the agreement rate is affected when, keeping the equal contract fixed, we make the unequal contract(s) more unequal.

In two-sided games where 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 (the drop is significant from Game 7 to 4; $p = 0.048$). Something similar happens when the equal contract is inefficient (Games 9-11), where the agreement rate in Game 9 is significantly larger than in Games 10 and 11 ($p < 0.05$). In the corresponding one-sided games (Games 18, 15 and 19 and 20-22) the effect however is not statistically significant. We summarize these findings in:

Finding 3. *In two-sided games, when the unequal contracts pay more to one player and the same to the other, the rate of disagreement increases. No significant effect is found in one-sided games.*

Once more, a natural explanation for these findings is that making the unequal contracts more unequal increases the conflict of interest in two-sided games, and this makes it harder to agree. In one-sided games, there is by design no potential for disagreement over which of the unequal contracts to settle on, and thus there is less disagreement.

4.2 Agreements

4.2.1 Question 1: How Does the Focality of the Equal Contract Vary With its Efficiency and Total-Earnings Properties?

In Games 1–6 and 12–17, more pairs settle on an equal and efficient contract than on an unequal and total-earnings maximizing contract (with the exception of Game 15). As we lower the total payouts of the equal contract there is in both two- and one-sided games a monotonic and statistically significant drop in the rate of agreement on the equal contract.[23] We also observe that as soon as the equal contract becomes inefficient (Games 5

[23]The relationship between the rates of agreement on equality for two-sided games is: Game 1 = Game 2 > Game 3 > Game 4 > Game 5 = Game 6 ($p < 0.05$); while for one-sided games it is: Game 12 = Game 13 > Game 14 > Game 15 > Game 16 > Game 17 ($p < 0.05$).

and 16), there is a dramatic drop in the focality of equality; in Game 5 (16), only 2 out of 55 (3 out of 56) bargaining pairs agree on the equal contract.

Finding 4. ***(Question 1)*** *In two- and one-sided games 1–6 and 12–17, almost all bargaining pairs agree on the equal contract when it is total-earnings efficient. The proportion of agreements on the equal contract falls as its total payouts decrease, but a majority of agreements remain on the equal contract as long as it remains efficient (with the exception of Game 15). As soon as the equal contract ceases to be efficient, the proportion of agreements on the equal contract drops dramatically.*

These findings show 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 (a majority in Game 15) from settling on an unequal allocation.[24]

The additional finding, that an equal but inefficient outcome lacks any salience, is, we think, equally important. It supports the fundamental and well-known claim made by most economists, that Pareto efficiency is a necessary condition for an outcome to be socially acceptable.[25]

4.2.2 Question 2: How is the Focality of the Equal Contract Affected When The Unequal Contract(s) Offer More Money to One of the Players?

Comparing Games 7, 4, 8 and 18, 15, 19 respectively show that when the unequal contracts offer more to one player and the same to the other (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

[24] Our 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."

[25] These findings differ from those reported in Herreiner and Puppe (2010) (cf. their Experiment R3), where there are two Pareto-efficient contracts, (66,40) and (46,75), and an equal and inefficient contract, (45,45). They observe that 22 out of 48 bargaining pairs agree on (45, 45) and 19 on (46, 75). This suggests a greater willingness to sacrifice efficiency for equality than in our experiment.

larger proportion agree on an unequal contract (Game 7 $>$ Game 4 $>$ Game 8, $p < 0.05$; Game 18 $>$ Game 15 $>$ Game 19, $p < 0.05$). If we consider the games where the equal contract is inefficient (Games 9–11 and 20–22), the agreement rate on equality is not significantly affected, but the agreement rate on inequality goes down in games 9–11.

Finding 5. ***(Question 2)*** *Consider the two- and one-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 and 18, 15, 19). In both groups of games, when y increases significantly fewer subjects agree on an equal and efficient contract, and significantly more agree on an unequal contract. When y is sufficiently high, a majority of agreements are on unequal contracts (Games 8 and 15 and 19).*

We can interpret these findings 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. As discussed earlier, we 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 sufficiently strongly. The data suggest that in the population as a whole the first and second effects dominate the third.

We think these findings are interesting for another reason. One might have conjectured that subjects would have responded to the increased conflict of interest over which unequal contract to agree on by being more likely to agree on an equal compromise, as a means to avoid conflict. This is not what the data show. Although the rate of disagreement does increase somewhat as we go from Game 7 to 4 and to 8, it remains quite low. The subjects deal quite well with more intense conflict and become more, not less, willing to settle on unequal outcomes.

When the equal contract is inefficient (Games 9–11 and 20–22), the rates of agreement on inequality are much higher than when the equal contract is efficient. For example, the rate of agreement on inequality in Game 9 is 94.6%, compared with 26.8% in Game 7. Perhaps because the agreement rate on inequality is already very high in Game 9, making the two unequal allocations even more unequal (Games 10 and 11) leads to a decrease of the rate of agreement on unequal contracts and an increase in the disagreement rate ($p > 0.1$). It is worth noting that the rate of agreement on inequality remains high in one-sided Games 20–22, suggesting that the effect observed in two-sided games is due to an increase in conflict over which of the two unequal allocations to choose rather than a dislike of inequality *per se*.

4.2.3 Question 3: Does the Focality of the Equal Contract Differ in Two- and One-Sided Games?

We next compare the focality of equality in two- and the corresponding one-sided game. Recall that Independence of Irrelevant Alternatives (IIA) states 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. This, however, is not what we typically observe.

Finding 6. ***(Question 3)*** *The average rate of agreement on the equal contract is significantly higher in two-sided than in one-sided games ($p = 0.022$).*[26] *It follows that the data violate Independence of Irrelevant Alternatives, as defined in Nash (1950).*

An equal outcome naturally 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. The equal contract can also be focal 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. To the best of our knowledge, we are the first to empirically observe a failure of the IIA axiom in bargaining due to a "compromise effect".[27]

[26]In one-to-one comparisons, 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 ($p = 0.048$).

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

4.3 Total-Earnings Efficiency

There are two reasons why average total-earnings efficiency (TEE) in a given game can be below 100%: either some bargaining pairs disagreed, or they agreed on an outcome that did not maximize total earnings.

Consider first how TEE varies across the games that differ in the equal contract, Games 1–6 and 11–17. As the equal contract becomes less efficient, TEE drops from almost 100% (Games 1 and 12) to around 70% – 80% of efficiency (Games 3, 4, and 14, 15), but then increases again when the equal-earnings contract becomes weakly or strongly dominated (Games 5, 6 and 16,17).[28]

Finding 7. *There is, in both one- and two-sided games 1–6 and 11–17, a u-shaped relationship between the rate of total-earnings efficiency (TEE) and the properties of the equal-earnings contract: TEE initially decreases as the equal contract becomes less than total-earnings efficient, but picks up again once the equal contract becomes inefficient.*

This u-shaped pattern can be explained by observing that when the equal contract is total-earnings maximizing, most bargainers agree on it, so TEE is close to 100%. As the total earnings offered by the equal contract fall there is a tendency to agree on unequal contracts more often. However, since many bargainers still settle on the equal contract as long as it remains efficient, the net effect is that TEE falls. As soon as the equal contract ceases to be efficient, almost all agreements are on an unequal and total-earnings maximizing contract, and this raises TEE again.

TEE also drops if the inequality of the alternative contract increases when the equal contract is efficient but not surplus maximizing (two-sided Games 7, 4, 8, and one-sided Games 18, 15, and 19).[29] The same is true for two-sided games with an inefficient equal contract (Games 9-11).[30]

setting where no contract is both equal and maximizes total earnings. Bone 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.

[28]TEE is significantly larger in Game 1 than Games 2-6 (Wilcoxon signed-rank tests, $p < 0.05$), in Game 2 than 3 and 4 ($p < 0.05$), in Game 5 than 3 and 4 ($p < 0.1$), in Game 6 than 4 ($p = 0.028$), in Game 12 than 14 and 15 (($p < 0.05$), in Game 13 than 14 and 15 ($p < 0.05$), in Game 16 than 13-15 ($p < 0.05$), and in Game 17 than 12-16 ($p < 0.05$).

[29]Game 7 > Game 4 = Game 8 ($p < 0.05$), and Game 18 > Game 15 > Game 19 ($p < 0.05$).

[30]Game 9 > Game 10 = Game 11 ($p < 0.05$), and Game 18 > Game 15 > Game 19 ($p <$

Finding 8. *Total-earnings efficiency decreases both in one- and two-sided games (except one-sided games where the equal contract is inefficient) when the unequal contract(s) offer more money to one subject while keeping the payoff to the other player the same.*

If we pool together the data of two-sided and one-sided games respectively, the percentage of total-earnings efficiency is statistically significantly higher for one-sided games ($p = 0.018$).

Finding 9. *Average total-earnings efficiency in one-sided games is significantly higher than in two-sided games.*

This can be attributed to the fact that bargainers tend to disagree and to settle on the equal contract more often in two-sided games.

4.4 Agreement Times

The agreement time is the number of seconds elapsed before an agreement is reached (cases where no agreement was reached are excluded). Although in our bargaining environment it does not matter when an agreement is reached ('time is not money', except of course if the deadline is exceeded), there are significant differences in agreement times across our games.

Table 2 shows that as the total earnings of an equal and efficient contract fall (Games 1–4, and 12–15), it takes longer to reach an agreement. When the equal contract becomes inefficient (Games 5–6 and 16–17), agreement times remain high in two-sided games but not in one-sided games.[31]

When the unequal contracts offer higher total payouts but also become more unequal (Games 7, 4, 8 and 18, 15, 19), the agreement times in the two-sided games increase significantly (Game 7 $<$ Game 4 $=$ Game 8, $p < 0.05$). For one-sided games there is no clear pattern. When the equal contract is inefficient, agreement times increase in both one-sided and two-sided games as the unequal contract becomes more unequal (Game 9 $<$ Game 10 $=$ Game 11, and Game 20 $=$ Game 21 $<$ Game 22, $p < 0.1$).

0.05). For one-sided games (Games 20-22), there is no significant difference in the rate of total-earnings efficiency between the games.

[31]For two-sided games, Game 1 $=$ Game 2 $=$ Game 3 $<$ Game 4 $=$ Game 5, Game 4 $=$ Game 6, and Game 5 $>$ Game 6 ($p < 0.05$). For one-sided games, Game 12 $=$ Game 16 $=$ Game 17 $<$ Game 13 $=$ Game 14 $=$ Game 15 ($p < 0.05$).

These findings appear intuitive – a more pronounced conflict over which unequal contract the players should settle on should make negotiations last longer. For one-sided games, increased agreement times for the game {(50,50),(60,240)} suggest that at least some subjects care about equality.

Comparing one- and two-sided games reveals that the agreement time is longer for two-sided games (p = 0.018).

We also observe that it takes significantly less time to agree on an equal than on an unequal contract (44.23 on an equal and 57.79 seconds on an unequal contract, p = 0.018). If we consider two- and one-sided game games separately, the averages are 48.34 and 74.40 (p=0.018) for two-sided, and 43.77 and 57.79 (p=0.398) for one-sided games. The (in)significant difference for (one-) two-sided games suggests that it is especially the conflict of interest over the two unequal contracts that prolongs the negotiations. The cumulative distributions of agreement times confirm these findings (see Figure 2).

Finding 10. *i) It takes significantly more time to reach an agreement in two- than in one-sided games. ii) It takes significantly more time to agree on an unequal than on an equal contract in two- but not in one-sided games.*

Besides looking at average agreement times, we can investigate the presence of a deadline effect in our data. The 'deadline effect' refers to the finding that in many bargaining experiments there is a surge in the number of agreements as the deadline approaches. For example, in Roth and Murnighan (1982), where subjects negotiate for twelve minutes, one-third of agreements are sealed during the last thirty seconds. See Roth et al. (1988) for a description of the findings from other experiments.

Figure 2 shows a steeper slope of the cumulative distributions of agreement times towards the end, suggesting a deadline effect. If we consider agreements reached during the last ten seconds, the most pronounced deadline effects are observed in games 5, 8-11 and 14, where more than 20% of agreements are reached during the last ten seconds (see Online Appendix B.5 for details). These are games where there is either a quite high conflict over the two unequal contracts (Games 5, 8–11), or where players disagree on whether they should settle on an efficient equal or on an unequal and total-earnings maximizing contract (Game 14). This supports the hypothesis that any deadline effect, and, more generally, late agreements, are primarily due to players engaging in a 'chicken type' of bargaining over which of two contracts the players should agree on. This

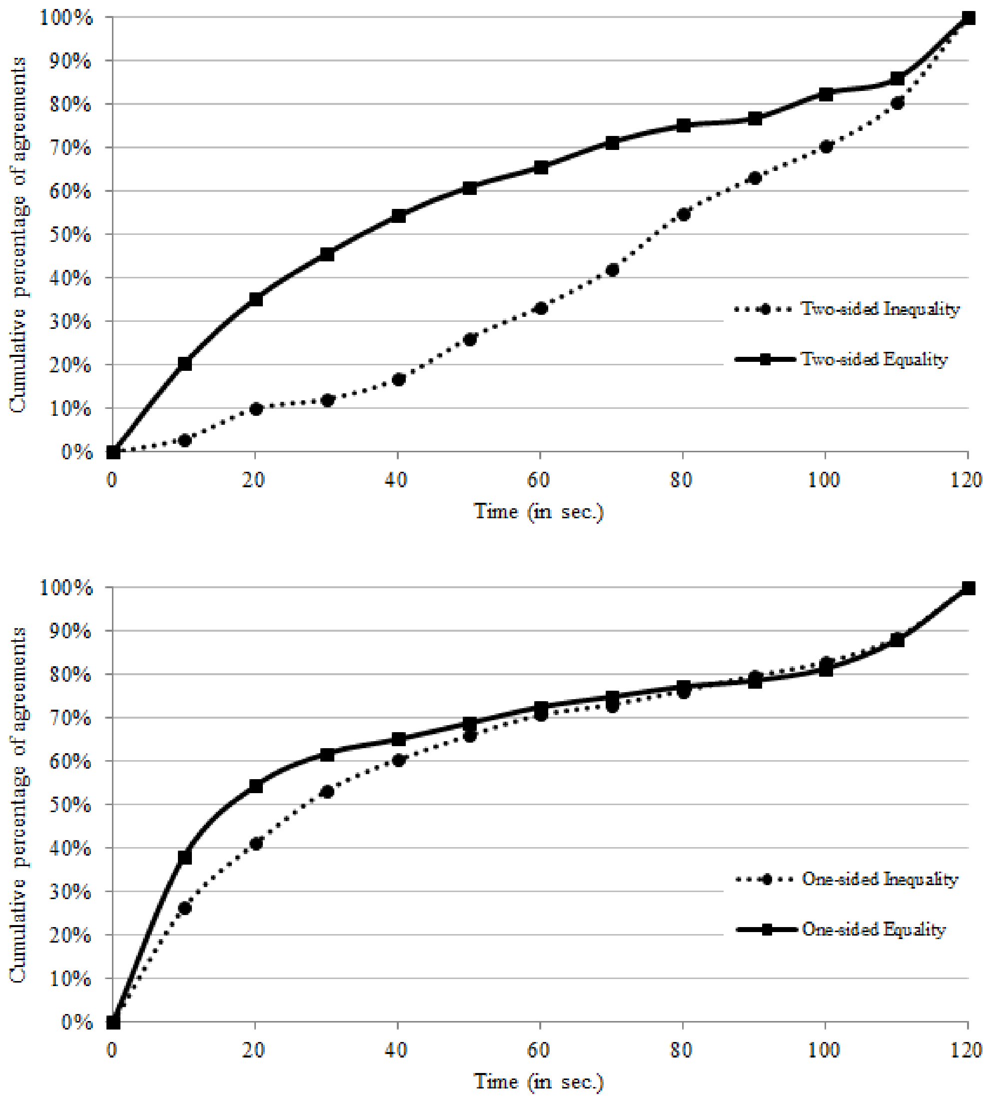

Figure 2: Cumulative distributions of agreement times: equal and unequal agreement times, in two- (top panel) and one-sided games (lower panel).

is also consistent with the games used and agreement times observed in

Isoni et al. (2014).[32]

5 Experimental Findings: Contract Proposals and Messages

In order to better understand the factors that shape the bargaining outcome, we analyzed the contract proposals and messages.

To begin with, we can look at how many contract proposals subjects make and how many messages they exchange. In what follows, 'contract proposal', or just 'proposal', refers to a subject clicking on a contract. In the majority of cases (78.88%), subjects make only one contract proposal during a bargaining period. In a very small number of cases (0.93%) subjects make more than three contract proposals.[33] If we look at the number of messages exchanged, subjects mostly send at least one message – typically between one and four messages – and in only 32.03% of cases they do not send any messages at all. Within each pair of bargainers, there is a strong and significant correlation between the number of messages sent by one bargainer and the other (Spearman $\rho = 0.864$, $p = 0.016$). In contrast, there is no correlation between the number of contract proposals made by one subject and those made by his/her matching partner ($\rho = 0.038$, $p > 0.999$).

5.1 Timing of Moves and Payoffs

Since our experiment has an endogenous timing of moves, it is interesting to investigate whether subjects start the bargaining by sending a message or by making a proposal. Overall, 78.51% of interactions started with a proposal (75.04% in two-sided games and 81.95% in one-sided games). We may also ask whether the order of moves matters for the final agreement. Interactions that started with a message resulted, although not significantly so (Wilcoxon signed-rank test, $p = 0.128$), in higher total earnings than interactions that started with a proposal (156.91 points versus 142.21 points in two-sided games; 166.85 points versus 158.02 points in one-sided

[32] In Online Appendix B.5 we compare our agreement time data with those in Roth and Murnighan (1982) and Isoni et al. (2014).

[33] Recall that a subject can switch back to a previously proposed contract.

games). This is primarily because pairs that started bargaining with a message more often agree on an unequal contract than pairs that started with a proposal ($p = 0.042$).[34]

Besides affecting overall earnings, the order of moves may also matter for the distribution of earnings between the two bargainers; in particular, the first player to make a proposal may have a first-mover advantage. We restrict the analysis of the first-mover advantage to interactions that started with a proposal; in interactions that started with a message there are potentially confounding factors, such as subjects using contract clicking to merely ratify an agreement that has already been reached via chat. In the two-sided games that started with a proposal, the first-mover earned on average 75.49 points compared with 66.72 points for the second mover, and the difference is weakly statistically significant (Wilcoxon signed-rank test, $p = 0.091$). In the case of one-sided games, we need to distinguish between the favored subject (the one that gets more than the other in the unequal contract) and the other subject. Subjects in both roles earned more as first movers than as second movers. Specifically, the favored player earned on average 111.90 points as a first mover compared with 98.22 points as a second mover and the difference is weakly statistically significant ($p = 0.091$); the other player earned 54.46 points as a first mover and 50.62 points as a second mover but the difference is not statistically significant ($p = 0.128$).[35]

[34] In two-sided games, pairs who started with a message more often agree on inequality (71.05% of pairs who started with a message agreed on inequality, compared with only 56.89% of pairs who started with a proposal; $p = 0.022$), less often agree on equality (28.95% versus 43.11%; $p = 0.022$) and, although not significantly so, less often disagree (4.61% versus 8.97%; $p = 0.499$). In one-sided games, pairs who started with a message also more often agree on inequality (72.07% versus 63.49%), and less often agree on equality (27.93% versus 36.51%), although none of the differences is statistically significant (p ¿ 0.1). They are also not less likely to disagree (4.50% versus 3.17%; $p = 0.866$).

[35] The overall conclusion would not change if we include all interactions, not just the ones that started with a proposal. Overall, the first bargainer to make a contract proposal earns 77.14 points on average in two-sided games compared to 68.73 points for the other bargainer. In one-sided games, both player types earn more on average when they make the first proposal (111.99 points versus 102.59 points for the favored player; 53.04 points versus 51.24 points for the other player).

5.2 First Moves and Final Agreements

Instead of looking at earnings, we may simply ask how often the final agreement coincides with the first proposal.

Two-Sided Games In the two-sided games that started with a proposal, the first mover proposed the unequal allocation that favored them in 69.58% of cases, the equal allocation in 28.45% of cases and the disadvantageous unequal allocation only in 1.97% of cases. Proposals of the equal division were eventually implemented 80.77% of the time, but only 38.68% of proposals of advantageous inequality were implemented – instead, the final outcome was the equal allocation in 28.93% of cases, the other unequal allocation in 20.44% of cases, and disagreement in 11.95% of cases. These findings suggest that the greater frequency of equal divisions in pairs that started with a proposal is not due to first movers disproportionately favoring the equal division (the first mover proposes the equal allocation in 28.45% of the cases, and this proportion mirrors the proportion of equal divisions agreed in the games that started with a message).

One-Sided Games We will distinguish between games where equality is efficient (12-15 and 18-19) and games where equality is inefficient (16-17 and 20-22). In the first case, selfish players would each prefer one of the contracts; in the second case, selfish players would both prefer the unequal allocation (weakly in the case of game 16).

In one-sided games with efficient equality, 81.85% of interactions started with a proposal (in 47.27% of those cases, the first proposal came from the favored player). Most proposals by the first mover are of the contract that gives them the most points (69.23% of first proposals by the favored player correspond to the unequal contract, and 95.86% of first proposals by the non-favored player correspond to the equal contract). All unequal proposals from the non-favored player were accepted, but only 54.44% of unequal proposals by the favored player were agreed upon. As for equal proposals, 85% of proposals by the favored player and 78.42% of proposals by the non-favored player were accepted. Overall, the identity of the first mover matters for the final outcome: games where the non-favored player makes the first proposal are more likely to result in an equal agreement (of all interactions that started with a proposal by the non-favored player, 75.17% result in an agreement on the equal contract, compared to

52.31% of interactions that started with a proposal by the favored player; $p = 0.018$).

As for the one-sided games with inefficient equality, 82.08% of bargaining pairs started with a proposal (of those, 62.45% started with a proposal by the favored player). Both types of players are more likely to propose the unequal contract (93.01% of favored players and 83.72% of non-favored players propose the unequal allocation), and over 99% of those proposals are accepted. In contrast, the very few proposals of the equal contract are not always accepted. Since most proposals by the non-favored player are of the unequal allocation, it is clear that many of our subjects are not strongly inequity averse: not only do they prefer the unequal outcome to disagreement (as evidenced by the final agreements data) but many also prefer the unequal allocation with higher earnings to the equal one.

6 Analysis of Chat Conversations

In order to analyze the content of conversations, we followed a similar approach to Brandts and Cooper (2007) and Cooper and Kühn (2014). We hired two coders and asked them to code the conversations separately 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.

The categories are as follows:

Rock-paper-scissors (RPS): the contract agreed upon is determined by the outcome of a game of RPS.[36]

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

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.

[36]This 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, 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.

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).

Sum of earnings: one or both subjects refer to some contracts as offering the largest sum of earnings.

Poor: one or both subjects make statements such as "I'm poor" or "I really need the money".

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

Rules/norms: references to how subjects should behave in general (i.e., in all the games), such as "We should always agree on an equal contract", "We should always agree to a contract that gives the largest sum of earnings", "We should alternate", or similar.

First: statements such as "I clicked on a contract first, so I should get it".

Been unlucky in the past: statements such as "I have been unlucky", "people have been nasty to me" or "I've been generous before".

Nice in the past: a statement such as "People have been nice to me in the past".

If nice now, then...: statements like "If you are nice to me now, then I will be nice to you next time we meet". The subject refers to a future meeting between him/her and *the same* other subject.

Pass it on: "If you are nice to me, I will pass it on" or a similar statement. The difference with the previous category is that the subject does not refer to a future meeting between him/her and the same other subject.

I send you good karma: subjects state that if the other subject is generous, then they will "send them karma" (or similar) such that other people will be generous to them.[37]

[37]The online Urban Dictionary (www.urbandictionary.com) states that karma is "The Buddhist belief that whatever you do comes back to you, e.g. if you do something good, something good will happen to you, and vice versa.".

I spread karma: a subject states that he/she is kind/generous, such that other people will be generous as well.

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.

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% of all games, and this is the value we use.[38] 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.

6.1 Results

The first thing to note is that some categories are much more common overall than others. There 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*. It is also worth noting that, even if we pool all the categories that might be related to supergame reasoning (*rules/norms*, *been unlucky in the past*, *nice in the past*, *if nice now then*, *pass it on*, *I send you good karma* and *I spread karma*) they only add up to 6.5%; of those, the only category that is above 1% is *been unlucky in the past*.

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). The full data set can be found in Online Appendix B.6.

In the remaining part of this section we focus on the five most prominent categories – *RPS*, *fairness/equality*, *standoff*, *standoff-equal* and *generous proposal*. We also consider a hybrid category (henceforth *supergame*) which is constructed by pooling together all the categories that are related

[38]These 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).

to supergame reasoning. For each of these six categories, we investigate in which games it is observed, and 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 in which period the game occurred in the session (period 1 to 22). The results are reported in Online Appendix B.6. 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 ran two multinomial logistic regressions – one for one-sided games and one for two-sided games – on the three possible outcomes of a bargaining situation (disagreement, agreement on equality, agreement on inequality).[39] Independent variables are a dummy for each of the 22 games, the period of the bargaining situation, and a dummy for each possible category of communication which was coded in at least 5% of the chats (we also controlled for situations where subjects did not chat at all). For all the regressions, we employed random effects at session level to control for session effects.[40] Table 4 reports the marginal effects of these regressions.

[39]We ran separate regressions for one- and two-sided games since certain categories of communication occurred only (or mostly) in either one- or two-sided games. For instance, *standoff* could only occur in two-sided games, while *RPS* was barely used in one-sided games. Pooling the data together would result in biased estimations of the effects of these categories.

[40]We also ran regressions, available upon request, where we include session dummy-variables instead of random effects. The results do not change.

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

Table 3: Percentage frequency of the most common chat categories

Notes: The table includes only categories that were assigned to at least 5 percent of the bargaining games (averaging across coders). κ = Cohen's kappa coefficient of agreement (Cohen, 1960) between the two coders. 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. [a] 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.

	Two-sided games			One-sided games		
	Disagreement	Equality	Inequality	Disagreement	Equality	Inequality
RPS	-0.113**	-0.311***	0.424***	-0.472	0.1	0.372
	(0.048)	(0.059)	(0.049)	(346.57)	(99.004)	(247.57)
Fairness	-0.031	0.081***	-0.05	0.005	0.148***	-0.153***
	(0.029)	(0.027)	(0.033)	(0.016)	(0.031)	(0.032)
Standoff	0.102***	-0.011	-0.091***	–	–	–
	(0.025)	(0.026)	(0.028)			
Standoff-equal	0.058*	0.03	-0.089**	0.049***	-0.057**	0.008
	(0.033)	(0.029)	(0.036)	(0.017)	(0.027)	(0.029)
Generous proposal	0.000	-0.290***	0.290***	-0.026	-0.310***	0.336***
	(0.059)	(0.090)	(0.068)	(0.049)	(0.117)	(0.109)
Both better off	-0.065	-0.048	0.113*	0.22	-1.825	1.606
	(0.052)	(0.066)	(0.058)	(38.764)	(330.79)	(292.02)
Cost benefit	-0.03	-0.072*	0.102**	0.061***	-0.169***	0.108***
	(0.041)	(0.041)	(0.042)	(0.021)	(0.038)	(0.041)
Be generous	-0.042	-0.073**	0.115***	-0.02	0.001	0.019
	(0.028)	(0.033)	(0.031)	(0.020)	(0.042)	(0.043)
Supergame	0.035	-0.100***	0.065*	0.045**	-0.023	-0.022
	(0.033)	(0.037)	(0.038)	(0.020)	(0.052)	(0.051)
off-topic	0.023	-0.032	0.009	0.046	-0.008	-0.039
	(0.032)	(0.038)	(0.043)	(0.036)	(0.056)	(0.064)
No chat	-0.062**	0.053*	0.009	-0.040***	0.143	-0.103
	(0.025)	(0.031)	(0.034)	(0.010)	(0.140)	(0.137)
Games and time controls	Yes			Yes		
N	609			615		
ll	-196.385			-169.825		

Table 4: Random-effects multinomial logit regressions

Notes: The table displays marginal effects. Random effects are at session level. Standard errors are in parenthesis. N = number of observations. ll = log likelihood. * $p<0.10$, ** $p<0.05$, *** $p<0.01$

6.1.1 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; $p < 0.001$).[41] 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,

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

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, $p < 0.05$) – 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, $p < 0.05$), 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 increases the likelihood of agreeing on inequality in two-sided games (by 31.10%, $p < 0.01$) (see Table 4), 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 (46.7%) of unequal agreements are the result of using RPS.

Finding 11. *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.*

6.1.2 Fairness

A comparison of each two-sided game with its corresponding one-sided game shows that fairness/equality references tend to be more common in two-sided games ($p = 0.01$); 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 ($p < 0.05$). 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 tends to increase the likelihood of agreeing on equality (as Table 4 indicates, by 8.10% in two-sided games and 14.8% in one-sided games, p < 0.01).

Finding 12. *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.*

6.1.3 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 result in more disagreement ($p < 0.01$). The likelihood of disagreeing increases by up to 16% for standoffs in two-sided games.

6.1.4 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 the probability of agreeing on inequality increases by as much as 33.6% ($p < 0.01$) (see Table 4).

Even though generous proposals are less frequent in two-sided games, they still account for 13.6% of agreements on inequality and contribute to an increase of 29% in the probability of agreeing on inequality ($p < 0.01$). 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.

6.1.5 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, suggesting that communication related to supergame reasoning was not a good strategy on average. In particular, we find that, in one-sided games, other things being equal, the probability of having a disagreement increases by 4.5% in situations which are categorized as *supergame* ($p < 0.05$). In two-sided games, supergame talks reduce, *ceteris paribus*, the probability of agreeing on equality by as much as 10% ($p < 0.01$).

7 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 perhaps 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 systematically vary the nature of the tradeoff between efficiency and equality, in order to understand how the typical focal agreement varies with the efficiency-equality tradeoff.

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 most 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, but less with ensuring equality of money earnings. Finally, the data reveal that equality, when also efficient, 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 axiom Nash (1950).

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, pp. 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., G. Nave, and A. Smith, 2015, Dynamic unstructured bargaining with private information and deadlines: Theory and experiment, working Paper.

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 XX, 37–46.

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.

Engelmann, D. and M. Strobel, 2004, Inequality aversion, efficiency, and maximin preferences in simple distribution experiments, The American Economic Review 94, pp. 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.

Fouraker, L. E. and S. Siegel, 1963, Bargaining Behavior (McGraw-Hill).

Gächter, S. and A. Riedl, 2005, Moral property rights in bargaining with infeasible claims, Management Science 51, 249–263.

Greiner, B., 2004, The online recruitment system orsee 2.0 - a guide for the organization of experiments in economics, Working Paper Series in Economics 10, University of Cologne, Department of Economics.

Güth, W., R. Schmittberger, and B. Schwarze, 1982, An experimental analysis of ultimatum bargaining, Journal of Economic Behavior and Organization 3, 367–388.

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.

Karagözoglu, E. and G. Bolton, 2015, On the influence of hard leverage in a soft leverage bargaining game: The importance of credible claims, working paper.

Karagözoglu, E. and A. Riedl, 2015, Performance information, production uncertainty, and subjective entitlements in bargaining, forthcoming in Management Science.

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).

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).

Selten, R., A. Sadrieh, and K. Abbink, 1999, Money does not induce risk neutral behavior, but binary lotteries do even worse, Theory and Decision 46, 213–252.