Preprints of the
Max Planck Institute for
Research on Collective Goods
Bonn 2016/20

Selling money on eBay: A field study of surplus division

Alia Gizatulina
Olga Gorelkina

MAX PLANCK SOCIETY

Electronic copy available at: https://ssrn.com/abstract=2884191

Preprints of the
Max Planck Institute
for Research on Collective Goods Bonn 2016/20

Selling money on eBay: A field study of surplus division

Alia Gizatulina / Olga Gorelkina

December 2016

Max Planck Institute for Research on Collective Goods, Kurt-Schumacher-Str. 10, D-53113 Bonn
http://www.coll.mpg.de

Electronic copy available at: https://ssrn.com/abstract=2884191

SELLING MONEY ON EBAY: A FIELD STUDY OF SURPLUS DIVISION *

Alia Gizatulina† and Olga Gorelkina‡

December 10, 2016

Abstract

We study the division of trade surplus in a competitive market environment by conducting a natural field experiment on German eBay. Acting as a seller, we offer Amazon gift cards with face values of up to 500 Euro. Randomly arriving buyers, the subjects of our experiment, make price offers according to eBay rules. Using a novel decomposition method, we infer offered shares of trade surplus and find that the average share proposed to the seller amounts to 29%. Additionally, we document: (i) insignificant effects of stake size; (ii) poor use of strategically relevant public information; and (iii) behavioural differences between East and West German subjects.

Keywords: *Field experiment, ultimatum game, surplus division, bargaining, Internet trade, eBay.*

JEL codes: *C72, C93, C57.*

*We are grateful to Anna Petukhova for the excellent research assistance. This paper has benefited enormously from our discussions with Brian Cooper, Christoph Engel, Thomas Epper, Alexandra Fedorets, Susann Fiedler, Dominik Grafenhofer, Ioanna Grypari, Werner Güth, Martin Hellwig, Oliver Himmler, Roland Hodler, Johannes Hörner, Guillaume Hollard, Laurent Lamy, Rida Laraki, John List, Ulrike Malmendier, Daniel Martin, Wanda Mimra, Humberto Moreira, Maryam Naghsh-Nejad, Aniko Oery, Nikita Roketskiy, Alvin Roth, Larry Samuelson, Franziska Tausch, and Shmuel Zamir. We gratefully acknowledge the financial support of the Max Planck Society and the Cowles Foundation for Research in Economics at Yale University. We used `Scilab 5.4.1` for numerical estimations and `Mathematica 10.0` for parametric testing. The program codes and data are available from the authors upon request.

†University of St. Gallen and Max Planck Institute for Research on Collective Goods; `alia.gizatulina@unisg.ch`

‡Max Planck Institute for Research on Collective Goods, Bonn; `gorelkina@coll.mpg.de`

Electronic copy available at: https://ssrn.com/abstract=2884191

1 Introduction

If trade generates a surplus, the price determines how the surplus is split. In thin markets without a common price benchmark, trading partners must agree on a mutually acceptable price. We run a *natural field experiment* to quantify the subjects' intentions to share the surplus from market trade. Acting as sellers of a good with a known value we collect price offers from randomly arriving buyers. We estimate the trade surplus from the data and map the received price offers into proposed shares of surplus. Our paper, hence, gives a unique account of sharing norms among anonymous traders in the field and contributes to the existing experimental studies on surplus division in a market setting.

The object we offer for sale is an Amazon gift card. Amazon gift cards are digital codes that any Amazon customer can use to add money to his Amazon account. Since the card's amount can be split, stored, and combined with other payment methods, Amazon gift cards are essentially cash to Amazon buyers and hence they provide an unprecedented opportunity to control the subjects' valuations in a natural field setting. On the seller side, we observe that the cards are typically received as presents, as well as rewards for participation in Internet surveys or lottery prizes that firms use to attract customers. If a card owner does not use Amazon and prefers cash to a credit on his Amazon account, then gains from trade can be generated by selling the card to an active Amazon customer. The existence of gains from trade of Amazon gift cards gave rise to a substantial secondary market on eBay, and in particular its German site (www.ebay.de), where we conduct the experiment.

In the experiment, we offer Amazon gift cards whose face value ranges between 5 and 500 Euros, with seven value treatments in total. As the trading protocol, we employ eBay's *Buy it now or best offer* format (BINBO). According to the rules of BINBO, the seller posts an initial ask price and invites any eBay buyer to acquire the object at that price or make their own price offer. The card is sold if some buyer accepts the posted price or if a buyer makes an alternative price offer that the seller accepts.

The BINBO format features two types of information frictions. First, at the moment of making his offer, the buyer observes only the *number* of offers made to the seller, but not the offered *prices*. This information is available only to the seller. Second, due to the market's dynamic structure, neither trading party observes the effective number of competing buyers bound to arrive by the end of the listing. In that sense, the environment is uncertain, and the buyers are unaware of the seller's true outside option.[1] Such uncertainty, while absent in any of the related lab experiments, is rather common to many real-world markets, be it an electronic trading platform or a Middle East bazaar.

Beyond being a realistic feature of many markets, the uncertainty of the buyer about the seller's outside option has an important implication for buyers' behaviour and the analysis in general. Namely, from the perspective of each individual buyer, there is a non-trivial surplus to be shared. This differs from the setting of the typical lab experiment on surplus sharing, where there is a fixed and commonly known degree of competition, with at least two buyers. To see this, note that the surplus generated in a trade between a given buyer and the seller is equal to the difference between the seller's outside option and the nominal value of the card. In a setting where the number of competitors is greater than one, certain and commonly known, the equilibrium surplus from trade between any individual buyer and the seller is zero.[2] In contrast, in a bilateral setting with only one buyer, the seller's outside option is his own usage value and so the trade surplus may be substantial. When the number of buyers is uncertain, the total surplus *perceived* by a given buyer depends on his beliefs about the seller's next-best option. If the seller's own valuation is zero, his next-best option is to

[1]Note that the *buyer's* outside option is irrelevant for the value of transaction with a given seller, as long as the buyer's preference for money is insatiable. Whether the buyer can purchase a gift card elsewhere, does not affect his valuation of the opportunity to get another card with a discount, i.e. to get money for free. See Section 5.1.

[2]With the commonly known degree of competition, the unique dominance solvable outcome is for both buyers to offer the entire nominal value to the seller. Hence from the perspective of buyer i, the surplus of his trade with the seller equals to zero, as the seller's outside option is equal to the nominal value of the card – the price that another buyer would offer.

trade with another buyer, who may or may not arrive. Hence, from the perspective of the buyer arriving first, his best strategy is to offer a price that is lower than the card's nominal value. Therefore, any subsequent buyer, when he offers his own price, would anticipate that the seller's outside option is not equal to the cash equivalent of the entire nominal value of the card. Put differently, here, similarly to a first-price sealed-bid auction, each buyer, when he makes an offer to the seller, weights off increasing the probability of being the winner against reducing the total price paid. This also implies that uncertainty makes the surplus that a given buyer and the seller avail themselves of to be non-degenerate.

As preliminary groundwork for the analysis of surplus division, we study whether the experimental data display any regularities by linking price offers to other subjects' observables. First, we find that the observed distributions of relative offers[3] do not vary with the amount of money at stake. The finding that stake size does not affect the relative offers permits us to pool the data across the different value treatments and to work subsequently only with the single dataset. Second, we find that the observed buyer behaviour is insensitive to the public information about the degree of competition or the public history of seller's responses to price offers. Third, we analyse the data on buyers' ZIP-codes to group observations by geographic regions, namely, East and West Germany. We notice that, while both groups of offers display clustering at exactly 50% of the card's value, the "naïve equal split" is more prevalent among the subjects coming from East rather than West Germany.[4] Furthermore, West German subjects are more likely to make competitive offers, that is, to comply with the theoretical prediction of a model with multiple buyers.

The vast heterogeneity of offers observed in the experiment suggests that the subjects' behaviour cannot be accounted for by a model with homogeneous beliefs, in particular equilibrium beliefs. Therefore, we have to

[3] A *relative offer* is the price offer divided by the card's face value. For example, the value of a *relative* offer that corresponds to 40 Euro for a 50-Euro gift card is 0.8. Transforming the data this way is necessary to make the offers comparable across treatments.

[4] A 50% offer can be viewed as an equal split of trade surplus only under a "naïve" assumption that the seller's outside option is 0.

consider a model of competition uncertainty where the buyers hold heterogeneous beliefs about the seller's outside option c. In this model, a buyer's subjective estimate of the seller's outside option summarizes the buyer's beliefs about the number of effective competitors and their offers. Given his estimate of c, the buyer places an offer at or above the estimate. The amount offered in excess of the buyer's estimate is the share in trade surplus offered to the seller, which we denote s. The price offer is then the sum of two elements, c and the share s in the trade surplus $1 - c$. The identification of c and s is equivalent to decomposing each price offer into its two parts.

It is clearly not feasible to decompose the relative offers at the level of individual observations, since each observed offer is a function of two unknowns and so there are infinitely many possible solutions. However, it turns out that we can implement an *aggregate decomposition* of the observed distribution of relative offers into distributions of shares s and cost estimates c. The aggregate decomposition is a novel statistical tool we bring to the analysis of field data.

Decomposing the observed distribution of relative offers in two underlying unobserved distributions is equivalent to solving an integral equation with two unknown functions. Since the problem is infinite-dimensional, the technical challenge is to reduce the dimensionality to obtain a computationally feasible program. To do so, we restrict the possible distributions to a family of finite polynomials. Starting with the uniform distribution as a candidate solution for both functions we subsequently raise the polynomial degree until the optimal solution within the respective class passes a pre-specified goodness-of-fit test; we use the bootstrapped Kolmogorov-Smirnov test statistic.

The results of this decomposition suggest that nearly nine out of twenty subjects offer between 40 and 50 percent of the total trading surplus to the counter-party. This finding confirms the focal preference for an egalitarian division in the subject population, in line with the theories of Boehm et al. (1993) and Henrich et al. (2006). More than one out of four subjects act consistently with the maximization of individual monetary payoffs, offering no more than 10 percent of the trade surplus to the seller. The average share offered by our eBay users is 29 percent, which is within the range of estimates obtained in the lab experiments on the ultimatum game (see the meta-study by Oosterbeek et al. (2004)).

To summarize, there are two main conclusions from our analysis. First, similarly to a large body of evidence from the lab,[5] the observed data cannot generated by a model where all players are selfish and there is common knowledge of this fact. Second, the uncertainty about the number of players results in large variations of behaviour on eBay as compared to the setting where the player set is common knowledge. Specifically, in contrast to the findings in comparable lab experiments on the ultimatum game with public information on the number of competing proposers, we observe that offers are not driven up to the full nominal value.[6] Therefore, for any buyer-seller pair, there is a non-trivial surplus from trade to be shared. Put differently, in a competitive setting with uncertainty, the buyers ("proposers") have a scope to express their preferences and they use it.

Our paper contributes to several strands of literature. First, it is directly linked to a wealth of papers testing the behaviour in the ultimatum game in different environments. For an excellent overview of this literature, we refer the reader to the recent work by Güth and Kocher (2013) and van Damme et al. (2014), as well as the earlier papers, such as Camerer (2003) and Bearden (2001). A study by Tisserand (2014) provides a meta-analysis of the ultimatum game experiments over the last thirty years (an earlier meta-study is Oosterbeek et al. (2004)).

Our results contribute to an important debate on the external validity of laboratory findings, contrasted, in particular, with field experiments. Comparing the findings to the studies of ultimatum games played in the lab, we establish that the difference between lab and field is fairly meagre when the focus is on surplus division. Concerning gift exchanges, in contrast, List (2006) finds less evidence for social preferences in the field as compared to the lab in otherwise equivalent settings.

[5] An important exception is Cherry et al. (2002).

[6] See, e.g., Roth et al. (1991).

The findings lend further support to the theoretical analysis of pro-social behaviour. The experimental evidence from the lab gave rise to several competing theories of social preferences. The first strand of this literature includes "consequentialist" theories, where agents have preferences only about consequences, that is, preferences over the distribution of final payoffs: Fehr and Schmidt (1999) and Ockenfels and Bolton (2000). The second strand of the theoretical analysis treats agents' behaviour as motivated by reciprocity in response to actions and intentions of opponents: Rabin (1993), Falk and Fischbacher (2006), Dufwenberg and Kirchsteiger (2004), Charness and Rabin (2002), to name a few prominent examples. Given its design, our experiment cannot distinguish between two types of theories. However our decomposition results can be used to calibrate some of these models.[7]

By the nature of data it studies, our paper contributes to the growing economic literature on eBay users' behaviour. Interestingly, this literature focuses predominantly on the effects of users' reputation on transaction prices (e.g., Resnick et al. (2006), Cabral and Hortascu (2010), Nosko and Tadelis (2015)) or the benefits from gaming the mechanism, such as snipe bidding (e.g. Roth and Ockenfels (2002), Ely and Hossain (2009)). An important exception to this literature is due to Bolton and Ockenfels (2014), who use the eBay platform to study the ultimatum game. They design a framed experiment where they hire students and match them via the eBay platform, while keeping the remaining features of the setup similar to the lab. In contrast, our paper reports on a natural field experiment; since seller's outside option cannot be controlled in this setting we use the novel statistical method to estimate the buyer's beliefs about the surplus from trade.

The next section describes the experiment, Section 2.2 presents the data and the regression analysis. The decomposition strategy and its results are presented in Section 4. Finally, Section 5 discusses various aspects of the experiment and Section 6 concludes.

[7] A further important contribution to the theoretical analysis of social preferences is the study of their origin. Alger and Weibull (2013) look at the evolutionary development and rationalization of such preferences. They show that the preferences described as a combination of selfishness and morality, akin to this paper's finding, are sustained in an evolutionary stable equilibrium.

2 Experiment

2.1 Setup

We set up a controlled environment within an existing secondary market for Amazon gift cards on the German site of eBay (`www.ebay.de`). Amazon gift cards are used primarily as presents.[8] A gift giver would buy a sixteen-digit code at the Amazon website in order to transmit it to a gift receiver. The receiver enters the code to top up his Amazon account with the amount paid earlier by the gift giver. The credit can be used for purchasing any goods offered on the Amazon website,[9] it can be split, combined with other payment methods and stored for up to 3 years. If the gift receiver does not intend to use the code, he can offer it for sale at a secondary market such as eBay.

Reselling Amazon gift cards is rather common on eBay.de (eBay Germany). For instance, 87 gift cards were on sale at 7 p.m. on June 13, 2014, and 1962 sales in total were posted within 114 days prior to that date. Nominal values of gift cards ranged from 5 to 2500 Euro.[10] The market turnover in Q4 2015 is estimated at 70 000 Euro.

In the experiment, we employ the *Buy-it-Now or Best Offer* (BINBO) format.[11] According to the rules of BINBO, the seller posts an initial ask price and invites the buyers to acquire the object at that price or to make their own price offer; the duration of sale is limited and the remaining time is public information. When a buyer makes an offer, the seller has 48 hours (unless

[8]Other uses include: rewards for participation in internet surveys, payments in consumer-to-consumer online purchases, for instance trading electronic train tickets.

[9]The goods can be bought from Amazon.com, Inc. / Amazon EU S.a.r.l. as well as any other seller, private or institutional, that uses Amazon as the selling platform.

[10]One may wonder that such expensive goods could be sold anonymously over the Internet. This is due, to a large extent, to efficient consumer protection services offered by eBay, as well as the importance of reputation (see Resnick et al. (2006)). Moreover, Germany ranks high on the level of trust between strangers (see Fukuyama (1995)). On the relation between culture and e-commerce diffusion see Gibbs et al. (2003).

[11]In German it is *"Sofortverkauf oder Preisvorschlag"*. The *Buy-it-Now* sales format is the alternative to the better-known eBay auction (essentially an English auction, where the seller specifies an initial price, typically 1 Euro, and buyers submit their bids within a fixed time period).

the listing expires earlier) to reply with acceptance, rejection, counter-offer or not reply at all (we discuss the role of counter-offers in Section 5). When a price offer is accepted, the card is sold.[12] A listing becomes inactive if it expires or if the card is sold to a buyer. EBay users can browse the history of inactive listings.

The subjects of our study are eBay.de users who open one of the listings we posted and make a price offer. When a buyer opens an active listing he observes the seller's posted price as well as the following information.[13] First, he can observe the total number of current offers received by the seller (0, 1, 2, etc.), and the timing and the current status of each offer (pending, rejected, counter-offer received). Second, the buyer observes the number of hours left before the listing expires, the seller's feedback record, and the other items the seller currently offers for sale. The buyers can also observe the history of previously offered items of this seller. The most immediate information about the seller's history is the list of all feedback entries (positive, negative, neutral) that are left to the seller for the items that he has sold. Importantly, only the seller can see the *amount* behind each offer that he has received. No buyer knows how much other buyers have offered.[14]

In our experiment, we offer Amazon gift cards with the face values: 5, 10, 20, 50 100, 200, and 500 Euro, i.e., in total seven treatments. In two waves of the experiment, March to July 2014 and March 2015, we post over 200 listings. The duration of listings is fixed at the minimal period of 3 days. We use five different seller accounts with feedback records ranging from no feedback on an account opened in 2014 to 430 stars on an account opened in 2004. Each seller account lists one or two gift cards at a time, and the nominal values of the gift cards rotate between seller accounts over the entire duration of the experiment.

We draft the listings in a way that complies with the common practice of

[12]Note that this last feature distinguishes BINBO from eBay's ascending auction, where if at least one bidder submitted a bid the seller is obliged to deliver the good regardless of the price.

[13]The announcement and URL to a screen-cast are given in the Appendix section A.1.

[14]Similarly, in Camerer (2003) and Roth et al. (1991), the proposers do not observe their competitors' offers.

gift card sales via eBay's BINBO. Relying on the history of similar posts, we use the typical wording and set the initial ask prices between 119 and 130 percent of the gift cards' nominal values (e.g., for a 100 Euro gift card we set the BIN price equal to 119 Euros). The choice to set the BIN price above the cards' face value comes across as a surprise, however, there are two reasons for this. First, as mentioned already, setting the BIN price above the face value is the common practice in the market for gift cards, and therefore serves the purpose of mimicking a typical seller. Second, this limits the number of actually executed transactions, since rational buyers should not accept the excessive BIN prices. While some subjects do accept the BIN price, we exclude those observations from the sample and concentrate only on the rational offers that do not exceed the card's nominal value. Finally, to eliminate any variation in the sellers' response to offers, we let all offers expire without a single answer from the seller (though we answered clarifying questions that some buyers were sending us).

As an illustration, we next describe one round of the experiment. We list a gift card with the nominal value of 100 Euro on June 1 at 1:08 p.m. We receive the first offer of 90 Euro from buyer *"lu..er"* with 6 eBay stars on June 3 10:16 a.m. and the second offer of 80 Euros from buyer *"xx...30"* with 60 stars at 1:23 p.m. We let both offers expire on June 4 at 1:08 PM. Thus we get two observations, with the offered amount, exact timing and order of the offer, buyer characteristics (eBay alias, reputation score and registered ZIP code), as well as the seller's information and the exact time the listing was posted. We also keep track of the announcements that did not receive any offers.

2.2 Data

72 percent of the listings receive at least one offer within the three-day period. The number of offers per listing ranges from 0 to 15, with 1.6 offers made on average. One offer per listing is both the median and the most frequent number of offers, and corresponds to the situation of bilateral trade.

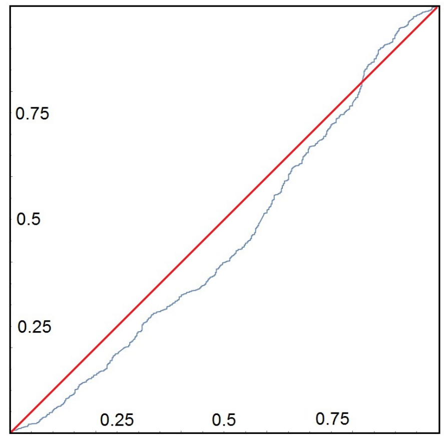

Figure 1: Cumulative distribution functions of the observed (blue), and uniform (red) arrival times, normalized to one.

The subjects come from all over the country, and 15 percent of the offers originate in East Germany or, more precisely, in the former GDR states.[15] The most experienced buyer in the sample was registered on eBay 16 years prior to our experiment; the average eBay experience amounts to about 8.5 years.

The distribution of the arrival times of price offers within the duration of sale is plotted in Figure 1 along with the uniform distribution. The times are normalized to one according to the formula $\frac{t_{\text{offer}}-t_{\text{listing}}}{3\times 24\times 60}$, where t is expressed in minutes. Contrary to the case of eBay's ascending auction, where the bidding frequencies spike at the end of sale (see, e.g., Roth and Ockenfels (2002)), we do not observe any such patterns in the arrival rates in the BINBO format we use in the experiment. This is not surprising since strategic waiting generates no payoff in a BINBO sale.[16]

The offers range from 1 Euro, the lowest admissible offer on eBay, to amounts exceeding the nominal value of the gift card. Since the latter offers are

[15] Five offers from Austria were excluded from this part of the analysis.

[16] In BINBO, the trade can take place any time within the duration of the listing; therefore waiting is risky since another buyer can strike a deal in the meantime. Second, much less information is revealed in BINBO compared to the ascending auction, and thus waiting does not create any real information advantage.

Voucher Value	All	5 €	10 €	20 €	50 €	100 €	200 €	500 €
No. of Listings	221	46	25	19	22	36	43	30
No. of Offers	358	42	45	38	57	60	74	43
Average Bid	0.73	0.71	0.77	0.73	0.77	0.70	0.72	0.71
Std. Err.	0.04	0.11	0.11	0.12	0.10	0.09	0.08	0.11
Std. Dev.	0.23	0.23	0.18	0.17	0.17	0.27	0.23	0.29
Median	0.80	0.80	0.80	0.75	0.80	0.80	0.79	0.82

Table 1: Descriptive statistics of relative offers "Offer / Nominal Value".

clearly irrational, we exclude them from the analysis. To make the data comparable across treatments, we normalize the offers by dividing each offer by the nominal value of the gift card. The pooled data display clustering: the relative offers concentrate around 0, 50, 80, and 90 percent of the gift card value (see the figures in the Appendix, A.3). The descriptive statistics, broken down with respect to nominal value treatments, are reported in Table 1. Overall, the distribution of offers is right-skewed: every second offer exceeds 80 percent of the nominal value. The average offer in our sample ranges from 71 to 77 percent of the nominal value and does not display monotonicity with respect to the gift card's nominal value. The same is true of median values that range from 75 to 82 percent of the nominal value.

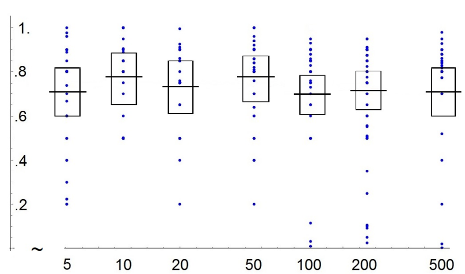

Figure 2: Relative offers (y-axis) vs. gift card nominal values (x-axis, log-scale). Data points in blue, the average and 1-standard error band in black.

We use a one-way ANOVA test to verify formally whether the empirical averages vary across the nominal value treatments. We find the null hypothesis of equal means is not rejected,[17] implying the relative offers' invariance in scale. The analysis of pairwise linear and log-linear regressions produces the same result: the nominal value does not have a statistically significant effect on the relative price offers (see Table 4). Figure 2 presents the data points and the fitted regression line,

[17] ANOVA *P-value* $= 0.50$.

where the logarithm of the nominal card value is plotted against the horizontal axis and the relative offer is the dependent variable on the vertical axis. Overall, our analysis indicates that stake size does affect the relative offers in the population of eBay buyers. Similar findings were obtained in the lab settings by Cameron (1999), Munier and Zaharia (2002), Hoffman et al. (1996), Slonim and Roth (1998).[18]

3 Effects of the Observables

In this section, we report on details of the statistical analysis that we conducted to relate all observable characteristics of the buyers to their offer. By the design of the experiment, we know exactly what information was available to the buyer at the time he made the offer, as well as several characteristics of the buyer himself. The dependent variable in all regressions is the relative offer b_i, defined as follows: If V_i is the face value op the gift card and B_i the price offer then $b_i = \frac{B_i}{V_i}$.

3.1 Competition Marks

We start by looking at the effects of information about the number of competitors for a gift card. In BINBO, buyers observe two signals informative of competition intensity. The first signal is amount of time remaining before the listing expires: the more time is left, the more buyers are expected to arrive by the end of sale, and the higher is the degree of competition. The second signal is the number of offers already outstanding by the time a given buyer makes his offer; naturally this signal conveys the information in a more direct way. Both indicators are displayed next to the BIN price and barely involve any search effort (see the screen-cast URL p. 29). Both indicators should make buyers update upwards their beliefs about the seller's outside option. This, in turn, should have an effect on the buyers'

[18] A meta-study by Oosterbeek et al. (2004) documents a small negative effect of increasing stake size. Slonim and Roth (1998) and Roth and Erev (1995) find a positive effect when the game is played in the first round or only played once.

relative offers. However the regression results demonstrate that both signals have insignificant effects on the subjects' relative offers (see Table 4 in the Appendix).

We take a further step to verify whether the simplest binary indicator of competition produces an effect. Specifically, we split the offers in two groups: those arriving *first* on a listing and those arriving when at least one other has been already made. The respective empirical distributions are presented in Figure 6 in the Appendix. Again, we find no significant difference in both groups' mean offers.[19]

Beside the obvious explanation that any updating, however simple, is inhibited by cognitive costs, we offer two alternative reasons why the competition marks fail to produce any significant effects in the experiment. First, the subjects may use a rule of thumb when making offers, drawing on their past experience of gift card sales. Such an approach is aimed at a longer-term market performance and substantiated by a large bulk of psychological literature, e.g., Tversky and Kahneman (1974), Newell et al. (1972), Gigerenzer (2007). Second, the buyers may (rationally) expect that if the gift cards remain unsold by deadline they go on sale again in the future. Therefore, the expected stream of future offers may outweigh any present competition, making the latter effect statistically insignificant.

3.2 Subjects' Learning

Next, we test whether the relative offers change over the course of our experiment, which may occur due to subjects learning about the environment. The subjects get information about the seller's response strategy from two sources. First, the history of sales was available through eBay's search engine at the time we conducted the experiment. For all five accounts used in the experiment, browsing the history of the account revealed that a number gift cards were listed exclusively in BINBO format and remained unsold.[20]

[19] ANOVA *P-value* $= 0.13$.

[20] Apart from a few cases when we received offers to pay the posted price that exceeds the card's value.

While this information could have stopped some buyers from making an offer, the regression analysis suggests that the *size* of relative offers was unaffected by history (see Tables 5, 6 in the Appendix). To proxy the opportunity for learning from history, we use calendar time: subjects participating at the start of the experiment have less evidence on our response behaviour than those arriving toward the end of the experiment. This finding implies that the subjects did not browse sale histories or did not take the information into account.

Second, buyers can learn from their own experiences of making an offer to a particular seller. To study the possible effect, we look at the sub-sample of the recurrent buyers. In total, 56 out of 277 buyers of our entire experiment made offers on multiple listings. We track how those buyers' offers change over time. Specifically, we calculate the increment of each subsequent offer relative to the previous offer that buyer made. This variable captures the subject's learning dynamics due to his or her experience with one of the seller accounts we use. Performing the Student's t-test we find that there is no statistically significant change in offers between two consecutive rounds of a subject's participation.[21]

3.3 Effects of Experience

Buyers' eBay experience, reflected by their feedback score,[22] has a mild positive effect on relative offer sizes (significance at 10%; see Tables 5, 6, 7 and 8 in the Appendix). Even if the effect is present, its size is extremely small. For instance, an inexperienced buyer with no feedback offers 1 percentage point less than the average buyer with 450 feedback entries.

A similarly sized effect, also statistically significant, is produced by the increase of the seller's feedback score. However, while the effect of buyer experience is linear, the marginal effect of the seller's experience decreases.

[21] $P\text{-}value = 0.317$.

[22] Note this is an imperfect measure of experience, due to the benevolent nature of ratings. Especially the sellers usually have very few incentives to report about buyers as by eBay's regulations the only feedback that they can leave is positive feedback.

Specifically, the effect on the relative offer of extra feedback that the seller receives decreases drastically after just 12 feedback entries.[23] The effects of buyer and seller experience do not display complementarity, i.e., there is no evidence that more experienced buyers make significantly higher or lower offers to more experienced sellers.

3.4 East and West Germany

Using data about buyers' registered ZIP-codes, we identify the impact of each subject's location on the size of the offer. Our sample contains about 15% of offers stemming from the states on the former GDR territory.[24] Our main observation here is that when subjects are grouped by region, there are important differences in the distribution of offers between East and West Germany. In the latter group, we observe a larger number of offers that are close to the competitive prediction – near 95 percent of the gift card value. By contrast, the buyers from East Germany make more offers in a close neighbourhood of 50 percent of the total value. (See Fig. 5 in the Appendix).

Can this variety in offers be attributed to the remaining cultural differences between East and West Germany? From the previous analysis we know that a buyer's experience affects the average size of his offer to the seller: more experienced buyers tend to make slightly higher offers. Since West German buyers in our sample have more experience with eBay, the regional difference we observe may be due to the difference in experience. In order to correct for the possible bias, we extract a sub-sample of buyers from West Germany that has the same the distribution of feedback scores as the East German sample. After the correction, the distributions of relative offers remain virtually unchanged and the same difference patterns emerge. The equal split of the "naive surplus" is a significantly more important focal

[23]There is a statistically significant effect of seller experience when we use the feedback score from 5 accounts jointly (*P-value* is equal to 1.9×10^{-4}). Splitting the sample into two groups, for sellers with zero feedback and at least 12 feedback entries results in *P-value* of 0.79 and 0.68, respectively.

[24]ZIP codes are only indicative of current residency, not the subject's origin.

point for East German subjects, while the competitive offers are more common among the West German subjects. In the next section, we also report on the difference between two regions at the level of offered shares of surplus.

For instance, in a large scale empirical study Alesina and Schuendeln (2005) found that East Germans displayed higher preference for equality and redistribution – something that could reinforce the prevalence of equal splits. In two waves of a public good experiment, Ockenfels and Weimann (1999) and Brosig-Koch et al. (2011) find important differences in East and West German behaviour, which persisted two decades past the reunification. Focusing on children and adolescents aged 10 to 18, John and Thomsen (2013) find more support for other-regarding preferences in East than in West Germany.

4 Decomposition

4.1 General Framework

In this section we study the incentives behind the subjects' offers in the experiment. Figure 3 presents the pooled data from BINBO sales: the relative offers $b_i = \frac{B_i}{V_i}$ are grouped into five-percent bins and plotted along the horizontal axis; each column's height corresponds to the number of offers falling into the bin.

The most striking feature of the experimental data is the vast heterogeneity and the irregular clustering of relative offers, and in particular, the spikes at around 0, 50 and 80-90 percent of the card's value. It is clear that the subjects' behaviour cannot be accounted for by any model with homogeneous beliefs, and in particular an equilibrium model. Consider for instance a standard model where the buyers believed that the environment was competitive on the buyer side and the beliefs were common knowledge. In this setting most offers would be confined to a close neighbourhood of 1 (this theoretical prediction is replicated in lab experiments: see Roth et al. (1991), Fischbacher et al. (2009)). Similarly, the model of BINBO as a pure bilat-

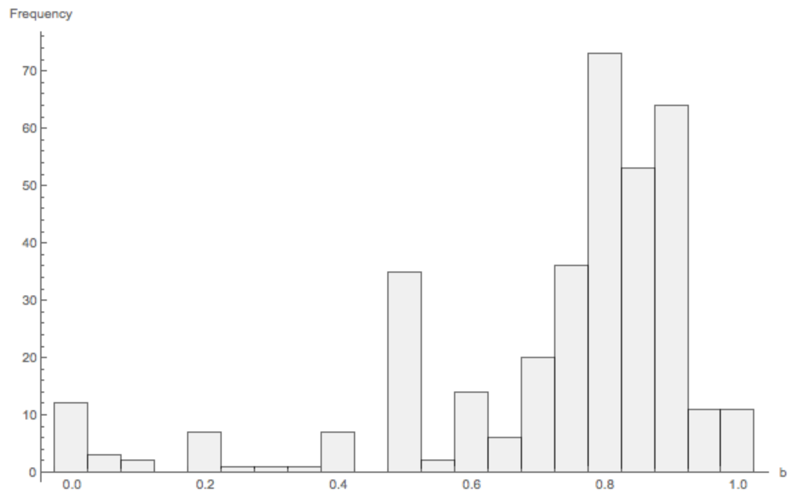

Figure 3: Empirical frequency of relative offers in the experiment

eral trade situation is hardly consistent with the data. Note that even if a "reasonable" intensity of social preferences is allowed for, the bilateral trade model cannot account for the majority of observations.[25] Overall, the observed pattern of offers demonstrates that there is a stark heterogeneity in subjects perceptions of the game and (or) their intentions when playing it.

Our analysis of the observed offers allows for two major sources of heterogeneity between subjects. First, the subjects may differ in their beliefs about what they are competing against. In particular, each subject's offer has a chance to be accepted only if it is larger than both the potential competitors' highest offer and the value that the seller can extract by using the card himself.[26] We impose no restrictions on the subjects' estimates, except that they are contained in the interval $[0, 1]$. The estimate, denoted c_i, is thus the seller's alternative cost of interacting with buyer i from the point of view of that buyer. It is also the theoretical lower bound on buyer i's offer.

We deliberately remain agnostic as to the way our buyers construct their estimates of c. In particular, we do not require rationality or any knowledge

[25] See the standard theories of social preferences, such as Fehr and Schmidt (1999) [Proposition 1], Ockenfels and Bolton (2000) [Statement 3].

[26] Formally, $c_i = \max\{u_{s,i}, \max\{b_{-i}\}\}$ where $u_{s,i}$ is the seller's own usage value from the viewpoint of subject i, and b_{-i} the set of competing bids, possibly discounted, with the convention $\max\{\emptyset\} = 0$.

of the distribution functions on the part of our subjects. This is an important feature of our identification strategy that grants full flexibility to our estimates. Our only assumption is that each player's beliefs about c are a degenerate probability distribution function. In the discussion section 5.4, we show that c_i can be derived as a payoff-maximizing bid in a standard auction setting; the same value corresponds to the mathematical expectation of the competing offer.

The second source of heterogeneity between subjects is the amount of money that subject i offers in excess of his estimate c_i. Consider a fixed c_i, the surplus from trade in the given between the seller and buyer i equals $1 - c_i$. A price offer $b_i \geq c_i$ translates into a division of surplus $1 - c_i$ between the buyer and the seller. When the card's face value is 1 (normalized), buyer i's payoff from trade at price b_i is $1 - b_i$.

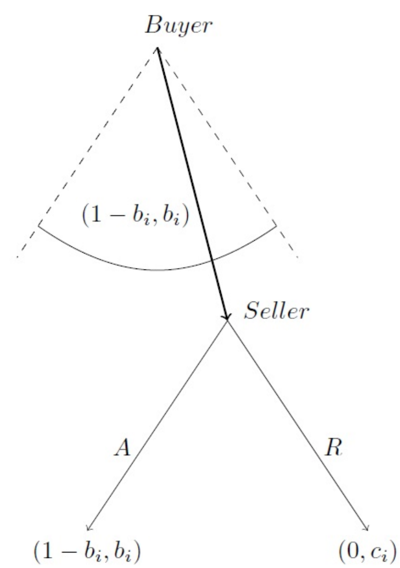

Figure 4: BINBO game tree.

We say that a buyer uses sharing rule s_i if his price offer leaves the seller with the share s_i of the surplus: $s_i = \frac{b_i - c_i}{1 - c_i}$, corresponding to the Kalai and Smorodinsky (1975) solution. Thus, an offer can be represented, or *decomposed*, as follows:

$$b_i = c_i + s_i (1 - c_i) . \tag{1}$$

Equation (1) demonstrates that the observed variation in offers has two sources. First, the subjects vary in the relative shares s_i they are willing to offer the seller. Second, subjects differ in their estimates c_i of the seller's cost c. To understand the prevalence of sharing rules in the data, we have to extract s_i and c_i from the observed relative offers b_i. Clearly c_i and s_i cannot be identified from a given relative offer b_i,i.e., decomposition at the level of an individual observation is infeasible. However, as we show below, one can decompose the offers *on aggregate*, by splitting the observed distribution of relative offers into a distribution of shares s_i and a distribution of cost estimates c_i. We drop the subscript i in what follows.

To define the decomposition problem, we let $f(s)$ and $g(c)$ denote, respectively, the unobserved distributions of the shares s and of the cost estimates c. Let $h(b)$ be the observed distribution of relative offers. Capitals F, G, and H denote the corresponding cumulative distributions. For simplicity we think of all three distributions as having continuous supports. We assume that the support of s is bounded above by 0.5, in line with the standard other-regarding preference theories (see footnote 25). The "selfish" preferences, corresponding to $s_i = 0$ for all i are possible under this specification.

Assuming that sharing rules and cost estimates are independently distributed in the population, the distributions are related by the following:

$$\begin{aligned} H(b) &= \Pr(c + s(1-c) < b) = \Pr\left(c < b, s < \frac{b-c}{1-c}\right) \\ &= \int_0^b \left[\int_0^{\frac{b-c}{1-c}} f(s)\, ds\right] g(c)\, dc = \int_0^b F\left(\frac{b-c}{1-c}\right) g(c)\, dc. \end{aligned} \tag{2}$$

The cumulative distribution function H on the left-hand side is given by the observations in our experiment. The right-hand side integrates over all c and s that generate an offer less or equal to b, according to (1). Our goal is to find $F(s)$ and $g(c)$ that best fit equation (2) given $H(b)$ given by the experimental data.

Two remarks are in order. First, as both $F(.)$ and $g(.)$ belong to infinite-dimensional spaces, problem (2) itself is infinite-dimensional and, therefore, is computationally hard. Hence, to find a solution, we need to reduce the problem's dimensionality to a point where optimization becomes feasible from a computational viewpoint. Second, we know that the solution to (2) exists but may not be unique unless we restrict the space of functions where $F(.)$ and $g(.)$ belong to.[27]

[27] There always exists a meaningless corner solution, where $h \equiv g$ and f is a Dirac delta

Our approach to decomposition addresses both issues. We employ three alternative methods to cross-check the findings, thereby we deal with the uniqueness problem. Two of our methods are parametric; dimensionality reduction is achieved by restricting the solution to belong to a space of parametric functions. In particular, in the first method reported below, we follow the standard approximation theory and look for a solution in the space of polynomials.[28] The second method, reported in the Appendix restricts f and g to the class of beta functions, leaving us with only four parameters to estimate. Finally, in we consider a non-parametric approximation method, where the dimensionality of the problem (2) is reduced by just discretizing the supports of f and g. All three methods yield similar results.

4.2 Polynomial Approximation

One of the most obvious ways to achieve dimensionality reduction is to look for a solution within a smaller set of functions and in particular within a space of parametric functions. Following the standard approximation theory, the best parametric class of functions is the space of polynomials. Hence, our approach is to identify $\hat{H}$ as a polynomial approximation of H. If $\hat{F}$ and $\hat{g}$ are polynomials with degrees n^F and n^g, then $\hat{H}$ is a polynomial itself and has $n^F + n^g + 2$ of parameters to be estimated.

We consider $\hat{f}$ and $\hat{g}$ to be equal to a linear combination of Chebyshev polynomials:

$$\hat{F}\left(\frac{b-c}{1-c}\right) = \sum_{k=0}^{n^F} \gamma_k T_k\left(\frac{b-c}{1-c}\right),$$

and

function with the entire probability mass concentrated at zero. (The symmetric solution does not exist since the support of f is bounded at 0.5.) Sadovnichy (1986) shows that, for a given kernel function $F(b,c)$, a Volterra integral equation such as (2) is a contraction mapping and hence has a unique solution g^*. Note that since both f and g are restricted to the class of probability measures, summing up to 1, Sadovnichy's result does not automatically imply that multiple solutions can be generated by just varying the kernel function.

[28]Following the Stone-Weierstrass theorem, the space of polynomials is dense in $\mathcal{C}[a,b]$ which means that any continuous function on a bounded interval can be approximated arbitrarily well by a polynomial (of arbitrary large degree).

$$\hat{g}(c) = \sum_{i=0}^{n^g} \delta_k T_k(c),$$

where $T_k(x)$ is a Chebyshev polynomial of the first kind of degree k, γ_k and δ_k are the weights attached to the corresponding polynomial of degree k.

Fixing the class of functions, our remaining goal is to find the degree of polynomial n^g and n^F, together with the values of the vector of parameters (γ, δ) that minimizes the Kolmogorov-Smirnov distance between $\hat{H}$ and H:

$$(\gamma^*, \delta^*) \equiv arg \min d_{KS}\left(\hat{H}, H\right)$$

where

$$d_{KS}\left(\hat{H}, H\right) \equiv \sup_{b \in [0,1]} \left|\hat{H}(b) - H(b)\right|, \tag{3}$$

and $\hat{H}$ is given by:

$$\hat{H}(b) := \int_0^b \int_0^{\frac{b-c}{1-c}} \hat{f}(s)\, \hat{g}(c)\, dsdc. \tag{4}$$

As the very first step of our search procedure, we generate a continuous version of $H(.)$.[29] For this, we derive a non-parametric kernel density function from the data, with the bandwidth 0.015. This bandwidth strikes a balance between the technical requirement of continuous $H(.)$ and the preservation of all information contained in the data.

Once $H(.)$ function is obtained we begin the main part of the search procedure. Namely, starting from the first degree Chebyshev polynomial, we iteratively increase the degrees in $F(.)$ and $g(.)$ and test whether the corresponding function $\hat{H}(.)$ is statistically indistinguishable from the empirical $H(.)$. In other words, we dismiss a solution if the respective Kolmogorov-Smirnov distance exceeds the bootstrapped critical value and we proceed to the next step where we add one extra polynomial term.[30] At $n^F = n^g = 4$ the hypothesis that $\hat{H} = H$ is no longer rejected. Following the standard

[29] The reason for generating continuous H is two-fold. First, our approach is then validated by the classical Stone-Weierstrass Theorem. Second, continuity is required by the software we use, in order to obtain a robust solution.

[30] We used `Mathematica 10` and its built-in differential evolution fitting method, allow-

approach to avoid overfitting we terminated the procedure at that point. However, we checked for polynomial degrees beyond 4 and up to 10; this never improved the distance significantly and more importantly, each subsequent solution was qualitatively equivalent to the above solution. Thus within the class of polynomials of up to the 10th degree, the estimated f and g correspond to the statistically optimal solution (which is also unique).

4.3 Decomposition Results

The estimated distributions are presented in Table 2.[31] Our main finding about the distribution of sharing rules f, is that 43-44 % of the subjects offer 40 to 50 % of the trade surplus to the seller. One third of all our subjects make "greedy" offers and propose no more than 10 % of the trade surplus to the seller. The remaining 25-30 % of the subjects make offers between 10 and 40 %. The average offer to the seller amounts to 30 % of the respective trade surplus.[32] Hence, our results are not completely aside from the result in the previous literature, the surplus share offered on eBay.de is within the range of estimates obtained in laboratory ultimatum games, however it yet stays close rather to the lower end of a cross-country distribution of offers (see, e.g., Oosterbeek et al. (2004)).

s:	0-10%	10-20%	20-30%	30-40%	40-50%
$\hat{f}^*(s)$:	33.2	3.7	3.6	16.6	43.0

c:	0-10%	10-50%	50-60%	60-70%	70-80%	80-90%	90-100%
$\hat{g}^*(c)$:	11.2	10.1	11.8	17.7	21.1	19.2	8.9

Table 2: Estimated distributions of sharing rules $\hat{f}^*(s)$ and cost reference points $\hat{g}^*(c)$ in the population of eBay buyers (in percentage points).

The found $g(.)$ function suggest that there are two types of beliefs that prevail in the population of our bidders. There is a minority of subjects, esti-

ing at each search round for 500 iterations. The differential evolution method is known for providing consistently the global solutions.

[31] The KS distance is 0.028. The corresponding Chebyshev coefficient estimates are given in Table 10 in the Appendix.

[32] 29.4 % according to the parametric and 29.8 % according to the non-parametric estimates (see Appendix A.4.1).

mated at 11-14%, who act as if they faced no competition from the other buyers; those subjects make their offers under the assumption that the seller's outside option is null. The beliefs of the remaining majority of the subjects, however, are described by a bell curve centered near the *empirically observed average offer*. This suggests presence of the "wisdom of the crowd" in the eBay market of Amazon gift cards – on average estimates in this group of subjects are correct, even if some noise around that estimate is still present.

4.4 Robustness Checks

To verify whether the decomposition we found is robust, we have conducted further decompositions, using two alternative approaches. One approach, reported in Appendix (A.4.1), is to discretize the support of distributions h, f, and g, in order to obtain a finite-dimensional problem. Namely, we partition the support of the three distributions into equal-length bins. Given the upper limit on the support of sharing rules of 50%, by letting the bin width to be equal to 10%, we obtain a substantially simplified problem with just 15 unknown values. We use an iterative numerical optimization to assign probability mass to each of the bins. Both approaches obtain similar distribution functions as the solution.

Our second robustness check, also explained in details in the Appendix A.4.2, employs another parametric procedure that confines both f and g to the class of beta probability distributions and thus reduces the entire problem to only four dimensions (each beta distribution is described by two parameters). While the resulting fit measured by the KS distance is unsatisfactory, the similarity of the beta functions solution to our main result suggests that the findings are robust.

5 Discussion

5.1 Competition with other sellers

As the sellers of gift cards, we face competition from other sellers present on eBay throughout the entire duration of the listings. However, for our experiment, such competition is hardly relevant, given that a gift card that is bought at a discount gives to an Amazon customer a payoff equivalent amount of cash. As long as the preferences for cash are at least locally insatiable, the demand for transaction is, by and large, independent from the fierceness of competition among the sellers.

5.2 First stage of a multi-stage bargaining game

The BINBO trading format allows a buyer and a seller to exchange up to 3 offers in total.[33] We concentrate on the first stage of the (potentially) multi-stage bargaining situation, namely on the initial offers made by buyers, for two reasons. First, the very first offers are most indicative of the subjects' true sharing intentions, i.e. the minimum of surplus that they would like to give to the seller. Second, our data do not support the hypothesis that subjects intend to play a multi-stage game. If this intention was systematically present, then (*i*) most offers would arrive on the first day of the listing, and (*ii*) low offers would be made more frequently than higher offers within the first hours. Figure 1 in Section 2.2 suggests that there is no systematic difference in the timing of arrival across 3 days of listing, therefore (*i*) does not hold. Table 4, and more generally the results reported in Section 3.1, imply that (*ii*) is not true either.

[33] If the seller makes a counter-offer to the buyer, then the buyer's initial offer is cancelled and the seller must wait until the buyer's reaction. The seller has no obligations vis-à-vis the buyer, that is, he can accept another buyer's offer and sell the item in the meantime. EBay rules allow for at most three exchanges of counter-offers between a buyer and a seller. In practice, the sellers barely ever respond with counter-offers, thus the buyers should rationally expect to have only one chance to call a price.

5.3 Risk-neutrality assumption

We have assumed risk neutrality of buyers; this is supported by our own data and for the size of stakes that we used, it is confirmed also by the existent literature. In particular, given the empirical evidence for the increased risk aversion at higher monetary stakes (see, e.g., Holt and Laury (2002) and references therein), if buyers were risk-averse we would observe a reduction, on average, in offers for gift cards of high nominal value, e.g. 200 or 500 Euro.[34] The regression results do not support this hypothesis (see Tables 4-9 in Appendix). Similarly, the paper by Fehr-Duda et al. (2010) demonstrates that for amounts similar to our stakes the risk aversion is not identifiable in the data.

5.4 Buyer's beliefs about c

We have assumed that the buyers hold degenerate beliefs about the seller's outside option, and we referred to the mass point of buyer i's beliefs as c_i. However the assumption that the buyer's beliefs about the seller's cost is degenerate can be dispensed of. Recall that the possible value range for c_i is $[0,1]$ and suppose that the distribution function $\Phi_i(\tilde{c}) = \tilde{c}^{\alpha_i}$, $\alpha_i > 0$, $\tilde{c} \in [0,1]$, represents buyer i's belief about the seller's random cost $\tilde{c}$. The mathematical expectation of the random variable $\tilde{c}$ is given by $\int_0^1 \tilde{c}\, d\Phi_i(\tilde{c}) = \frac{\alpha_i}{\alpha_i+1}$. Therefore, if we set $c_i = \frac{\alpha_i}{\alpha_i+1}$ we obtain the first interpretation of c_i as buyer i's expectation of the seller's cost. Furthermore, suppose now buyer i is "selfish"; thus, his sharing motive s_i is zero in the model. The selfish buyer maximizes his expected payoff $(1-b_i)\Phi_i(b_i)$ with respect to b_i. The solution is $b_i^* = \frac{\alpha_i}{\alpha_i+1} = c_i$. Thus, the same c_i can be interpreted as either the expected cost and payoff-maximizing bid of a selfish bidder. Since c_i is a monotonically increasing in α_i, any heterogeneity in the belief parameters α_i generates a variety of c_i in the model.

[34] In 2014, the monthly disposable median net income per capita was 1644 Euro in Germany (European Commission data: http://ec.europa.eu/eurostat/web/gdp-and-beyond/quality-of-life/median-income)

6 Conclusion

Although the importance of the field experiments is often emphasized,[35] an overwhelming majority of experiments on surplus division in a bargaining context are conducted in the lab, following the ultimatum game of Güth et al. (1982).[36, 37] One possible reason for the relative scarcity of field experiments is the difficulty of finding a setup that would be as simple as the classical ultimatum game. Our experiment proposes one such natural environment. The setup features a non-standard (non-student) subject pool, a field context in the commodity and in the information set that the subjects can use; most importantly, our subjects naturally undertake the task and do not know that they are in an experiment. Also, by contrast to the lab, our subjects (buyers or "proposers") match with the experimenter's agents (sellers or "responders") in a way that we do not control. On the other hand, while uncontrolled elsewhere, our setup features perfect control of the buyers' valuations of gift cards and so we have a reduced number of unobserved characteristics that affect the behaviour of participants.

The main distinction of our setup from a typical lab experiment on the ultimatum game is the uncertainty of the buyer about the overall number of competitors.[38] We consider this feature of our experiment to be rather opportune. Compared to the settings with the commonly observed degree of competition, we do not observe at the buyers' side rallying up of offers. Hence, there is no unravelling of the surplus that is created in a match of a given buyer and the seller (remember that the surplus in our setting is defined *relative to the competitor's offer*). The presence of a non-trivial surplus to share allows buyers, in turn, to express more saliently their social preferences.

On the other hand, because of uncertainty about the degree of competition,

[35] See Levitt and List (2009), List (2011), and Galizzi and Navarro Martinez (2015).

[36] Overall, laboratory experiments still constitute about three quarters of all experimental literature, see Card et al. (2011).

[37] As discussed in the Introduction, notable exceptions from this literature are Güth et al. (2007) and Bolton and Ockenfels (2014).

[38] In contrast, Roth et al. (1991), Fischbacher et al. (2009) look at the behaviour within the ultimatum game where the number of competing proposers is commonly known.

the beliefs of our subjects are less predictable than in the lab. This brought us to search for a new methodological approach to handle the experimental data and to measure social preferences. As the result, our paper comes up with a new estimation technique that permits us to decompose the observed distribution of offers into the distribution of unobserved beliefs about the seller's opportunity cost and the distribution of unobserved sharing rules as proposed by the buyers.

Our main findings suggest that even at a large one-shot interaction market like eBay participants offer an equal splitting of surplus. In particular, we find that up to 44% of players offer roughly a half of surplus, while selfish players (offering 0-10% of surplus) constitute less than a third of all our subjects. The estimated frequencies of the various shares of surplus offered to the counter-party are consistent with the large body of estimates emerging from the lab studies. This implies, that despite all discrepancies between the lab and the field, lab experiments on surplus division provide a valid insight into the field behaviour, both qualitatively and quantitatively.

We have also collected new and puzzling evidence about the distribution of attention of market participants to the strategically relevant public information. About 55% of buyers, who are given public information that there is a competitor who has made already an offer, still offer an amount well below the theoretical prediction for the ultimatum game with competing proposers.[39] This is also reflected in our findings during the decomposition stage, in the estimated distribution of beliefs. Specifically, we document a significant heterogeneity of subjects' beliefs, with a small but visible cluster at 0-10% of the face value and a large majority scattered around the average relative offer observed in our data.

All in all, we believe our study sheds some first light on the scope of surplus sharing among actual participants of large and competitive markets. The eBay platform, featuring sequential arrival of buyers, is a fair representation of a generic market place found all over the world and across time.

[39]Note also that in a theoretical model with uncertain competition, second and later arriving buyers still offer *on average* more than first arriving buyers. We do not observe such increase in average offers.

Most of those markets are dynamic and impregnated with uncertainty – the population of traders evolves over time, some participants quit the market, while new participants enter at random. Traders do not possess perfect and common information, instead they learn about going prices and the fierceness of competition via their own private experiences of bargaining with different sellers or through observations of how markets clear; all in the same way as it happens on eBay. Our analysis suggests that because of such uncertain and evolving degree of competition, even a large market offers a scope to share the surplus non-trivially. Of course, more research is needed to in order to understand how exactly competition and strategic uncertainty interact with individual preferences for pro-social behaviour. However, we believe that our paper offers a first and important step in this direction.

A Appendix

A.1 Experiment

An example text used in the description box of a listing (German):

> `"Biete hier einen Amazon-Gutschein im Wert von 50 Euro an, gültig bis zum xx.xx.20xx. Der Gutscheincode wird nach Zahlungseingang auf meinem Konto am gleichen Tag via eBay Mitteilung versendet. Es handelt sich um echte Geschenkgutscheine (keine Aktionsgutscheine!), d.h. sie haben keinen Mindestbestellwert, es können mehrere Gutscheine kombiniert werden, und eventuelles Restguthaben verbleibt auf dem Amazon-Kundenkonto."`

English translation:

> `"I offer here an Amazon gift card worth 50 Euros, valid until xx.xx.20xx. The card code is sent to you via eBay message on the same day I receive payment. These are actual gift certificates (not promotion coupons!), that is, there is no minimal purchase requirement, multiple coupons can be combined, and any residual credit would remain on the Amazon account."`

Watch a video illustration at:[40]

https://www.youtube.com/watch?v=cDeTSl0cAXk.

[40] ATTENTION Anonymous Referee: YouTube tracks viewers' information, please take the necessary precautions to protect your anonymity.

A.2 Tables

A.2.1 Regression Analysis

Variable	Description	Range
relative offer	Offer divided by the card's nominal value	0.02 .. 1
nominal	Gift card's nominal in euros	5 .. 500
time to deadline	Time left to listing expiry when offer is made divided by listing duration	0.0007 .. 0.9988
order	1, if the offer arrives first on the listing, 2, if second etc.	1 ..15
trend	Time elapsed since the arrival of the first offer, in days	0 .. 359.20
buyer_exp	Number of buyer's eBay stars when he makes offer	1.. 8847
seller_exp	Number of seller's eBay stars when he receives offer	1 .. 511

Table 3: Regression variables.

Dep.: relative offer	Model 1	Model 2	Model 3	Model 4	Model 5
constant	0.736***	0.751***	0.737***	0.741***	0.762***
nominal	$-6.5 \cdot 10^{-5}$		$-6.5 \cdot 10^{-5}$	$-6.8 \cdot 10^{-5}$	$-6.4 \cdot 10^{-5}$
log(nominal)		$-5.8 \cdot 10^{-3}$			
time to deadline			$-1.8 \cdot 10^{-3}$	$2.1 \cdot 10^{-3}$	$4.6 \cdot 10^{-3}$
order				$-2.1 \cdot 10^{-3}$	$-4.1 \cdot 10^{-3}$
trend					$-1.2 \cdot 10^{-4}$
N obs	358	358	358	358	358
R sq.	0.001	0.001	0.002	0.002	0.006

Table 4: Regression models 1-5.

Dep.: relative offer	Model 6	Model 7	Model 8	Model 9	Model 10
constant	0.733***	0.747***	0.751***	0.763***	0.753***
buyer_exp	$2.9 \cdot 10^{-5}$*	$3.0 \cdot 10^{-5}$**	$2.8 \cdot 10^{-5}$ *	$2.9 \cdot 10^{-5}$*	$2.8 \cdot 10^{-5}$ *
trend	$-1.24 \cdot 10^{-4}$	$-1.1 \cdot 10^{-4}$	$-1.4 \cdot 10^{-4}$	$-1.5 \cdot 10^{-4}$	$-1.1 \cdot 10^{-4}$
order		$-5.0 \cdot 10^{-3}$	$-5.8 \cdot 10^{-3}$	$-5.2 \cdot 10^{-3}$	$-5.9 \cdot 10^{-3}$
time to deadline			$8.9 \cdot 10^{-3}$	$8.9 \cdot 10^{-3}$	$11 \cdot 10^{-3}$
nominal			$-4.8 \cdot 10^{-5}$		$-5.0 \cdot 10^{-5}$
log(nominal)				$-4.55 \cdot 10^{-3}$	
offer number					$-2.0 \cdot 10^{-5}$
N obs	358	358	358	358	358
R sq.	0.01	0.01	0.01	0.02	0.02

Table 5: Regression models 6-10.

Dep.: relative offer	Model 11	Model 12	Model 13	Model 14	Model 15
constant	0.684***	0.695***	0.69***	0.701***	0.716***
log(buyer_exp)	$1.2 \cdot 10^{-2}$*	$1.2 \cdot 10^{-2}$*	$1.2 \cdot 10^{-2}$ *	$1.21 \cdot 10^{-2}$ *	$1.2 \cdot 10^{-2}$*
trend	$-1.18 \cdot 10^{-4}$	$-1.3 \cdot 10^{-4}$	$-1.4 \cdot 10^{-4}$	$-1.3 \cdot 10^{-4}$	$-1.43 \cdot 10^{-4}$
order		$-4.2 \cdot 10^{-3}$	$-4.7 \cdot 10^{-3}$	$-4.9 \cdot 10^{-3}$	$-4.5 \cdot 10^{-3}$
time to deadline			$11.6 \cdot 10^{-3}$	$10.3 \cdot 10^{-3}$	$10.3 \cdot 10^{-3}$
nominal				$-5.8 \cdot 10^{-5}$	
log(nominal)					$-6.16 \cdot 10^{-3}$
N obs	358	358	358	358	358
R sq.	0.01	0.012	0.012	0.014	0.014

Table 6: Regression models 11-15.

Dep.: relative offer	Model 16	Model 17	Model 18	Model 19	Model 20
constant	0.703***	0.736***	0.742***	0.714***	0.723***
buyer_exp	$2.41 \cdot 10^{-5}$	$2.68 \cdot 10^{-5}$*	$2.55 \cdot 10^{-5}$	$2.43 \cdot 10^{-5}$*	$2.53 \cdot 10^{-5}$*
seller_exp	$9.9 \cdot 10^{-5}$*	$1.09 \cdot 10^{-4}$ *	$1.17 \cdot 10^{-4}$ *		
log(seller_exp)				$1.60 \cdot 10^{-2}$***	$1.52 \cdot 10^{-2}$***
trend		$-1.2 \cdot 10^{-4}$	$-1.2 \cdot 10^{-4}$	$-1.37 \cdot 10^{-4}$	$-1.43 \cdot 10^{-4}$
order		$-7.5 \cdot 10^{-3}$	$-8.55 \cdot 10^{-3}$	$-10.9 \cdot 10^{-3}$	$-10.2 \cdot 10^{-3}$
time to deadline			$1.11 \cdot 10^{-2}$	$1.75 \cdot 10^{-2}$	$1.76 \cdot 10^{-2}$
nominal			$-7.04 \cdot 10^{-5}$	$-8.01 \cdot 10^{-5}$	
log(nominal)					$-4.68 \cdot 10^{-3}$
N obs	358	358	358	358	358
R sq.	0.016	0.022	0.025	0.036	0.034

Table 7: Regression models 16-20.

Dep.: relative offer	Model 21	Model 22	Model 23	Model 24
constant	0.636***	0.666***	0.669***	0.669***
log (buyer_exp)	$0.98 \cdot 10^{-2}$	$1.07 \cdot 10^{-2}$	$1.05 \cdot 10^{-2}$	$1.1 \cdot 10^{-2}$
log (seller_exp)	$1.32 \cdot 10^{-2}$ **	$1.51 \cdot 10^{-2}$ ***	$1.61 \cdot 10^{-2}$ ***	$1.52 \cdot 10^{-2}$ ***
trend		$-1.41 \cdot 10^{-4}$	$-1.39 \cdot 10^{-4}$	$-1.47 \cdot 10^{-4}$
order		$-8.7 \cdot 10^{-3}$	$-1.04 \cdot 10^{-2}$	$-9.66 \cdot 10^{-3}$
time to deadline			$2.11 \cdot 10^{-2}$	$2.13 \cdot 10^{-2}$
nominal			$-8.4 \cdot 10^{-5}$	
log(nominal)				$-5.55 \cdot 10^{-3}$
N obs	358	358	358	358
R sq.	0.023	0.032	0.036	0.034

Table 8: Regression models 21-24.

Dep.: relative offer	NL Model 1	NL Model 2	NL Model 3	NL Model 4
constant	0.63***	0.66***	0.703***	0.67***
log(seller_exp)	$1.32 \cdot 10^{-2}$**	$5.8 \cdot 10^{-3}$		
log(buyer_exp)	$9.83 \cdot 10^{-3}$	$5.1 \cdot 10^{-3}$		
*log(buyer_exp) * log(seller_exp)*		$1.4 \cdot 10^{-3}$		
seller_exp			$9.76 \cdot 10^{-5}$	$1.16 \cdot 10^{-3}$*
buyer_exp			$2.34 \cdot 10^{-5}$	$5.75 \cdot 10^{-5}$*
*seller_exp*buyer_exp*			$2.78 \cdot 10^{-9}$	
seller_exp²				$-2.25 \cdot 10^{-6}$
buyer_exp²				$-6.54 \cdot 10^{-9}$
N obs	358	358	358	358
R sq.	0.023	0.024	0.016	0.026

Table 9: Regression models 25-28.

A.2.2 Decomposition

Coefficient	$\hat{\gamma}_0$	$\hat{\gamma}_1$	$\hat{\gamma}_2$	$\hat{\gamma}_3$	$\hat{\gamma}_4$
Estimate	-43.9533	73.5298	-52.4472	22.4486	-8.4939
Coefficient	$\hat{\delta}_0$	$\hat{\delta}_1$	$\hat{\delta}_2$	$\hat{\delta}_3$	$\hat{\delta}_4$
Estimate	8.39309	-11.4105	4.53588	0.639369	-2.1579

Table 10: Chebyshev polynomial coefficients.

s:	0-10%	10-20%	20-30%	30-40%	40-50%
$\hat{f}^*(s)$:	25.9	7.4	3.7	18.5	44.4

c:	0-10%	10-50%	50-60%	60-70%	70-80%	80-90%	90-100%
$\hat{g}^*(c)$:	14.3	7.1	5.4	17.9	23.2	21.4	10.7

Table 11: Non-parametric estimate of the distributions of sharing rules ($\hat{f}$) and beliefs ($\hat{g}$).

s:	0-10%	10-20%	20-30%	30-40%	40-50%
$\hat{f}^\beta(s)$:	23.5%	12.9%	12.0%	15.2%	36.4%

c:	0-10%	10-50%	50-60%	60-70%	70-80%	80-90%	90-100%
$\hat{g}^\beta(c)$:	0.4%	22.1%	11.9%	14.4%	16.6%	17.9%	16.7%

Table 12: Parametric estimates when f and g are restricted to the class of β distributions. $\alpha^f = 0.230$, $\beta^f = 0.031$; $\alpha^g = 2.470$, $\beta^g = 1.121$, KS distance = 0.066

Iteration	0	1	2	3	4	5	6	7	8	9	10	11	12	13	...	50
$f(b_1)$	20%	14.3%	11.1%	9.1%	14.3%	20.0%	22.2%	26.3%	28.6%	27.3%	26.1%	28.0%	26.9%	25.9%		25.9%
$f(b_2)$	20%	14.3%	11.1%	18.2%	14.3%	13.3%	11.1%	10.5%	9.5%	9.1%	8.7%	8.0%	7.7%	7.4%		7.4%
$f(b_3)$	20%	14.3%	11.1%	9.1%	7.1%	6.7%	5.6%	5.3%	4.8%	4.5%	4.3%	4.0%	3.8%	3.7%		3.7%
$f(b_4)$	20%	28.6%	33.3%	27.3%	28.6%	26.7%	27.8%	26.3%	23.8%	22.7%	21.7%	20.0%	19.2%	18.5%		18.5%
$f(b_5)$	20%	28.6%	33.3%	36.4%	35.7%	33.3%	33.3%	31.6%	33.3%	36.4%	39.1%	40.0%	42.3%	44.4%		44.4%
$g(b_1)$	10%	7.1%	5.6%	4.8%	8.3%	7.7%	10.0%	9.4%	10.5%	11.4%	12.5%	12.5%	13.5%	14.3%		14.3%
$g(b_2)$	10%	7.1%	5.6%	4.8%	4.2%	3.8%	3.3%	3.1%	2.6%	2.3%	2.1%	2.1%	1.9%	1.8%		1.8%
$g(b_3)$	10%	7.1%	5.6%	4.8%	4.2%	3.8%	3.3%	3.1%	2.6%	2.3%	2.1%	2.1%	1.9%	1.8%		1.8%
$g(b_4)$	10%	7.1%	5.6%	4.8%	4.2%	3.8%	3.3%	3.1%	2.6%	2.3%	2.1%	2.1%	1.9%	1.8%		1.8%
$g(b_5)$	10%	7.1%	5.6%	4.8%	4.2%	3.8%	3.3%	3.1%	2.6%	2.3%	2.1%	2.1%	1.9%	1.8%		1.8%
$g(b_6)$	10%	7.1%	5.6%	4.8%	4.2%	3.8%	3.3%	3.1%	5.3%	6.8%	6.3%	6.3%	5.8%	5.4%		5.4%
$g(b_7)$	10%	14.3%	16.7%	19.0%	16.7%	15.4%	16.7%	15.6%	15.8%	15.9%	16.7%	16.7%	17.3%	17.9%		17.9%
$g(b_8)$	10%	14.3%	16.7%	19.0%	20.8%	23.1%	23.3%	25.0%	23.7%	22.7%	22.9%	22.9%	23.1%	23.2%		23.2%
$g(b_9)$	10%	14.3%	16.7%	19.0%	20.8%	23.1%	23.3%	25.0%	23.7%	22.7%	22.9%	22.9%	23.1%	21.4%		21.4%
$g(b_{10})$	10%	14.3%	16.7%	14.3%	12.5%	11.5%	10.0%	9.4%	10.5%	11.4%	10.4%	10.4%	9.6%	10.7%		10.7%
KS-stat	0.1364	0.1658	0.0667	0.0398	0.0327	0.0264	0.0196	0.0171	0.0143	0.0127	0.0106	0.0089	0.0081	0.0077		0.0077

Table 13: The consecutive iterations of the binary search procedure.

Iteration	0	1	2	3	4	5	6	7	8	9	10	11	12	...	24	...	50
$f(b_1)$	16.7%	12.5%	9.1%	8.3%	8.3%	13.3%	16.7%	19.0%	21.7%	22.2%	19.4%	21.2%	22.9%		23.7%		23.7%
$f(b_2)$	16.7%	12.5%	9.1%	8.3%	8.3%	6.7%	5.6%	4.8%	4.3%	3.7%	6.5%	6.1%	5.7%		8.5%		8.5%
$f(b_3)$	16.7%	12.5%	9.1%	8.3%	8.3%	6.7%	5.6%	4.8%	4.3%	3.7%	3.2%	3.0%	2.9%		1.7%		1.7%
$f(b_4)$	16.7%	12.5%	18.2%	16.7%	16.7%	20.0%	22.2%	19.0%	17.4%	18.5%	19.4%	21.2%	22.9%		16.9%		16.9%
$f(b_5)$	16.7%	25.0%	27.3%	33.3%	33.3%	33.3%	33.3%	33.3%	34.8%	33.3%	32.3%	30.3%	28.6%		37.3%		37.3%
$f(b_6)$	16.7%	25.0%	27.3%	25.0%	25.0%	20.0%	16.7%	19.0%	17.4%	18.5%	19.4%	18.2%	17.1%		11.9%		11.9%
$g(b_1)$	10.0%	7.1%	5.6%	4.5%	7.7%	7.4%	9.4%	11.1%	10.8%	12.2%	13.3%	13.0%	13.7%		15.6%		15.6%
$g(b_2)$	10.0%	7.1%	5.6%	9.1%	7.7%	7.4%	6.3%	5.6%	5.4%	4.9%	4.4%	4.3%	3.9%		2.6%		2.6%
$g(b_3)$	10.0%	7.1%	5.6%	4.5%	3.8%	3.7%	3.1%	2.8%	2.7%	2.4%	2.2%	2.2%	2.0%		1.3%		1.3%
$g(b_4)$	10.0%	7.1%	5.6%	4.5%	3.8%	3.7%	3.1%	2.8%	2.7%	2.4%	2.2%	2.2%	2.0%		1.3%		1.3%
$g(b_5)$	10.0%	7.1%	5.6%	4.5%	3.8%	3.7%	3.1%	2.8%	2.7%	2.4%	2.2%	2.2%	2.0%		2.6%		2.6%
$g(b_6)$	10.0%	7.1%	11.1%	9.1%	7.7%	7.4%	6.3%	5.6%	5.4%	4.9%	4.4%	4.3%	3.9%		2.6%		2.6%
$g(b_7)$	10.0%	14.3%	16.7%	18.2%	19.2%	18.5%	18.8%	19.4%	18.9%	19.5%	20.0%	19.6%	19.6%		19.5%		19.5%
$g(b_8)$	10.0%	14.3%	16.7%	18.2%	19.2%	18.5%	18.8%	19.4%	21.6%	22.0%	22.2%	23.9%	23.5%		27.3%		27.3%
$g(b_9)$	10.0%	14.3%	16.7%	18.2%	19.2%	22.2%	21.9%	22.2%	21.6%	22.0%	22.2%	21.7%	21.6%		19.5%		19.5%
$g(b_{10})$	10.0%	14.3%	11.1%	9.1%	7.7%	7.4%	9.4%	8.3%	8.1%	7.3%	6.7%	6.5%	7.8%		7.8%		7.8%
KS		0.1250	0.0550	0.0411	0.0361	0.0322	0.0271	0.0231	0.0190	0.0171	0.0161	0.0151	0.0129		0.0066		0.0066

Table 14: The consecutive iterations of the binary search procedure, with restriction 0.6.

A.3 Figures

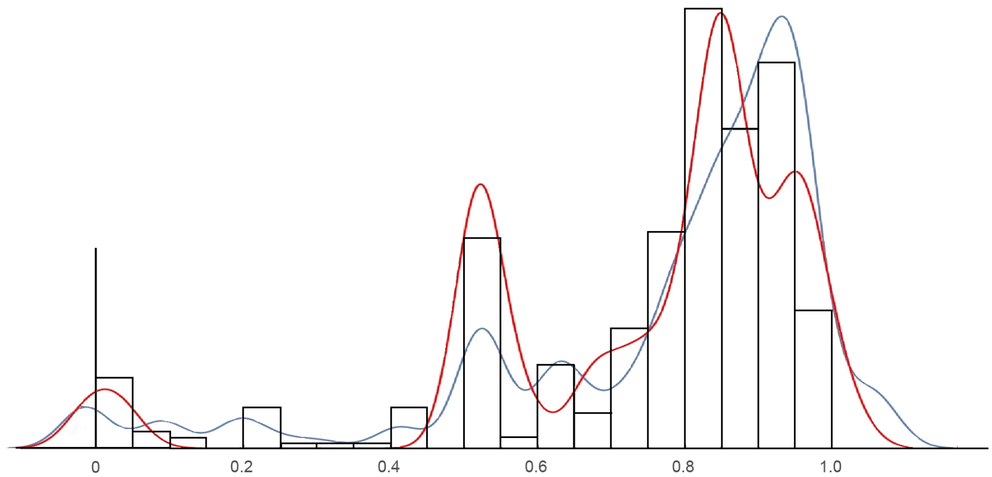

Figure 5: Smooth kernel histograms for the offers made by the East (red) and the West German participants (blue).

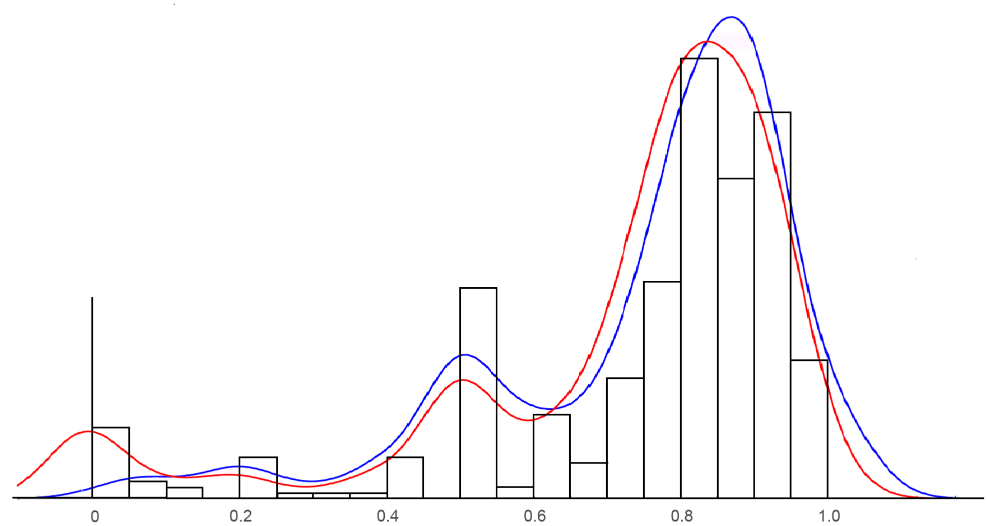

Figure 6: The histogram of the relative offers (grey bars) and smooth kernel histograms for the offers made first (blue) and all the higher-order offers (red).

A.4 Notes

A.4.1 Note: Non-parametric Decomposition

As an alternative way to solve for f and g, we discretize the support of distributions in (2) and we search for finite solution approximations using a non-parametric approach. In an iterative procedure with the initial state where both f and g are uniform, we gradually increase precision until the solution cannot be improved. The estimates of f and g are chosen to minimize the KS distance as a goodness-of-fit criterion, adapted to the discrete case:

$$d_{KS}\left(\hat{H},H\right) \equiv \sup_{\{b_k\}_{k=1,..,10}} \left|\hat{H}\left(b_k\right) - H\left(b_k\right)\right|, \tag{5}$$

where both $\hat{H}$ and H are defined on a set of bins $(b_1, b_2, .., b_{10}) = (0.1, 0.2, .., 1)$. In the discrete problem, we look for $\hat{f}$ and $\hat{g}$ that minimize (5), subject to $\hat{f}\left(b_k\right) = 0$ for $k = 6, 7, .., 10$, and $\sum_k \hat{f}\left(b_k\right) = \sum_k \hat{g}\left(b_k\right) = 1$. (Recall that $\hat{f}$ and $\hat{g}$ define $\hat{H}$ according to (4)).

We estimate discretized versions of f and g simultaneously at each iteration. As a starting point of recurrence, we consider uniform f and g that correspond to the maximal entropy in both c and s (See Table 13 in the Appendix). That is, we assign equal mass to each of the 5 bins of f and 10 bins of g at the first iteration. The candidate solution can be represented as a vector of ones: $v_1 = (1, 1, .., 1)$; it corresponds to the distribution of probability mass across bins before normalization. v_1 is the unique candidate solution at the first iteration, thus we set $v_1^* = v_1$ At the second iteration, we consider all elements $e_{2,k}$ of the set $\{0,1\}^{15}$ and the respective $v_{2,k} = v_1 + e_{2,k}$. This gives $2^{15} = 32768$ candidate solutions (f, g) and we choose the one that minimizes the KS distance (after the appropriate normalization) – v_2^*. Generally, in iteration t, we go through all possible constellations of adding one unit or not changing the mass in the bins of distributions defined by v_{t-1}^* selected at iteration $t-1$. Among the pairs of f and gn we select the one that minimizes the KS distance and take it to iteration $t+1$. The process is repeated until the result at a subsequent iteration stays unchanged. Note that as the depth of the binary search tree increases, the estimates become increasingly precise. The results of the non-parametric approach are presented in Table 11.

Define $\hat{H}^*$ as in (4) by plugging in $\hat{f} = \hat{f}^*$ and $\hat{g} = \hat{g}^*$ (Table 11). $\hat{H}^*$ is an extremely good fit for H: the Kolmogorov-Smirnov distance is 7.7×10^{-3}, meaning that the maximal divergence between $\hat{H}^*$ and H across bins is less than 1 percentage point. The corresponding bootstrap test does not distinguish between H and $\hat{H}^*$ at the conventional significance levels,[41] implying that H and $\hat{H}^*$ can be regarded as equivalent.

As a further robustness check, we relax the restriction on the support of f allowing six bins; similar results are obtained (see Table 14).

A.4.2 Note: Parametric Decomposition with Beta Distributions.

As a further robustness check, we look for estimates $\hat{f}^\beta$ and $\hat{g}^\beta$ within the class of the beta distributions. The beta class is chosen due to its support on [0, 1], small number of parameters and flexibility, as Beta distributions can have one or two modes.[42] Since each of the distributions f and g is pinned down by two parameters, the problem to find the suitable $\hat{H}(.)$ in (3) is now (only) four-dimensional. We estimate the parameters by random grid search. More precisely, we fix a grid size for each of four parameters and randomly choose the grid position. For every intersection of the grid (a combination of four parameter values), we compute the discretized versions of f and g, estimate the integral (4), then we calculate the KS distance and reiterate to find the best-performing combination of parameters. By performing the procedure multiple times, we refine the search and narrow down the parameters ranges. A random grid search permits us to trace out local minima and find the global solution. We report the parametric estimates of distributions $\hat{f}^\beta$ and $\hat{g}^\beta$ Table 12. The parametric Beta estimates $\hat{f}^\beta$

[41] Bootstrap critical values corresponding to 358 observations are 0.049 for significance at 10%, 0.057 for significance at 5%, and 0.071 for significance at 1%.

[42] Excluding the uniform distribution.

and $\hat{g}^{\beta}$ are dominated by both the parametric and the non-parametric solutions reported in the main text of the paper (Kolmogorov-Smirnov distance $= 0.066$). The distributions have a shape similar to the non-parametric estimates $\hat{f}^{*}$ and $\hat{g}^{*}$, with the exception of the lowest bin in the distribution of cost estimates $\hat{g}$.

References

ALESINA, A. AND N. F. SCHUENDELN (2005): "Good bye Lenin (or not?): The effect of Communism on people's preferences," Tech. rep., National Bureau of Economic Research.

ALGER, I. AND J. W. WEIBULL (2013): "Homo moralis – preference evolution under incomplete information and assortative matching," *Econometrica*, 81, 2269–2302.

BEARDEN, J. (2001): "Ultimatum Bargaining Experiments: The State of the Art," Working paper, INSEAD – Decision Sciences.

BOEHM, C., H. B. BARCLAY, R. K. DENTAN, M.-C. DUPRE, J. D. HILL, S. KENT, B. M. KNAUFT, K. F. OTTERBEIN, AND S. RAYNER (1993): "Egalitarian behavior and reverse dominance hierarchy [and comments and reply]," *Current Anthropology*, 34, 227–254.

BOLTON, G. E. AND A. OCKENFELS (2014): "Does laboratory trading mirror behavior in real world markets? Fair bargaining and competitive bidding on eBay," *Journal of Economic Behavior & Organization*, 97, 143–154.

BROSIG-KOCH, J., C. HELBACH, A. OCKENFELS, AND J. WEIMANN (2011): "Still different after all these years: Solidarity behavior in East and West Germany," *Journal of Public Economics*, 95, 1373–1376.

CABRAL, L. AND A. HORTASCU (2010): "The dynamics of seller reputation: Evidence from eBay," *The Journal of Industrial Economics*, 58, 54–78.

CAMERER, C. (2003): *Behavioral game theory: Experiments in strategic interaction*, Princeton, NJ: Cambridge University Press.

CAMERON, L. A. (1999): "Raising the stakes in the ultimatum game: Experimental evidence from Indonesia," *Economic Inquiry*, 37, 47–59.

CARD, D., S. DELLAVIGNA, AND U. MALMENDIER (2011): "The Role of Theory in Field Experiments," *Journal of Economic Perspectives*, 25, 39–62.

CHARNESS, G. AND M. RABIN (2002): "Understanding Social Preferences with Simple Tests," *The Quarterly Journal of Economics*, 117, 817–869.

CHERRY, T. L., P. FRYKBLOM, AND J. F. SHOGREN (2002): "Hardnose the dictator," *American Economic Review*, 1218–1221.

DUFWENBERG, M. AND G. KIRCHSTEIGER (2004): "A theory of sequential reciprocity," *Games and Economic Behavior*, 47, 268 – 298.

ELY, J. C. AND T. HOSSAIN (2009): "Sniping and Squatting in Auction Markets," *American Economic Journal: Microeconomics*, 1, 68–94.

FALK, A. AND U. FISCHBACHER (2006): "A theory of reciprocity," *Games and Economic Behavior*, 54, 293 – 315.

FEHR, E. AND K. M. SCHMIDT (1999): "A Theory of Fairness, Competition, and Cooperation," *The Quarterly Journal of Economics*, 114, 817–868.

FEHR-DUDA, H., A. BRUHIN, T. EPPER, AND R. SCHUBERT (2010): "Rationality on the rise: Why relative risk aversion increases with stake size," *Journal of Risk and Uncertainty*, 40, 147–180.

FISCHBACHER, U., C. FONG, AND E. FEHR (2009): "Fairness, Error and the Power of Competition," *Journal of Economic Behaviour and Organization*, 72, 527–545.

FUKUYAMA, F. (1995): *Trust: The social virtues and the creation of prosperity*, Free Press New York.

GALIZZI, M. M. AND D. NAVARRO MARTINEZ (2015): "On the external validity of social-preference games: A systematic lab-field study," .

GIBBS, J., K. L. KRAEMER, AND J. DEDRICK (2003): "Environment and policy factors shaping global e-commerce diffusion: A cross-country comparison," *The information society*, 19, 5–18.

GIGERENZER, G. (2007): *Gut feelings: The intelligence of the unconscious*, Penguin.

GÜTH, W. AND M. KOCHER (2013): "More than Thirty Years of Ultimatum Bargaining Experiments: Motives, Variations, and a Survey of the Recent Literature," Working paper, Max Planck Instute of Economics and CESifo.

GÜTH, W., C. SCHMIDT, AND M. SUTTER (2007): "Bargaining outside the lab – a newspaper experiment of a three-person ultimatum game," *The Economic Journal*, 117, 449–469.

GÜTH, W., R. SCHMITTBERGER, AND B. SCHWARZE (1982): "An experimental analysis of ultimatum bargaining," *Journal of Economic Behaviour and Organization*, 3, 367–388.

HENRICH, J., R. MCELREATH, A. BARR, J. ENSMINGER, C. BARRETT, A. BOLYANATZ, J. C. CARDENAS, M. GURVEN, E. GWAKO, N. HENRICH, ET AL. (2006): "Costly punishment across human societies," *Science*, 312, 1767–1770.

HOFFMAN, E., K. MCCABE, AND V. SMITH (1996): "On expectations and the monetary stakes in ultimatum games," *International Journal of Game Theory*, 25, 289–301.

HOLT, C. AND S. LAURY (2002): "Risk Aversion and Incentive Effects," *American Economic Review*, 92, 1644–1655.

JOHN, K. AND S. L. THOMSEN (2013): "Environment and other-regarding preferences," Tech. rep., Discussion Paper 4, Niedersächsisches Institut für Wirtschaftsforschung, Hannover.

KALAI, E. AND M. SMORODINSKY (1975): "Other solutions to Nash's bargaining problem," *Econometrica: Journal of the Econometric Society*, 513–518.

LEVITT, S. D. AND J. A. LIST (2009): "Field experiments in economics: The past, the present, and the future," *European Economic Review*, 53, 1 – 18.

LIST, J. A. (2006): "The Behavioralist Meets the Market: Measuring Social Preferences and Reputation Effects in Actual Transactions," *Journal of Political Economy*, 114.

——— (2011): "Why Economists Should Conduct Field Experiments and 14 Tips for Pulling One Off," *Journal of Economic Perspectives*, 25, 3–16.

MUNIER, B. AND C. ZAHARIA (2002): "High stakes and acceptance behavior in ultimatum bargaining," *Theory and Decision*, 53, 187–207.

NEWELL, A., H. A. SIMON, ET AL. (1972): *Human problem solving*, vol. 104.

NOSKO, C. AND S. TADELIS (2015): "The limits of reputation in platform markets: An empirical analysis and field experiment," Tech. rep., National Bureau of Economic Research.

OCKENFELS, A. AND G. E. BOLTON (2000): "ERC: A Theory of Equity, Reciprocity, and Competition," *American Economic Review*, 90, 166–193.

OCKENFELS, A. AND J. WEIMANN (1999): "Types and patterns: an experimental East-West-German comparison of cooperation and solidarity," *Journal of Public Economics*, 71, 275–287.

OOSTERBEEK, H., R. SLOOF, AND G. VAN DE KUILEN (2004): "Cultural Differences in Ultimatum Game Experiments: Evidence from a Meta-Analysis," *Experimental Economics*, 7, 171–188.

RABIN, M. (1993): "Incorporating Fairness into Game Theory and Economics," *American Economic Review*, 83, 1281–1302.

RESNICK, P., R. ZECKHAUSER, J. SWANSON, AND K. LOCKWOOD (2006): "The value of reputation on eBay: A controlled experiment," *Experimental Economics*, 9, 79–101.

ROTH, A., V. PRASNIKAR, M. OKUNO-FUJIWARA, AND S. ZAMIR (1991): "Bargaining and Market Behaviour in Jerusalem, Ljubljana, Pittsburgh and Tokyo," *American Economic Review*, 81, 1068–1095.

ROTH, A. E. AND I. EREV (1995): "Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term," *Games and Economic Behavior*, 8, 164 – 212.

Roth, A. E. and A. Ockenfels (2002): "Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet," *American Economic Review*, 92, 1093–1103.

Sadovnichy, V. (1986): *Theory of Operators (in Russian)*, Moscow State University.

Slonim, R. and A. E. Roth (1998): "Learning in high stakes ultimatum games: An experiment in the Slovak Republic," *Econometrica*, 569–596.

Tisserand, J.-C. (2014): "Ultimatum game: A meta-analysis of the past three decades of experimental research," Working paper, International Institute of Social and Economic Sciences.

Tversky, A. and D. Kahneman (1974): "Judgment under uncertainty: Heuristics and biases," *science*, 185, 1124–1131.

van Damme, E., K. G. Binmore, A. E. Roth, L. Samuelson, E. Winter, G. E. Bolton, A. Ockenfels, M. Dufwenberg, G. Kirchsteiger, U. Gneezy, M. G. Kocher, M. Sutter, A. G. Sanfey, H. Kliemt, R. Selten, R. Nagel, and O. H. Azar (2014): "How Werner Güth's ultimatum game shaped our understanding of social behavior," *Journal of Economic Behavior and Organization*, 108, 292–318.