Working Paper 10/2015

Fairness preferences in a bilateral trade experiment

A. Ciccone, R. Bråten and
O. Rogeberg

The CREE Centre acknowledges financial support from The Research Council of Norway, University of Oslo and user partners.
ISBN: 978-82-7988-204-6
ISSN: 1892-9680
http://www.cree.uio.no

Abstract in Norwegian:

CREE - Working Paper 10/2015

Ulikhetsaversjon i et tosidig handelseksperiment

A. Ciccone, R. Bråten and O. Rogeberg

Siden søttitallet har adferdsforskerne brukt laboratoriestudier og psykologisk forskning til å påvise mangler i menneskers evne til å innhente og bearbeide informasjon, utforme og gjennomføre planer. De har også vist at vi mennesker ikke bare bryr oss om hvor mye penger vi selv får - men også hvor mye andre får og den relative fordelingen av goder. Mennesker som i tillegg til egen vinning også bryr som om utfall for andre sies å ha sosiale preferanser. Et eksempel på sosiale preferanser er ulikhetsaversjon, som uttrykker et ønske om at goder til en viss grad skal deles likt.

I hvilke situasjoner har disse laboratoriefunnene betydning? Når det gjelder sosiale preferanser tyder noen studier på at vi legger fra oss sosiale preferanser når vi går inn i et marked. En studie viser for eksempel at folk bidrar mindre til fellesskapet hvis sitasjonen kalles et "markedsspill" enn hvis situasjonen kalles et "fellesskapsspill". Det har blitt foreslått at handelssitusjoner får oss til å fokusere kun på egen gevinst, mens det spiller liten eller ingen rolle hvor mye kjøpmannen eller motparten tjener på handelen.

For å undersøke betydningen av handel for sosiale preferanser, utviklet vi forskerne ved Frischsenteret og CREE et eksperiment der deltagere ble delt inn i grupper på to. Hver av deltagerne i et par fikk utdelt ressurser de kunne handle seg imellom til en oppgitt pris eller beholde. På slutten av hver periode fikk de penger etter hvor mange ressurser de hadde og kunne bruke til "produksjon". Den ene parten i et par fikk mange ressurser, men tjente lite på å produsere. Den andre fikk få ressurser, men tjente mye på å produsere. Begge tjente på å handle ressursene seg imellom - og begge tjente mest når den ressursrike parten solgte alle sine ressurser til medspilleren.

Vi ville undersøke hvorvidt folk som stod overfor en slik handelssituasjon brydde seg om at handelen skapte ulikhet. De varierte derfor prisen som partene fikk lov til å handle med: Noen av deltagerparene fikk beskjed om at handel måtte skje til en gitt lav pris. Dette går i kjøperens favør, som dermed fikk mesteparten av gevinsten. Selgeren tjente også på å selge snarere enn å beholde ressursen, men mye mindre enn kjøperen. Andre deltagerpar fikk oppgitt en veldig høy pris, der kjøper tjente lite (men noe) på å kjøpe, mens selger "håvet inn" gevinst. Atter andre par fikk oppgitt en pris som gjorde at begge parter kom akkurat like godt ut av det.

Spillet varte i 10 runder. Deltagerne satt ved hver sin dataskjerm (uten å vite hvem i rommet de spilte mot), og fikk se hvordan ulike handelsnivå ville øke deres og motpartens inntekt. På denne måten ønsket vi å undersøke om folk var villige til å si fra seg sin egen og motpartens gevinst ved å redusere handel, for dermed å redusere skjevheten i inntektsfordelingen.

Vi fant at den totale handelen var 15% lavere blant parene med skjevfordelt gevinst enn blant parene med lik fordeling av handelsgevinsten. Det var først og fremst de som tjente lite på handel som ikke ønsket å handle, mens de som tjente mye stort sett ønsket å handle. Dette tyder på at mange bryr seg om hvor de mye får i forhold til andre, også i en handelssituasjon.

Fairness preferences in a bilateral trade experiment *

A. Ciccone† R. Bråten‡ O. Rogeberg§

Working Paper

Abstract

The importance of social preferences is well documented in the experimental economics literature. Several studies, however, indicate that framing an experiment as involving trade or market institutions substantially reduces the impact of fairness concerns on behavior. We present laboratory evidence of fairness concerns affecting behavior in a bilateral game with a stylized trade setting and a buyer-seller frame. The study examines the extent to which individuals are willing to participate in a Pareto-improving trade, where fixed prices create increasing payoff inequality. Results show that about 60 percent of pairs withhold trade, reducing trade quantities by, on average, 15 percent, as compared to the control treatment, where gains from trade are split equally. We conclude that fairness concerns may affect outcomes and reduce economic efficiency even in a trade game where strategic or retaliation motives for restricting trade are absent.

Keywords: Fairness; Inequity aversion; Trade; Markets; Experiment; Social preferences.

JEL: C90, D02, D03, D63

*__Aknowledgments:__ We gratefully acknowledge Oslo Centre for Research on Environmentally friendly (CREE), the Norwegian Research Council, the Department of Economics (UiO) and the project "International cooperation on emission abatement" (Norwegian Research Council, project number 193703). We wish to thank all those who helped us with feedbacks and valuable comments: Ernesto Reuben, Kjell Arne Brekke, Fortuna Casoria, Alberto Antonioni; and participants in workshops organized by the mentioned NRC project; We also gratefully acknowledge the help of Emily Oswald with language and copy editing.

†University of Oslo, Department of Economics. CREE: Oslo Centre for Research on Environmentally friendly Energy. E-mail: alice.ciccone@econ.uio.no

‡The Ragnar Frisch Centre for Economic Research and Oslo Economics. CREE: Oslo Centre for Research on Environmentally friendly Energy. Oslo. E-mail: rbr@osloeconomics.no

§The Ragnar Frisch Centre for Economic Research. CREE: Oslo Centre for Research on Environmentally friendly Energy. E-mail: ole.rogeberg@frisch.uio.no

1 Introduction

Public concern with fairness in consumer and labor markets has been widely documented by surveys (Kahneman et al. (1986), Fehr and Gächter (2000), Bewley (1999)), boycotts or in increased demand for "fair trade" and charity-linked products (Hainmueller et al. (2014), Elfenbein and McManus (2010), Castaldo et al. (2009) and Elliott and Freeman (2003)). A significant share of consumers report avoiding some products that were perceived to reflect "unfair trade with developing countries" (MMI (2010)) or were willing to pay extra for products made in ethical ways (Elliott and Freeman (2003)). While economists tend to see markets as an efficient mechanism that mutually benefits all participants, others are more inclined to judge some trading prices as potentially unfair. For example, the former would argue that a low-paid job in a developing country may represent an improvement of the worker's condition relative to available alternatives, while the latter might view a low wage as exploitative and unethical. In other words, some people may dislike the perceived unequal or "unfair" distribution of gains from trade so much that they prefer to abstain from trade, even if this imposes an economic cost on the disadvantaged party.

Another source of evidence for the relevance of fairness concerns and social preferences is the field of experimental economics. Despite the abundance of studies finding other-regarding preferences to be a fundamental driver of social interactions (Camerer (2003), Fehr and Schmidt (2006) and Rabin (1993)), many studies have found that fairness concerns are unlikely to affect market outcomes (Roth et al. (1991) and Prasnikar and Roth (1992)).[1] More importantly, previous experimental studies have also found that the expression of fairness and inequity concerns is weakened simply by framing an experiment as involving markets or trade: changing the framing of a public good game to a market game with an automated seller is sufficient to reduce cooperation rates (Cherry et al. (2013)). Similarly, Hoffman et al. (1994) found lower mean offers and rejection rates in an ultimatum game when the game was given a "seller-buyer" context instead of being framed as a simple "split $10" task.[2] The lesson drawn from such findings is that market or trade setting and framing can eradicate or, at least, significantly reduce the impact of social preferences on behavior.

We present experimental evidences of fairness concerns such as inequity aversions in a simple game with a trade setting and a buyer-seller frame. Specifically, we examine the extent to which people are willing to participate in

[1] A general result is that experimental markets quickly converge to the equilibrium predicted by a model with exclusively self-interested players (Smith (1962), Smith (1991), Davis (1993) and Großer and Reuben (2013)).

[2] Outcomes still deviated from the predictions of the standard model.

a fixed trade that makes both players better off in monetary terms, while also resulting in increased payoff inequality. We design a stylized two-person game with imposed terms of trade. High-productivity players with few resources – the buyers – can offer to buy resources from low-productivity players with many resources – the sellers – at fixed and constant prices.[3] Prices are set such that trade is Pareto-improving in monetary terms and both players' payoff increases linearly with trade volumes. The experiment is implemented with a partner matching design where different buyer-seller pairs face different constant prices across treatments. We distinguish between three treatments: (i) the "buyer treatment" has a low price, distributing most of the gains from trade to the buyer; (ii) the "seller treatment" has a high price, distributing most of the gains to the seller; (iii) the "Symmetric Treatment" has a price set to equalize the gains between the parties (see Table 1 for payoff calculations). The experiment is designed to reflect everyday trading contexts where prices are perceived as fixed and exogenously given, with both sellers and buyers acting as price takers. For instance, this set-up can mimic the dynamics by which both low wage workers and supermarket consumers see the wages and prices they face as take-it-or-leave-it offers.

Many experimental studies of other-regarding preference have used institutions where strategic considerations play an important role (Rutström and Williams (2000)).[4] In posted-offer experiments, buyers were found to withhold Pareto-improving trade when posted prices implied an uneven distribution of the gains (Franciosi et al. (1995); Ruffle (2000); Tyran and Engelmann (2005)). However, this effect decreased substantially in the final rounds, suggesting that the reduced trading in earlier rounds partly reflected a strategic motive to force fairer offers in later rounds. This highlights the importance of a design that separates fairness-motivated restriction of trade from profit-oriented strategic action. Disentangling strategic motives from social preferences is also difficult in gift-exchange experiments (Fehr et al. (1993)), where "employees" can punish unfair "wage offers" by providing low effort, thus providing an incentive for higher wage offers from the employer's side. Moreover, comparisons between ultimatum and dictator games show that equality concerns may play a role in both games, but that players are far more generous in the ultimatum game than in the dictator game. Offers are higher in ultimatum games because of the high rates of responders that turn down small offers. Hence, Forsythe et al. (1994) reject the fairness hypothesis as a primary explanation for the generous offers typically observed in ultimatum games. Similar observations in sequential bargaining where

[3] Our game is framed as a bilateral trade game with buyers and sellers, but it is distinct from standard market games since prices are fixed by the experimenters and there is no competition by design. Hence, players have no possibility to affect the terms of trade.

[4] Non-self-interested behavior with non-strategic interactions can be found for instance in dictator games such as Kahneman et al. (1986) and Forsythe et al. (1994).

there is a strategic motive are found in Bolton and Ockenfels (2000). They illustrate the important interplay between equity concerns and strategic considerations, but do not reject the presence of some individuals motivated or influenced by fairness considerations.

Our contribution to the literature is threefold. First, we show that fairness concerns significantly affect behavior in experiments with a trade setting and a buyer-seller frame. Second, we present a clean and simple experimental design where such concerns are separated from strategic incentives or retaliation motives. Finally, we isolate Inequity Aversion as in Fehr and Schmidt (1999) from alternative theories such as preference for efficiency or for the worse-off party. The experiment is created to give no immediate or strategic motive for restricting trade. Specifically, in each round both players suggest how many resources they wish to trade between zero and ten. The realized trade is, then equal to the lowest proposal within a trading pair. Since participants are given no information on the trading preferences of their co-player beyond observed trade in earlier rounds, there is no signalling motive for the non-constraining participant to choose any trading level below what he or she would actually prefer. Also, the distribution of the gain is determined by the constant price each pair faces and cannot be affected in any way. Strategic motives are absent as none of the parties can influence future terms of trade and there is no retaliation motive whereby sellers or buyers would seek to punish their opponent for price-offers. Given this, any trade restriction observed in our experiment should be due to the asymmetry of price treatments and the unequal distribution of gains they imply.[5]

In the proposed experiment, trading prices, productivity parameters and endowments are set such that trade restrictions are predicted under the preference structure and average parameter values proposed for Inequity Aversion in Fehr and Schmidt (1999). It is important to mention the presence of different theories of social preference in the literature. Sometimes, expressions of inequity aversion can be confounded with social welfare preferences or with reciprocity (Charness and Rabin (2002)). Also, inequity aversion and inefficiency aversion are found to be important drivers in other-regarding considerations even though it is not yet clear under what conditions one dominates the other (Engelmann and Strobel (2004)).[6] The present study

[5] An alternative explanation for our results could be that participants act as if they were in a market game, where future trading terms could be influenced by reducing their desired trade, simply because this reflects the habits and rules of thumb that tend to be useful in real world trading contexts. However, such an interpretation would also imply that "irrational" restrictions in trade could take place in real market contexts as well, e.g., when the same habits and rules of thumb are triggered in situations where the trading partner are unlikely to be affected by our individual actions.

[6] For more aspects of this debate, see Bolton and Ockenfels (2000), Bolton and Ockenfels (2006), Engelmann (2012), Engelmann and Strobel (2006), Fehr et al. (2006) and Fehr and Schmidt (2010).

does not aim to validate one theory rather than the other, but to merely test the expression of social preferences in a trade setting. However, in our experiment, theories of social preferences such as maximin preferences and a preference for efficiency predict full trade and cannot hence explain the observed trade restrictions.

Experimental data show that full trade is rapidly implemented in the symmetric treatment. This implies that participants understood the game quite well and that the trading game in itself did not trigger reductions in trading volumes. In contrast, the majority of trading pairs substantially restrict trade in the asymmetric treatments. Specifically, we observe that about 60% of couples restrict trade and the average number of resources traded is significantly lower, by about 15%, in both asymmetric treatments as compared to the symmetric treatment. We run non-parametric tests over completely independent observations to test differences between trade in the asymmetric treatments and in the Symmetric control treatment considering different rounds to avoid possible learning effects. The null hypothesis is always rejected with high statistical significance (p-values between .0015 and .0002). Moreover, we find no significant difference in terms of trade volumes between the buyer and seller treatments, excluding the possibility of a framing effect one party is called the "seller" and the other the "buyer". The trade restrictions observed in the both asymmetric treatments are mostly driven by the worse-off players who systematically restrict trade. Finally we do observe a small number of favored players who actively restrict trade. In conclusion, our results support the hypothesis that unequal distribution of gains from trade may result in reduced trading volumes in a rudimentary trade setting with a buyer-seller frame.

The paper is organized as follows: Section 2 presents the game and gives a detailed description of the experimental design. Section 3 describes the data and reports the results. Finally we conclude by discussing the findings and their implications.

2 Experimental design

A total of 128 students from Columbia University were recruited with the Online Recruitment System for Economic Experiments (ORSEE) and divided in six sessions over different days. The participant group is diverse in terms of gender (55% women), discipline of study and country of origin, although a majority of participants are from the United States. The experiment was programmed with z-tree software (Fischbacher (2007)) and was run in April 2014 at Columbia Experimental Laboratory in the Social Sci-

ences.[7] The experiment used a fixed-partner and between-subjects design. Take home earnings ranged from USD 5.4 to USD 45.6 with an average of USD 24.3, plus a flat show up fee of 5 USD. The payment was made in cash immediately after the experiment as the students left the lab. Subjects were asked to answer a short survey immediately after playing the game. The experiment lasted in total about one hour. At the beginning of each experimental session, a written set of instructions was handed out to the participant and read aloud (reported in the appendix).

2.1 The trade game

Participants were randomly assigned to be *buyers* or *sellers* and then randomly matched into trading pairs that remained fixed for ten identical trading rounds. Subjects were told that they would face a trading price that was given and constant across all rounds, and that this price would determine how the gains of trade were split between the two parties. Players were also given a fixed number of resources as endowment every round, and a productivity parameter that translated resources into the Experimental Currency Unit (ECU). The ECU had a fixed and known exchange rate to US dollars.[8] Buyers converted resources to monetary payoffs with a high productivity, but were given a low initial resource endowment. In each round, they received one resource and at the end of each round they earned 20 ECU for each held resource. The sellers faced the opposite situation, with low productivity and high resource endowments: in each round, they received 20 resources and they earned 1 ECU per resource held at the end of the round. Note that the no-trade payoffs were identical (20 ECU per round) for both buyers and sellers. The total gain from trading was determined by the difference in productivity, in that each resource held at the end of a round was worth 19 ECU more for the buyer than for the seller. Price, resource endowments and productivity for both types was common knowledge.

During the experiment, the buyer-seller couples were free to trade resources at a fixed and given price to increase their payoffs. The instructions specifically reminded them that they could trade resources with their partner "to increase earnings". Trade in each round took place as follows: both players simultaneously stated their desired trading level (ranging from zero to ten units per round). Since trade was voluntary, actual trade was determined by the lowest offer. Before making their choice, both parties were visually informed, with bar graphs and numbers, of their own and their co-player's (1) secured incomes in the present round (i.e. the players' earnings in the absence of trade); (2) possible gains from different levels of trade in that

[7] C.E.L.S.S. Columbia University, New York - USA.
[8] The exchange rate was set to 50 ECU = 1 USD.

round; and (3), from the second round onwards, cumulative payoffs from all previous trading rounds. For players in all treatment groups, this visualization made one aspect of trade particularly relevant to the experiment clear, that is, that trade increased both parties' payoff. The visualization received by players under the asymmetric treatment also made it clear that trade made outcomes highly unequal.

Prices were known, exogenously given, and constant over all rounds within each buyer-seller pair. Each pair faced one of three price treatments. The buyer treatment favors the buyer, with a fixed price of 3 ECU. For each unit sold, the seller made a 2 ECU profit relative to keeping the unit and producing ECUs with a low productivity, while the buyer made a 17 ECU profit by receiving 20 ECU from production minus the purchasing price of 3 ECU. In the treatment favoring the seller – the "seller treatment"– the price was 18 ECU. Each unit traded in this treatment gave the seller a profit of 17 ECU (18 in sales price minus 1 in foregone production income), while the buyer netted a profit of 2 ECU. Finally, in the "symmetric treatment," the price was 10.5 ECU, implying an equal distribution of the gains of trade between the two players. They both earned a profit of 9.5 for each resource traded (the seller earned the sales price 10.5 minus 1 in foregone production and the buyer earned the production "revenue" of 20 minus a purchasing price of 10.5 ECU). Payoffs are shown in Table 1. For instance, in the "buyer treatment", if the buyer suggested full trade while the seller offered to trade only 8 resources, the actual number of resources traded would be 8 and the final payoffs would be 156 ECU for the buyer and 36 ECU for the seller.

Table 1: Payoffs for Buyers and Sellers by Treatments

	Buyer T Price: 3		**Seller T** Price: 18		**Symmetric T** Price: 10.5	
Resources	Buyer	Seller	Buyer	Seller	Buyer	Seller
0	20	20	20	20	20	20
1	37	22	22	37	29.5	29.5
2	54	24	24	54	39	39
3	71	26	26	71	48.5	48.5
4	88	28	28	88	58	58
5	105	30	30	105	67.5	67.5
6	122	32	32	122	77	77
7	139	34	34	139	86.5	86.5
8	156	36	36	156	96	96
9	173	38	38	173	105.5	105.5
10	190	40	40	190	115	115

Experimental parameters such as prices, endowments, and productivity were determined using the inequity aversion model of Fehr and Schmidt (1999), along with their suggested calibration of parameters. Their suggested distribution of (α, β) in the population implies average values $(\alpha, \beta) = (0.85, 0.32)$. With these average parameters and the chosen parameters in the experimental design, the Fehr and Schmidt model predicts that the favored party will always prefer full trade, while the disfavored party will always prefer to trade zero resources. This result follows from the linear inequity preferences assumed by Fehr and Schmidt, which implies "threshold" inequality effects. When the distribution of the gains from trade passes some threshold (which varies among people in a population), the individual suddenly shifts from full to no trade. We find this literal interpretation implausible and hypothesize that participants may also want to trade at intermediate levels if given the chance. Competing models such as Charness and Rabin (2002) that focus on maximin preferences, reciprocity, or efficiency would all predict full trade in our game, in line with the predictions of the standard self-interest model.

Finally, it should be noted that the splitting of the experiment into a number of rounds is in some sense arbitrary. Each player receives the same amount of resources in each round, and the trading terms are the same for each pair of players over all rounds. Consequently, the ten-stage game is theoretically equivalent to a one-stage game where the players receive all their resources and trade only once. Our choice of ten rounds is in therefore arbitrary and was chosen to give people more time to learn the mechanics of the game, as well as the opportunity to observe increasing differences in the distribution of cumulative payoffs.

At the end of the experiment, all participants answered a short survey, including general comments on the experiment. The survey answers reveal that nearly all participants understood the experiment. Some participants state that they used the first few rounds to test the outcomes of different trade amounts, before deciding on a strategy in the remaining rounds. As a result of this, we look at behavior in all ten rounds and over the last seven, six, and five rounds respectively when analyzing the results. Only one participant revealed that she felt uncertain about how the experiment worked, even after preliminary testing. Hence, the trading pair involved was removed from the data.

3 Results

This section analyses trading behavior of 126 students from Columbia University who participated in six experimental sessions run at similar times

during a two-week period in April 2014.[9] The experiment is composed of three treatments: 22 couples took part in the "buyer treatment", while 20 couples in the "seller treatment" and 21 couples in the "symmetric treatment". We present our results in terms of resources traded per couple or in terms of individuals' offered-for-trade resources. These observations are averaged over all ten rounds played or over the last seven, six, and five rounds respectively in case learning effects influenced the results. We also look at whether individuals restrict trade analyzing a binary variable that takes values one when the suggested trade is less than ten resources and zero otherwise. Our results are estimated with a between-subjects design and non-parametric tests are performed over completely independent observations (trading pairs) and over individuals. Note that all figures reported below show outcomes per round averaged over couples and therefore are not plotting independent observations. The propose of the figures is to illustrate the average trends over time.

Result 1: Buyer and seller treatments are not significantly different from each other.

We find no framing effect of calling one party the "seller" and the other the "buyer". By construction the Seller and buyer treatments are mirror images of each other (see Table 1). Consequently, we would not expect any difference in the buyer-favoring and seller-favoring price treatments in so far as trade restrictions were caused by the payoff-differences. However, some differences in behavior could occur if attitudes towards inequality-increasing trade are influenced by the framing of one party as seller and the other as buyer. Figure 1 visually suggests no statistical difference in the number of resources traded between the buyer and seller treatments. Moreover, a Wilcoxon rank-sum non-parametric test run over the average trade suggests that the null hypothesis cannot be rejected (p-value 0.7505) when considering completely independent observations (within-pair trade averaged over the last last ten, seven, six, or five rounds respectively). In other words, whether the seller or the buyer was the favored or the disfavored party did not affect traded volumes.

Given Result 1, we are able to merge the data from buyer and seller treatments in one unique asymmetric treatment and we discuss the next results comparing the asymmetric treatment with the symmetric treatment. Similar results are reported separately for buyer and seller treatments in the Robustness checks in Appendix A.

[9]The original sample was 128, but one couple is discarded because one subject reported no understanding of the game even after some rounds.

Figure 1: The figure shows the mean actual trade per round in the two asymmetric price treatments – buyer treatment and seller treatment – with standard errors.

Result 2: Full trade in the symmetric treatment.

In the symmetric treatment, the price was set such that gains from trade were split equally between buyer and seller within each trading pair. When payoffs were symmetric and inequality between parties is absent, subjects quickly maximized their payoffs by trading the maximum number of resources possible (i.e. ten resources). Figure 2 shows how, in the symmetric treatment, an upwards trend reaches full trade after a few rounds. The survey run at the end of the experiment suggests that some students used few rounds at the beginning of the experiment to understand the game and to test different trading volumes. As a consequence, we run non-parametric tests considering all ten rounds as well as discarding the first four, five, and six rounds to ensure robustness of the results. The average number of resources traded in the symmetric treatment was 9.65 when considering all ten rounds and 9.93, 9.96 and 9.98 when considering the last seven, six, and five rounds respectively (Table 2). Only one third of the 21 couples in the symmetric treatment ever deviated from full trade. Of these seven buyer-seller pairs, only one failed to reach full trade within the fourth round (Figure 6 in Appendix B). In other words, only 1 out of 21 participant pairs traded less than the full amount after the forth round, and even this pair chose full trade in the last three rounds.

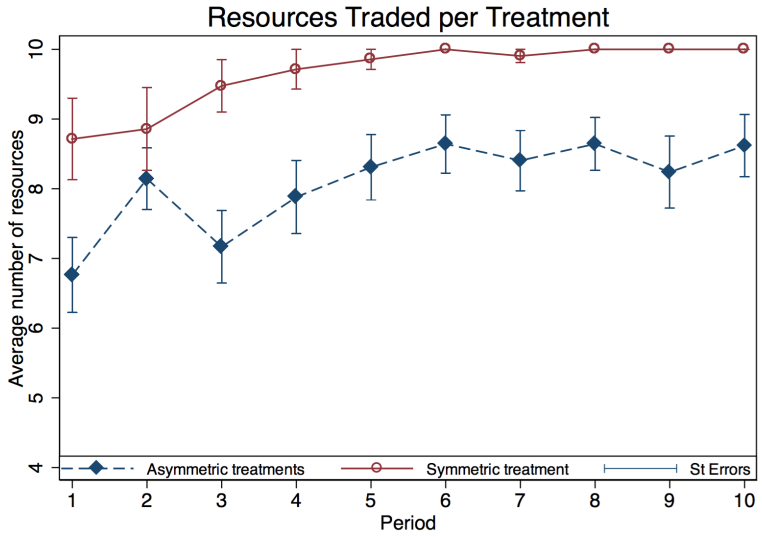

Figure 2: Mean actual trade per round in the Symmetric and Asymmetric price treatments, with standard errors.

Table 2: Average Actual Trade in Asymmetric and Symmetric Sreatments

Obs	**Asymmetric T** (N=42)	**Symmetric T** (N=21)	**Wilcoxon p-val**
1st-10th round	8.08	9.65	**(0.0002)**
4th-10th round	8.39	9.93	**(0.0002)**
5th-10th round	8.47	9.96	**(0.0006)**
6th-10th round	8.51	9.98	**(0.0015)**

Non-parametric two-sample Wilcoxon rank-sum (Mann-Whitney) tests are used to test the null hypothesis of equality between the distributions of the asymmetric treatment and the symmetric treatment. All reported p-values imply the rejection of the null hypothesis.

Result 3: Trade is significantly restricted in the asymmetric treatment compared to the symmetric treatment.

There is significantly less trade in the asymmetric treatment compared to the symmetric treatment (Figure 2). Specifically, when considering all rounds, the mean trading volume is 8.08 in the asymmetric treatment and 9.65 in the symmetric treatment (Table 2). These trade restrictions are the result of 32 couples (76%) in the asymmetric treatment restricting trade at least once in the ten rounds played, and only seven couples (33%) are trading less than ten resources in the symmetric treatment. If we consider only the last six rounds of the game to eliminate possible learning effects, the number of couples trading less than ten resources in the asymmetric treatment decreases to 18 (43%) bringing the average trade in the asymmetric treatment down

to 8.51. In the symmetric treatment the average trade is 9.98 because only one couple (4.7%) restricted trade in the last six rounds of the experiment (Table 2).

Non-parametric two-sample Wilcoxon rank-sum tests reported in Table 2 suggests that the null hypothesis, equality of the means between the asymmetric treatment and the symmetric treatment, can be rejected with a high statistical significance independent of the number of rounds considered (p-val between .0015 and .0002). Moreover, the reduction observed in actual trading volumes in the asymmetric treatment creates an efficiency loss of about 15% (16.3% considering all rounds and 14.7% considering the last six rounds) compared to trade occurring in the symmetric treatment.

Result 4: The worse-off party within a pair is the main driver of the observed trade restrictions in the asymmetric treatment.

The disfavored party within a trading pair is found to be the main constraint on trading volumes.[10] The average number of resources suggested for trade by the worse-off party is between 12% and 9% lower than the average trading proposals of the better-off party in all rounds (Figure 3). Table 3 reports the individual suggested mean trade when all ten rounds are considered and when only the last seven, six, and five rounds respectively are considered, together with the p-values from the non-parametric test in parenthesis. Average offers of disfavored players in the asymmetric treatment are significantly lower than average offers of players in the symmetric treatment, no matter the number of rounds considered (two-sample Wilcoxon non-parametric tests, p-values between 0.0038 and 0.0004). On average, Favored players suggest lower or equal trade in every round, but these differences are not statistically significant (Table 3). In the non-parametric tests reported in Table 3, the unit of observation for the asymmetric treatment is the average suggested trade for each type of player (favored or disfavored) and this is compared to average suggestion of each couple in the symmetric treatment.

[10] We define buyers and sellers in the symmetric treatment as "symmetric" players.

Table 3: Average trade proposals by type of player

Obs	**Favored Pl.** (N=42)	**Disfavored Pl.** (N=42)	**Symmetric Pl.** (N=21)
1st-10th round	9.46	8.33	9.82
Wilcoxon p value	(0.5372)	**(0.0009)**	
4th-10th round	9.58	8.57	9.96
Wilcoxon p value	(0.3304)	**(0.0038)**	
5th-10th round	9.60	8.65	9.98
Wilcoxon p value	(0.3304)	**(0.001)**	
6th-10th round	9.60	8.72	9.99
Wilcoxon p value	(0.1816)	**(0.0004)**	

The non-parametric Wilcoxon rank-sum tests the null hypothesis of equality of means respectively between the trade proposals of the favored and disfavored players in the asymmetric treatment, and the symmetric players in the symmetric treatment. P-values in bold imply the rejection of the null hypothesis.

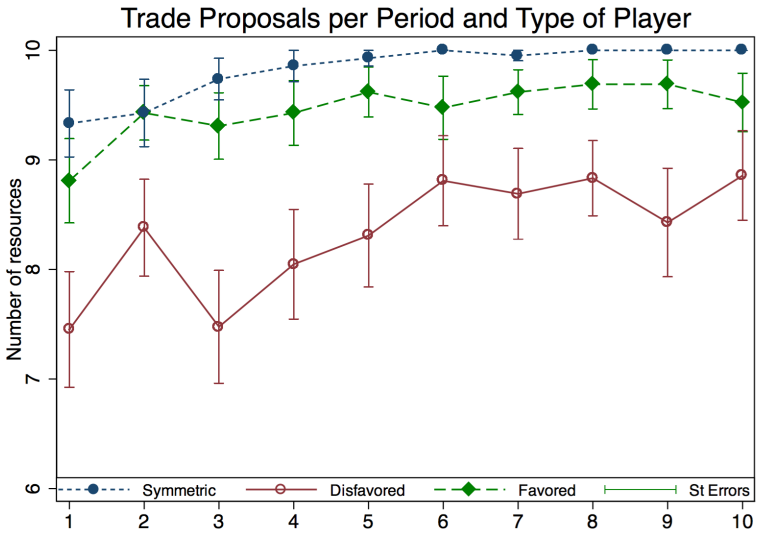

Figure 3: Mean trade proposals per round and type of player in the Symmetric and Asymmetric price treatments with standard errors.

Result 5: A small share of favored subjects also restrict trade.

If favored players are unaffected by the price treatments, their probability of proposing less than full trade should be the same as in the symmetric treatment. There are some indications that the favored party in the asymmetric treatments sometimes reduce trade (Figure 3). The experimental data show that a small number of favored players also suggest to trade less than ten resources and this number is higher than in the symmetric treatment. To test the independence of the observed distribution we use the Chi squared test

with one degree of freedom. The results are weakly significant (10% level) and are sensible to the number of rounds considered (Table 4). Given the small number of observations we also run Fisher's exact test, which delivered similar results.

Table 4: Trade restrictions by type of player and treatment

	Favored Pl. (N=42)	**Symmetric Pl** (N=42)	**Pearson Chi2** DF=1
1st-10th round	16 (38%)	9 (21%)	2.79 (0.095)
4th-10th round	7 (17%)	1 (2%)	4.97 **(0.026)**
5th-10th round	5 (12%)	1 (2%)	2.87 (0.090)
6th-10th round	5 (12%)	1 (2%)	2.87 (0.090)

Number of individuals who suggested trading less than ten resources at least once. We consider all ten rounds and the last seven, six, and five rounds. The last column reports Person Chi square with one degrees of freedom with p-values in parenthesis.

It is worth noticing that among the 16 better-off players who suggested at least once to trade less than ten resources in all ten rounds, twelve actually restricted trade, i.e. they were the party who suggested the smaller number of resources in at least one round. If we look at the last five rounds only, four favored players, out of five who suggested less than ten resources, actually constrained trade. In the symmetric treatment, eight out of nine restricted trade in the first few rounds and only one player actively restricted trade after the fourth round.

4 Conclusions

Fairness preferences have previously been shown to affect outcomes in a variety of economic games, such as the ultimatum game, the gift exchange game and the prisoner's dilemma (see Fehr and Schmidt (2006) for an overview). However, introducing a trade or market frame is enough to drastically reduce the expression of other-regarding preferences ((Cherry et al. (2013) and Hoffman et al. (1994)). In the previous market experiments where preferences for equality do seem to play a role, the design cannot distinguish fairness preferences from strategic or retaliation motives (Franciosi et al. (1995), Ruffle (2000) and Forsythe et al. (1994)). The experimental evidence presented in this article is the first to show that fairness preferences can be triggered in a trade setting with buyer-seller frame in the absence of strategic or retaliation motives.

When the gains from trade are split equally between buyer and seller, trade quickly converges to its maximum after a few rounds of practice. In contrast,

when the terms of trade generate substantial inequality, trade is significantly restricted. Specifically, about 60% percent of pairs in treatments with uneven distribution traded less than ten resources in our experiment, reducing traded quantities on average by 15% percent, compared to the symmetric treatment group. These trade restrictions are found to be mostly driven by the disfavored players, who suggests trading about 10% less resources compared to favored players. Specifically, 28 out of 42 disfavored players, restricted trade at least once in the asymmetric treatments. Results are robust even when we discard the first three, four, and five rounds from the analysis to avoid learning effects. These results are in line with predictions from the preference structure and average parameter values proposed for Inequity Aversion in Fehr and Schmidt (1999). Other theories of social preferences such as maximin preferences and a preference for efficiency predict full trade in our setting and hence cannot explain the observed trade restrictions.

One issue regarding the generalization of our results is the high salience, by design, of the outcome inequality in our experiment. In a real world setting, trading partners will rarely know the exact gains from trade or how they are distributed. Fair trade products, however, can be viewed as a way of improving consumer information about how the gain is distributed, that is used in order to increase purchases from consumers who care about outcome inequality. It has been argued that by removing this asymmetry of information it is possible to increase consumers' welfare by facilitating products differentiation (Hainmueller et al. (2014)). Despite this limitation we decided to have a high saliency design to ensure the clarity of the game and to simplify payoffs calculations for the participants.

Another limitation to the external validity of this study is the fixed price feature in our design, which determines the distribution of gains-of-trade between the two players. One could argue that there are no free markets which are determined by fixed prices, and that the results therefore are irrelevant for real markets. If there is a scope for Pareto-improving trade between two parties, they should always be able to negotiate a price which is acceptable for both of them. However, this may not be how markets are perceived by people at both ends of the value-chain. Workers in producing factories are seldom able to negotiate their wage. Similarly, consumers in the grocery stores are not able to pay a higher price in order to increase the salary of the workers producing the good. For these agents, market terms may appear fixed to the individuals, and their only choice is whether to take part in the market or not. Our experimental design aim to reproduce such a situation where the gains from trade are determined by the experimenters and cannot be affected in any way. Because of this feature we are able to present a simple and clean design which allow us to isolate fairness concerns from strategic or retaliation motives. The results show that in such situations, inequality aversion may severely restrict trade and reduce efficiency.

A Robustness Checks

In both buyer and seller treatment, we observe a significant reduction in trading volumes compared to the symmetric treatment (Figure 4). Specifically, in the buyer treatment the average number of resources traded over all rounds is 8.20, while in the seller treatment it is 7.96 (Table 5). We then test for the equality of each of these means with the mean trade volume in the symmetric treatment (9.65) using the non-parametric two-sample Wilcoxon rank-sum test. The non-parametric test is performed for the within-participant-pair trading volume averaged over multiple rounds, providing one independent observation for each participant pair. The within-couple trading volumes are averaged over the last five, six, and seven rounds, as well as for all ten rounds, and the resulting P-values are between 0.0001 and 0.005 (reported in Table 5). Varying how many of the early rounds we include helps us assess whether the groups mainly differed in the early rounds where learning effects might still dominate. These results imply a very high significant difference in the number of resources traded between each asymmetric treatment and the symmetric treatment.

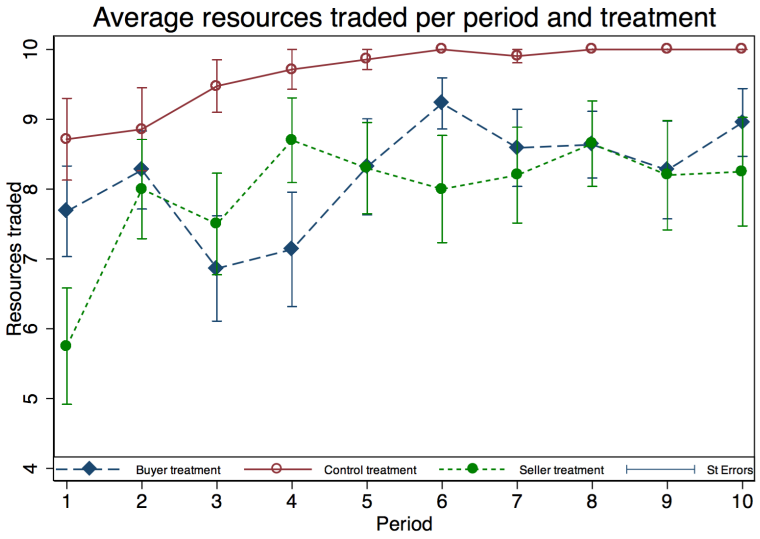

Figure 4: Average trade per round in the three price treatments with standard errors

Table 5: **Realized trade per couple and non-parametric test**

Treatment	**1-10th round**	**4-10th**	**5-10th**	**6-10th**
Symmetric T (N=21)				
mean	9.65	9.93	9.96	9.98
sd	(0.75)	(0.34)	(0.18)	(0.09)
Buyer T (N=22)				
mean	8.20	8.45	8.67	8.74
sd	(2.00)	(2.07)	(1.97)	(1.91)
Wilcoxon (p-val)	0.0011	0.0001	0.0007	0.0017
Seller T (N=20)				
mean	7.96	8.33	8.27	8.26
sd	(2.41)	(2.55)	(2.70)	(2.79)
Wilcoxon (p-val)	0.0013	0.0026	0.0023	0.005

Non-parametric two-sample Wilcoxon rank-sum (Mann-Whitney) tests are used to test the null hypothesis of equality between the distributions of the asymmetric treatments and the symmetric treatment. All reported p-values imply the rejection of the null hypothesis.

The disfavored party is found to be the main constraint on trading volumes. In fact, the average trade proposed by the disfavored party is lower or equal to the average trading proposals of the favored party in all rounds for both symmetric treatments (Figure 5). Table 6 reports the individual suggested trade averaged over the last five rounds as well as over all ten rounds by type of player in each treatment. A two-sample Wilcoxon non-parametric test finds the differences between the average offers of disfavored players and their peers in the symmetric treatment is statistically significant. Specifically, in the buyer treatment, sellers send significantly lower offers than sellers in the Symmetric price treatment (p-value 0.0035 in the data from the last five rounds). The reverse holds in the seller treatment (p-value of 0.0039 when considering only the last five rounds).

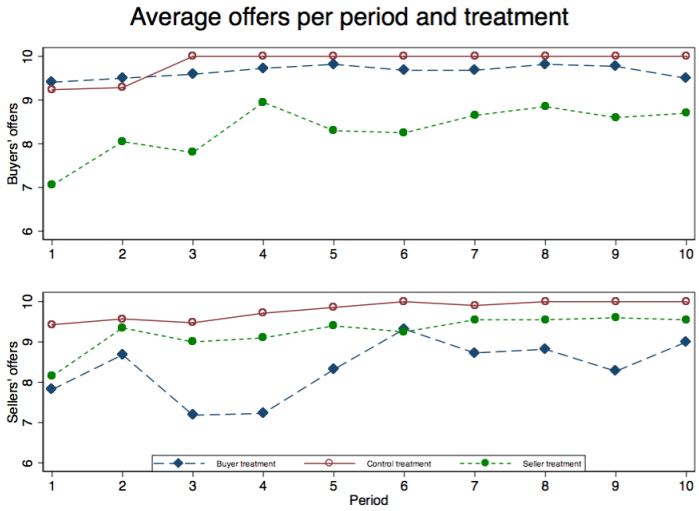

Figure 5: Average offers per round in the three price treatments with standard errors

Table 6: **Mean suggested trade and p-values Wilcoxon test**

	Buyer T (p-val)	**Seller T (p-val)**	**Symmetric T**
1-10th round			
Buyer	9.65 (0.5399)	**8.32 (0.006)**	9.85
Seller	**8.34 (0.0002)**	9.25 (0.076)	9.80
6-10th round			
Buyer	9.69 (0.0832)	**8.61 (0.0035)**	10.00
Seller	**8.83 (0.0039)**	9.50 (0.488)	9.98

The reported p-value from the Wilcoxon rank-sum tests are calculated over individuals' offers averaged over ten rounds or over the last five rounds. The disfavored party (the seller in the buyer treatment and the buyer in the seller treatment) significantly restricts trade compared to buyers and sellers in the symmetric treatment.

B Additional Figures and Tables

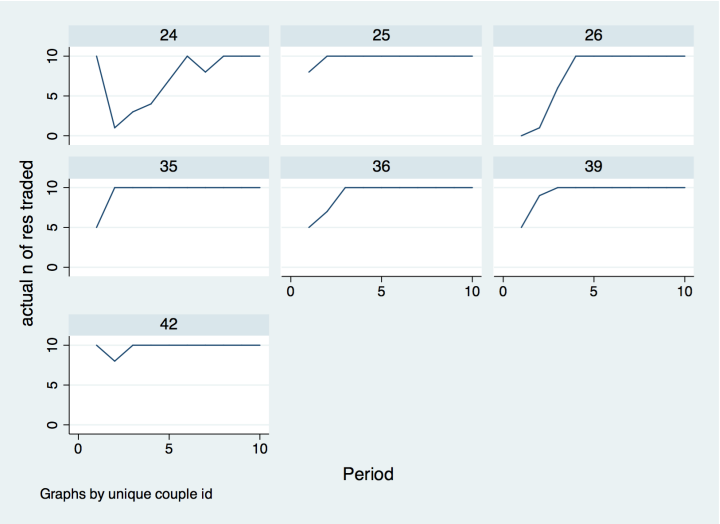

Figure 6: Resources traded by the couples who traded less than ten resources in the control treatment

Instructions

General information

Welcome to this economics experiment. You will now make anonymous decisions on a computer – neither the researchers nor the other participants will ever know which decisions were made by which participant. Communication with other participants during the experiment is not allowed. If you have questions or need help with the computer, please raise your hand and one of us will approach you and assist you privately.

It is important to us that you trust the information we give in this experiment, so we want to emphasize that *all the information you receive is true*. It is also important that you fully understand how the experiment works, so we ask you to read the instructions carefully.

Instructions

During the experiment you will be asked to make decisions regarding trade, in terms of ECU points (Experimental Currency Unit) that will be exchanged into US dollars and paid out at the end of the experiment.

Exchange rate: 50 ECU= $1

Before the experiment starts you will be randomly placed into pairs. Within each pair, one individual will be randomly assigned to be a *seller* while the other one will be a *buyer*. Nobody can learn the identity of their co-players, and no communication with the co-player is possible. The pairs remain fixed throughout the experiment.

The experiment consists of 10 identical periods. At the beginning of each period you will receive a fixed sum of resources. Buyers and sellers differ in the number of resources they receive and in their productivity, which determines how much money they can earn from production using these resources.

During the experiment, resources can be traded with your co-player to increase earnings (details regarding the trade follow). All trade will occur at set price determined at the beginning of the experiment. This price will be the same in all rounds and it will be presented on the first screen.

At the end of each period you will receive a certain number of ECU based on the number of resources you are currently holding, your productivity and the payments or earnings due to trade. Buyers earn 20 ECU for each resource unit held after trading while sellers earn 1 ECU for each resource held after trading. In other words, a buyer produces 20 times as many ECU per resource unit then a seller.

The following table summarizes the relevant information.

	Buyer	**Seller**
Number of resources received every period	1	20
Productivity (Number of ECU earned for each resource held at the end of a period)	20	1

Briefly told, you will be given a certain amount of resources at the beginning of each period. You and your co-player will then choose how many resources to trade. When trading is over, each player's earnings for that period will be determined by the amount of resources they hold, their productivity, and the payments made or earned from the traded resources. For instance, imagine that a buyer can buy resources at a price of 10 ECU. Each resource would then produce 20 ECU to the buyer while costing 10 ECU. *The net payoff for the buyer would then be 10 ECU per unit bought.* In the same trade, the co-player, the seller, would earn 10 ECU per resource sold, while losing 1 ECU in foregone production. *The net gain for the seller in this period would be thus 9 ECU per resource.*

Trade is organized as follows:

Both players in a pair state how many resources they wish to trade between 0 and 10. Since trade is voluntary, the lowest of the two numbers will determine the actual amount traded. For instance: if the buyer wants to buy 5 units while the seller wants to sell 6 units, then the seller will only be able to sell the 5 units that the co-player wants to buy.

To help you decide, you will be able to visualize the outcome resulting from possible trades by clicking on the "visualize" button. This will display two bars representing your total payoff in the case the proposed trade is going to be the actual trade. The bars highlight your accumulated payoff from previous periods, the payoff resulting from the number of resources hold in a period and the payoff resulting from trade in that period. An example is illustrated below.

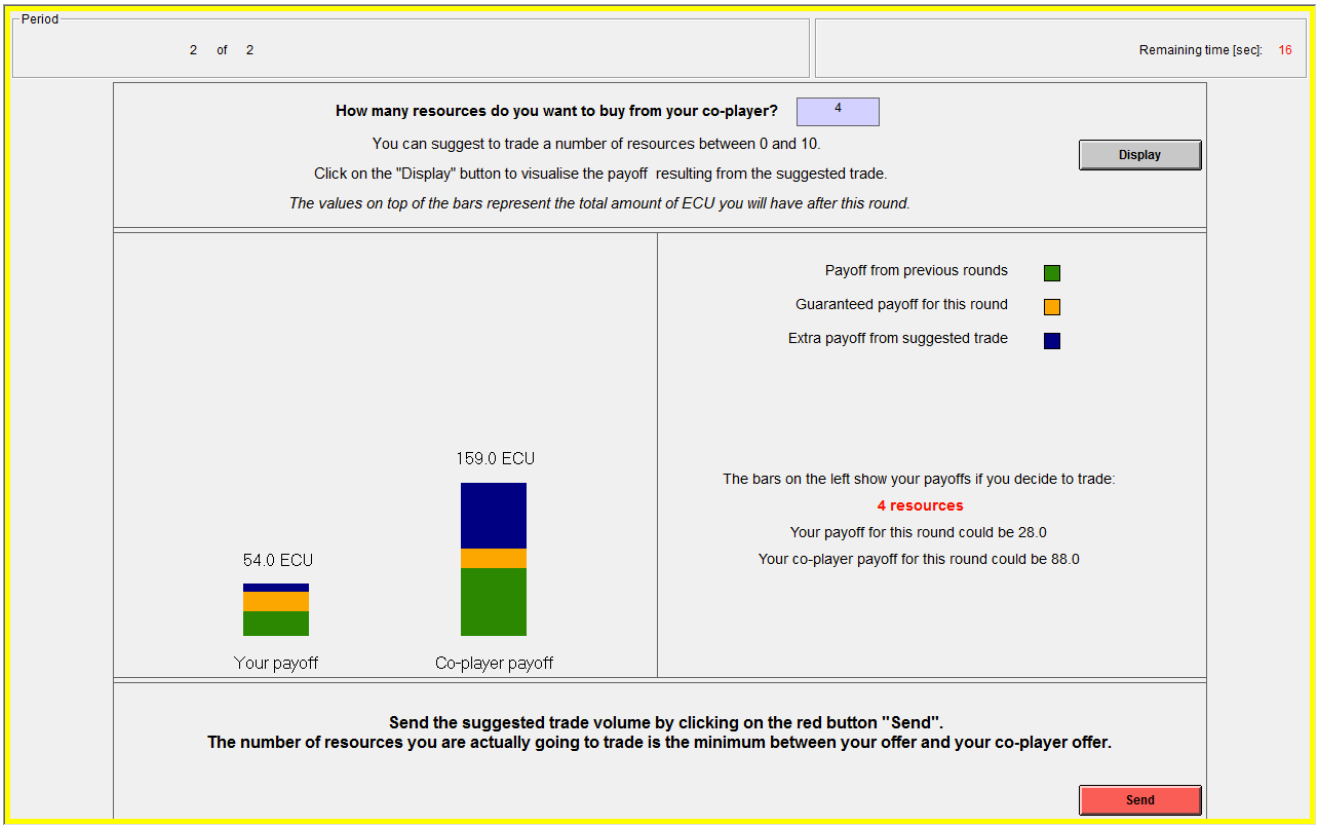

You have to visualize the outcome of at least one proposed trading volume before you can send your actual decision. Your co-player will not know which trading level (or how many) you examine before making your actual choice. When you have decided how much you wish to trade, click the "send" button. In a screen like the one below, you will then learn how many resources were traded and how the trade has affected the earnings you and your co-player have made so far in the experiment.

It is important to note that the *axes on the columns changes from round to round,* since they also display how much you have already earned. To improve clarity, above each column you will also see the **total** number of ECU displayed by the column.

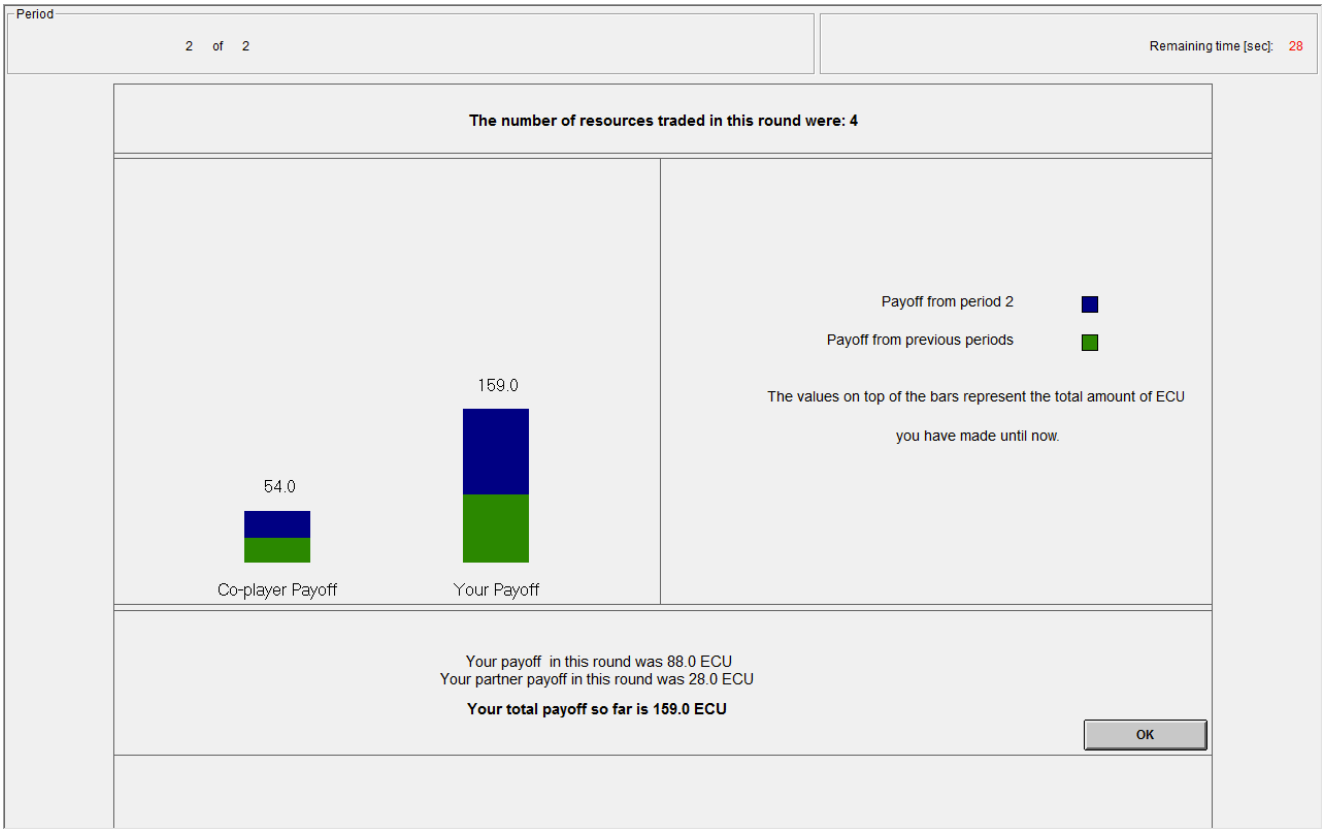

Periods

This is how one period of the experiment is run. There will be 10 such periods.

When the last round is over, we will ask you to answer a survey. Note that this survey, like the experiment, is anonymous.

Questions?

If you have any questions please raise your hand and we will assist you privately.

References

Bewley, Truman F (1999), *Why wages don't fall during a recession*. Harvard University Press.

Bolton, Gary E and Axel Ockenfels (2000), "Erc: A theory of equity, reciprocity, and competition." *American economic review*, 166–193.

Bolton, Gary E and Axel Ockenfels (2006), "Inequality aversion, efficiency, and maximin preferences in simple distribution experiments: comment." *The American economic review*, 1906–1911.

Camerer, Colin (2003), *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press.

Castaldo, Sandro, Francesco Perrini, Nicola Misani, and Antonio Tencati (2009), "The missing link between corporate social responsibility and consumer trust: The case of fair trade products." *Journal of business ethics*, 84, 1–15.

Charness, Gary and Matthew Rabin (2002), "Understanding social preferences with simple tests." *Quarterly journal of Economics*, 817–869.

Cherry, Todd L, Steffen Kallbekken, Stephan Kroll, and David M McEvoy (2013), "Cooperation in and out of markets: An experimental comparison of public good games and markets with externalities." *Economics Letters*, 120, 93–96.

Davis, Douglas D (1993), *Experimental economics*. Princeton University Press.

Elfenbein, Daniel W and Brian McManus (2010), "A greater price for a greater good? evidence that consumers pay more for charity-linked products." *American Economic Journal: Economic Policy*, 2, 28–60.

Elliott, Kimberly Ann and Richard Barry Freeman (2003), "Can labor standards improve under globalization?" *Washington, DC: Institute for International Economics.*

Engelmann, Dirk (2012), "How not to extend models of inequality aversion." *Journal of Economic Behavior & Organization*, 81, 599–605.

Engelmann, Dirk and Martin Strobel (2004), "Inequality aversion, efficiency, and maximin preferences in simple distribution experiments." *American economic review*, 857–869.

Engelmann, Dirk and Martin Strobel (2006), "Inequality aversion, efficiency, and maximin preferences in simple distribution experiments: Reply." *The American economic review*, 1918–1923.

Fehr, Ernst and Simon Gächter (2000), "Fairness and retaliation: The economics of reciprocity." *The journal of economic perspectives*, 159–181.

Fehr, Ernst, Georg Kirchsteiger, and Arno Riedl (1993), "Does fairness prevent market clearing? an experimental investigation." *The Quarterly Journal of Economics*, 108, 437–459.

Fehr, Ernst, Michael Naef, and Klaus M Schmidt (2006), "Inequality aversion, efficiency, and maximin preferences in simple distribution experiments: Comment." *The American economic review*, 1912–1917.

Fehr, Ernst and Klaus M Schmidt (1999), "A theory of fairness, competition, and cooperation." *The quarterly journal of economics*, 114, 817–868.

Fehr, Ernst and Klaus M Schmidt (2006), "The economics of fairness, reciprocity and altruism–experimental evidence and new theories." *Handbook of the economics of giving, altruism and reciprocity*, 1, 615–691.

Fehr, Ernst and Klaus M Schmidt (2010), "On inequity aversion: A reply to binmore and shaked." *Journal of Economic Behavior & Organization*, 73, 101–108.

Fischbacher, Urs (2007), "z-tree: Zurich toolbox for ready-made economic experiments." *Experimental economics*, 10, 171–178.

Forsythe, Robert, Joel L Horowitz, Nathan E Savin, and Martin Sefton (1994), "Fairness in simple bargaining experiments." *Games and Economic behavior*, 6, 347–369.

Franciosi, Robert, Praveen Kujal, Roland Michelitsch, Vernon Smith, and Gang Deng (1995), "Fairness: effect on temporary and equilibrium prices in posted-offer markets." *The Economic Journal*, 938–950.

Großer, Jens and Ernesto Reuben (2013), "Redistribution and market efficiency: An experimental study." *Journal of Public Economics*, 101, 39–52.

Hainmueller, Jens, Michael J Hiscox, and Sandra Sequeira (2014), "Consumer demand for the fair trade label: Evidence from a multi-store field experiment." *Review of Economics and Statistics*.

Hoffman, Elizabeth, Kevin McCabe, Keith Shachat, and Vernon Smith (1994), "Preferences, property rights, and anonymity in bargaining games." *Games and Economic Behavior*, 7, 346–380.

Kahneman, Daniel, Jack L Knetsch, and Richard Thaler (1986), "Fairness as a constraint on profit seeking: Entitlements in the market." *The American economic review*, 728–741.

MMI (2010), "Den store norske klima- og milj£unders£kelsen 2010. the nor-

wegian climate and environmental survey." *Ipsos MMI and CICERO Center for climate research.*

Prasnikar, Vesna and Alvin E Roth (1992), "Considerations of fairness and strategy: Experimental data from sequential games." *The Quarterly Journal of Economics*, 107, 865–888.

Rabin, Matthew (1993), "Incorporating fairness into game theory and economics." *The American economic review*, 1281–1302.

Roth, Alvin E, Vesna Prasnikar, Masahiro Okuno-Fujiwara, and Shmuel Zamir (1991), "Bargaining and market behavior in jerusalem, ljubljana, pittsburgh, and tokyo: An experimental study." *The American Economic Review*, 1068–1095.

Ruffle, Bradley J (2000), "Some factors affecting demand withholding in posted-offer markets." *Economic Theory*, 16, 529–544.

Rutström, E Elisabet and Melonie B Williams (2000), "Entitlements and fairness:: an experimental study of distributive preferences." *Journal of Economic Behavior & Organization*, 43, 75–89.

Smith, Vernon L (1962), "An experimental study of competitive market behavior." *The Journal of Political Economy*, 111–137.

Smith, Vernon L (1991), *Papers in experimental economics*. Cambridge University Press.

Tyran, Jean-Robert and Dirk Engelmann (2005), "To buy or not to buy? an experimental study of consumer boycotts in retail markets." *Economica*, 72, 1–16.