CESifo WORKING PAPERS

9280
2021

September 2021

Ultimatum Game Behavior in a Social-Preferences Vacuum Chamber

Volker Benndorf, Thomas Große Brinkhaus, Ferdinand A. von Siemens

CESifo

Impressum:

CESifo Working Papers
ISSN 2364-1428 (electronic version)
Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH
The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute
Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de
Editor: Clemens Fuest
https://www.cesifo.org/en/wp
An electronic version of the paper may be downloaded
· from the SSRN website: www.SSRN.com
· from the RePEc website: www.RePEc.org
· from the CESifo website: https://www.cesifo.org/en/wp

CESifo *Working Paper No. 9280*

Ultimatum Game Behavior in a Social-Preferences Vacuum Chamber

Abstract

We study strategic interaction in an experimental social-preferences vacuum chamber. We mute social preferences by letting participants knowingly interact with computers. Our new design allows for indirect strategic interaction: there are several waves in which computer players inherit the behavior of human players from the previous wave. We apply our method to investigate trembling-hand perfection in a normal-form version of the ultimatum game. We find that behavior remains far off from a trembling-hand perfect equilibrium under selfish preferences even towards the end of our experiment. The likely reasons for our findings are strategic uncertainty and incomplete learning.

JEL-Codes: C920, C720, D910.

Keywords: social preferences, induced-value theory, learning, ultimatum game, strategic interaction.

*Volker Benndorf**
Goethe University Frankfurt
Frankfurt am Main / Germany
benndorf@econ.uni-frankfurt.de

Thomas Große Brinkhaus
Goethe University Frankfurt
Frankfurt am Main / Germany
thomas.gb@t-online.de

Ferdinand A. von Siemens
Goethe University Frankfurt
Frankfurt am Main / Germany
vonsiemens@econ.uni-frankfurt.de

*corresponding author.

This research was funded by Goethe University Frankfurt. Declarations of interest: none.

1 Introduction

Economic experiments offer incomparable control over the strategic situation and are thus the best method to empirically study concepts from game theory. However, full control over the players' payoffs – how experimental participants evaluate the outcomes of the game – eludes experimenters if their participants have social preferences. As Smith (1976) himself points out, induced value theory does not apply if participants care for the wellbeing of other participants because their evaluations are then not necessarily strictly increasing in their own monetary payoffs. More recently, Weibull (2004) concludes that this lack of full control over what game participants actually play makes testing game theory challenging, even with laboratory experiments.

In this paper, we try to take back some control by creating an experimental social-preferences vacuum chamber. We aim to create a decision environment with reduced social preferences while still allowing for strategic interaction between the participants. We are much like physicists who use vacuum chambers to study gravity in the absence of air friction. And as these physicists do not question the existence of air, we do not question the importance of social preferences for strategic interaction. Instead, we want to study strategic situations "less confounded" by social preferences to understand better the fundamental forces driving strategic behavior.

Our social-preferences vacuum chamber is based on the following design principle. Applying a well-established method from the literature, we reduce social preferences by having players knowingly interact with computers. Our design innovation is how we program the computers to allow indirect strategic interaction between the human players. Our experiment consists of several waves, where computer players always inherit the behavior of human players from the previous wave.

We use our social-preferences vacuum chamber to study the empirical relevance of trembling-hand perfection in a normal-form version of the ultimatum game (Güth et al., 1990). Because the game is played in the normal form, subgame perfection cannot be applied to reduce the number of Nash equilibria. However, alternative refinements such as trembling-hand perfection or weak dominance generate the same predictions. Put differently, corresponding to the two subgame-perfect equilibria of the extensive-form game, there are two trembling-hand perfect equilibria of the normal-form game. In the first, responders demand one euro, and proposers offer one euro. In the second, responders demand zero euros, and proposers offer zero euros.

When humans interact with humans, we expect responders to reject low offers and proposers to make generous offers. Unfortunately, such behavior does not reveal much about the empirical relevance of trembling-hand perfection because it is consistent with a trembling-hand perfect equilibrium if participants have social preferences. However, if trembling-hand perfection is empirically relevant, we should observe the following behavior in our social-preferences vacuum chamber. Human responders should be more willing to accept low offers made by computer proposers. Human proposers should figure this out and become more aggressive when playing against the more accepting computer responders. In the following, every wave of more aggressive proposers should make the next wave of responders more lenient. And every wave of more lenient responders induces the next wave of proposers to become more aggressive. Behavior should converge to a trembling-hand perfect equilibrium with selfish preferences.

Consistent with the existing literature, we find that human responders are much more willing to accept low offers when playing against computers than human proposers. We conclude that our social-preferences vacuum chamber successfully reduces social preferences. However, the proposers playing against computers hardly adjust their behavior and continue to make cautious, generous offers. Even after nine waves, proposer behavior continues to resemble the behavior in the classic ultimatum experiments. We conclude that trembling-hand perfection is not empirically relevant in our setup.

A more in-depth analysis reveals that incomplete learning prevents the players from arriving at a trembling-hand perfect equilibrium. We find some learning within each wave if humans interact with computer players: responders become more lenient and proposers become more aggressive. But learning is slow, in particular for the proposers. This finding is consistent with Gale et al. (1995), Roth and Erev (1995), and Levine (2012), who argue that such strategic uncertainty and slow learning could explain the experimental evidence even if participants have purely selfish preferences. Finally, we observe that proposers and responders do not adjust their behavior within the wave if they interact with other humans. Therefore, learning depends on whether participants interact with humans or computers. This finding emphasizes the usefulness of our social-preferences vacuum chamber for studying strategic interaction while reducing social preferences' confounding effects.

2 Related literature

In their seminal contributions, Houser and Kurzban (2002), Ferraro et al. (2003), and Ferraro and Vossler (2010) use computer players to analyze the effects of social preferences in public goods games. They find that confusion contributes substantially to the amount of cooperation, but social preferences also matter. The programming of the computer players remains the same throughout the experiment, so that indirect strategic interaction as in our experiment is ruled out. Sanfey et al. (2003) and van't Wout et al. (2006) focus on responder behavior in the ultimatum game. Using functional magnetic resonance imaging and skin conductance measurements, they find that responders are less emotionally involved and more accepting when playing against computers. Johnson et al. (2002) let human players knowingly interact with computer players in a three-stage bargaining game. They tell the participants that the computer opponents follow the subgame-perfect equilibrium strategy with selfish preferences.[1]

Our social-preferences vacuum chamber also uses the virtual player approach to reduce social preferences. However, our paper is the first to allow indirect strategic interaction between the human players across different waves. The computer players in all previous studies follow pre-determined algorithms, which precludes any interactive learning of human proposers and human responders. However, we believe that such interactive learning is necessary for trembling-hand perfection to become empirically relevant over time.

Our findings on learning contribute to the literature on learning in the ultimatum game. In experiments with human-human interaction – where social preferences may play a role – List and Cherry (2000) and Cooper et al. (2003) show that responders learn to accept lower offers. Cooper and Dutcher (2011) find that responders learn to accept relatively high offers but to reject very low offers. Our experiment reveals that responders learn to accept low offers, but only if they interact with computers. This result confirms that learning and social preferences interact.

[1] Harrison and McCabe (1996) and Winter and Zamir (2005) use computer players to control expectation formation or study evolutionary adjustment dynamics in the ultimatum game. Their participants interact with both human and computer players. However, the human players either do not know that some of their opponents are computer-simulated, or they receive only limited information on this.

3 Experimental Design

3.1 Strategic Situation

We consider an ultimatum game in the normal form. A proposer and a responder bargain over the division of ten euros. The proposer makes the responder an offer between zero and ten euros, in steps of one euro. The responder simultaneously chooses a demand between zero and ten euros in steps of one euro. The offer of the proposer is accepted, and the proposed allocation is implemented, if it is weakly larger than the demand. Otherwise, the offer is rejected, and the responder and proposer receive zero euros. Note that the responder determines his demand before he learns the actual offer of the responder.

The above ultimatum game has eleven Nash equilibria in pure strategies if proposers and responders maximize their expected monetary rewards. In each equilibrium, responders set a specific demand, and proposers exactly match this demand. Because the game is played in the normal form, subgame perfection cannot be applied to reduce the number of Nash equilibria. But corresponding to the two subgame-perfect equilibria of the extensive-form game, there are two trembling-hand perfect equilibria of the game. In the first, responders demand one euro, and proposers offer one euro. In the second, responders demand zero euros, and proposers offer zero euros.

3.2 Interaction in the Social-Preferences Vacuum Chamber

We dampen the social preferences of players by letting them knowingly interact with computer players. This virtual player approach is a well-established method to reduce social preferences, and we will see that it is also effective in the present setup. Our design innovation is that we invite several waves of participants. These participants play against computers, which are programmed to behave like the human players in their role from the previous wave. Behavior can settle down over the waves. Next, we describe our wave design in more detail.

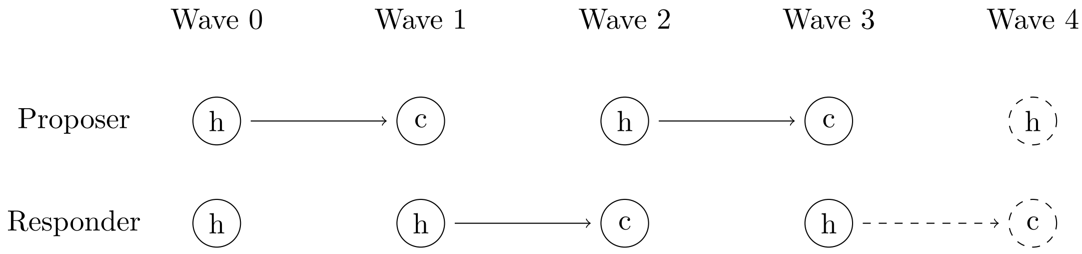

Figure 1: Structure of the Experimental Design.

Note: Structure of the experimental design. The experiment contains a number of waves. The circles represent players, proposers and responders, and h and c indicate whether the players are human or computer-simulated. The arrows show that human players determine the behavior of computer players with the same role in the next wave.

Figure 1 illustrates our experimental design. The experiment has several waves, each with new participants. In Wave 0, human proposers and human responders interact over ten periods. In Wave 1, human responders interact with computer proposers. Computer proposers are programmed to behave exactly like the human proposers in the previous wave. Human responders know they interact with computer proposers, which should dampen their social preferences. In the next wave, human proposers interact with computer responders, who inherit their behavior from the previous wave's human responders. We repeat this procedure until we run out of money, until we deplete the subject pool, or hopefully until the behavior becomes stable in the sense that we no longer expect behavior to change substantially over further waves.

Once the behavior has become stable, we observe human proposers, who interact with computer responders, who behave like human responders, who interact with computer proposers, who behave like the human proposers themselves. We equally see human responders, who interact with computer proposers, who act like human proposers, who interact with computer responders, who behave like the human responders themselves. We consider this behavior to be an equilibrium in the social-preferences vacuum chamber. If trembling-hand perfection has empirical relevance in our setup, then this equilibrium behavior should correspond to a trembling-hand perfect equilibrium under selfish preferences.

3.3 Implementation

We implement our experimental design as follows. Upon entering the laboratory, participants receive written instructions on all parts of the experiment. We first measure participants' risk preferences as proposed by Gneezy and Potters (1997). Participants receive 2 euros. They can invest any amount in steps of 20 euro cents into a risky project. This risky project yields 2.5 times the invested amount with a probability of 50% and otherwise zero. Participants learn the outcome of their investment decision at the end of the experiment.

Participants then play an ultimatum game as described before. Participants are randomly assigned the role of proposer or responder and keep their role throughout the experiment. They play the ultimatum game over ten rounds with a perfect stranger matching. We use the turn-pike matching protocol by Cooper et al. (1996), which rules out that behavior can directly or indirectly affect the behavior of opponents in any future round. The procedure avoids contagion effects discussed in Kandori (1992) and makes sure that players face a sequence of one-shot interactions. The procedure also limits the number of interactions to ten given the size of the laboratory. We do not explain to participants the complicated matching algorithm but inform them that each player meets another player at most once and that players cannot affect in any way their future opponents' behavior.

The instructions include explicit information on whether participants interact with other human or computer-simulated participants. In Wave 0, we inform participants that other participants are assigned the other role in the experiment, that they directly interact with these other participants, and that their behavior affects their own and other human participants' payoffs in the experiment today. In the next waves, we inform people that the players in the other role are simulated by the computer. The precise wording is as follows.

> "A computer simulates the other role in the experiment. Therefore, you do not interact with another human participant in the experiment, but a computer simulates your opponents. This means that your behavior does influence only your payment, but does not directly affect the payment of other human participants. The behavior of the computer-simulated opponents corresponds to the behavior of human participants in earlier sessions of the experiment."

The computer-simulated participants follow the human participants' behavior in the same role from the previous wave, round for round. We provide what might be called natural feedback:

at the end of each period, responders learn the offer by their proposers, and proposers learn whether their offers were rejected or accepted by their responders. The proposers do not learn the strategy, that is, the demand of their responders.

Participants must answer control questions to ensure understanding before the experiment starts. Our experimental manipulation only works if participants understand that they are in a social-preferences vacuum chamber. Therefore, we explicitly ask participants whether they can affect other human participants' payoffs in the same session. We incentivize the answer to this question; participants earn one euro for a correct answer. An error message pops up if participants give the wrong answer, which repeats the relevant information. Thus, we are reasonably confident that participants understand – at least after receiving feedback and before the experiment starts– whether they can affect other participants' monetary payoffs in that session. We also include a short questionnaire that asks participants for their gender, field of study, and age.

3.4 Behavioral predictions

Participants in any experiment might need time to get familiar with the setup and the decision environment. They might also have to experience the material consequences of their decisions to trigger learning and understanding. We are interested in the behavior of experienced participants who have had some learning opportunities and understand the consequences of their actions. Thus, we focus on behavior in the last period in the following.

In Wave 0, the social-preferences vacuum chamber is inactive, and we expect to observe the usual ultimatum game behavior. However, we reduce social preferences in Wave 1. The existing literature suggests that human responders are more willing to accept unfair offers made by computers than human proposers. The reasons might be that it makes no sense to punish computer proposers, to voice anger to a virtual opponent, or to avoid payoff inequity with the computer. Therefore, we expect last-period demands in Wave 1 to be lower than in Wave 0.

In Wave 2, the proposers know that they interact with computers. Relieved of fairness norms and unaffected by social preferences, these proposers should maximize their monetary payoffs. Although they do not know what demands their computer opponents make, they can use the initial periods to experiment with their offers. Learning should lead them to match lower demands with lower offers. Therefore, we expect responders to make lower last-period offers in Wave 2 than in Wave 0.

After Wave 0, the first period of all waves is a priori identical for the participants: they receive the same instructions with identical information about their computer opponents. However, learning and thus the adjustments of behavior should depend on the choice frequencies with which we endow the computer players. Proposers should set lower offers when confronted with lower demands. Responders facing lower offers more frequently experience that high demands can lead to zero payoffs, and they should further lower their demands. Through this interplay, we expect offers and demands to fall with each wave so that last-period behavior approaches the trembling-hand-perfect Nash equilibrium for the game without social preferences.

3.5 Procedures

We programmed the experiment in z-tree by Fischbacher (2007) and conducted all sessions in the Frankfurt Laboratory for Experimental Economic Research (FLEX) at Goethe University Frankfurt. We recruited participants from a standard student subject pool via ORSEE by Greiner (2015). We had 487 participants in 28 sessions. Each participant participated only once and only in one of nine waves. The sessions lasted for around 25 minutes, and participants earned, on average, 12.29 euros, including 5.00 euros show-up fee. We include a translation of the originally German instructions in the appendix.

4 Empirical Results

Figure 2 shows the average last-period demands and offers over the waves. We hypothesized that responders set lower last-period offers in Wave 1 than in Wave 0. We indeed observe a sharp drop in average demands as responders interact with computers in Wave 1 rather than with human proposers in Wave 0. The average last-period demands are 3.45 in Wave 0 and 1.87 in Wave 1. This change in demands is consistent with our expectations. We also hypothesized that proposers in Wave 2 set lower last-period offers in Wave 2 than in Wave 0. We do not observe this: average last-period offers drop only from 4.34 in Wave 0 to 4.13 in Wave 2. Proposers do not substantially change their behavior when interacting with computers rather than human responders.

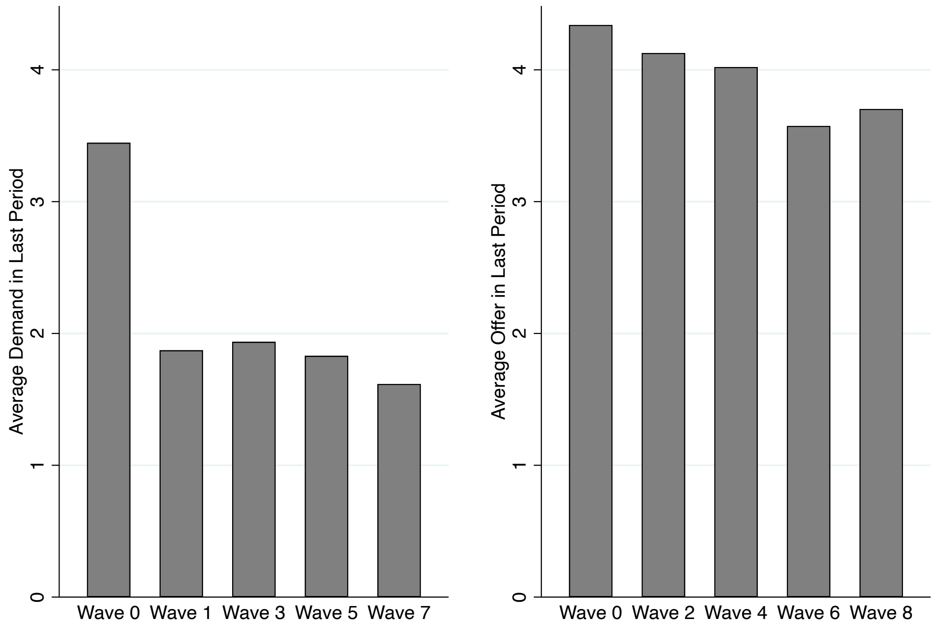

Figure 2: Average Demands and Offers in the Last Period.

Regression analyses confirm these first impressions. Concerning the responder behavior, we find that demands are significantly lower in Wave 1 than in Wave 0 (p-value smaller than 0.001). Note that proposer behavior is, by definition, the same in Waves 0 and 1. The drop in demands shows that our social-preferences vacuum chamber is functional. Concerning the proposer behavior, we find that the differences in offers between Wave 0 and Wave 2 are not statistically significant (p-value of 0.40). Conclusions are the same when including gender, age, and risk-aversion as control variables. See Table 1 in the appendix for more details.

We hypothesized that every wave of more aggressive proposers should make the next wave of responders more lenient. And every wave of more lenient responders should induce the next wave of proposers to become more aggressive. Behavior should converge to a trembling-hand perfect equilibrium with selfish preferences. However, proposers in Wave 2 are not substantially more aggressive than in Wave 0, which suggests that the convergence process might be very slow. In the following, we study how demands and offers evolve over the waves in our social preferences vacuum chamber.

Figure 2 indeed shows that demands and offers do not decline enormously over the waves in our social-preferences vacuum chamber. Concerning the responder behavior, the figure suggests that demands might be particularly low in the last two waves, which might constitute a time trend. But the overall decline in demands is small after Wave 0. Regression analyses reported in the appendix shows that demands do not significantly decrease after the initial wave. We find that last-period demands do not significantly decrease over time in our social-preferences vacuum chamber, also when allowing for a non-linear trend (p-values weakly larger than 0.38). Concerning the proposer behavior, Figure 2 might hint at a slightly more pronounced downward trend of the last-period offers. Our regression analysis shows that the time trend in offers is significant in the linear specification (p-value of 0.06). However, the decline in offers is not statistically significant once allowing for a quadratic time trend. Our conclusions are the same when including gender, age, and risk-aversion as control variables. Contrary to our expectations, we observe that demands and offers do not decline substantially in our social-preferences vacuum chamber. See Table 2 in the appendix for more details.

Because demands and offers do not fall very sharply over the waves, the behavior should remain far off from any trembling-hand perfect equilibrium with selfish preferences. Figure 3 reports last-period behavior in the last two waves of the experiment. Whereas 64% of the responders demand zero or one, only 13% of the proposers make offers consistent with trembling-hand perfection.

To provide a more statistical perspective, we code a dummy variable that equals one if and only if participants make final demands or offers that are not consistent with trembling-hand perfection – if final demands or offers are larger than one. A deviation from trembling-hand perfection could be considered an error. We can now use two-sided binomial tests to study what error probabilities are consistent with our data at a 10% significance level. For responders, we reject error probabilities that are smaller than 0.25 or larger than 0.48. For proposers, the same tests reject error probabilities that are smaller than 0.79 or larger than 0.92. These results confirm our impression that behavior, even towards the end of the experiment, is still far off any trembling-hand equilibrium with selfish preferences. Trembling-hand perfection has no empirical relevance in our setup.

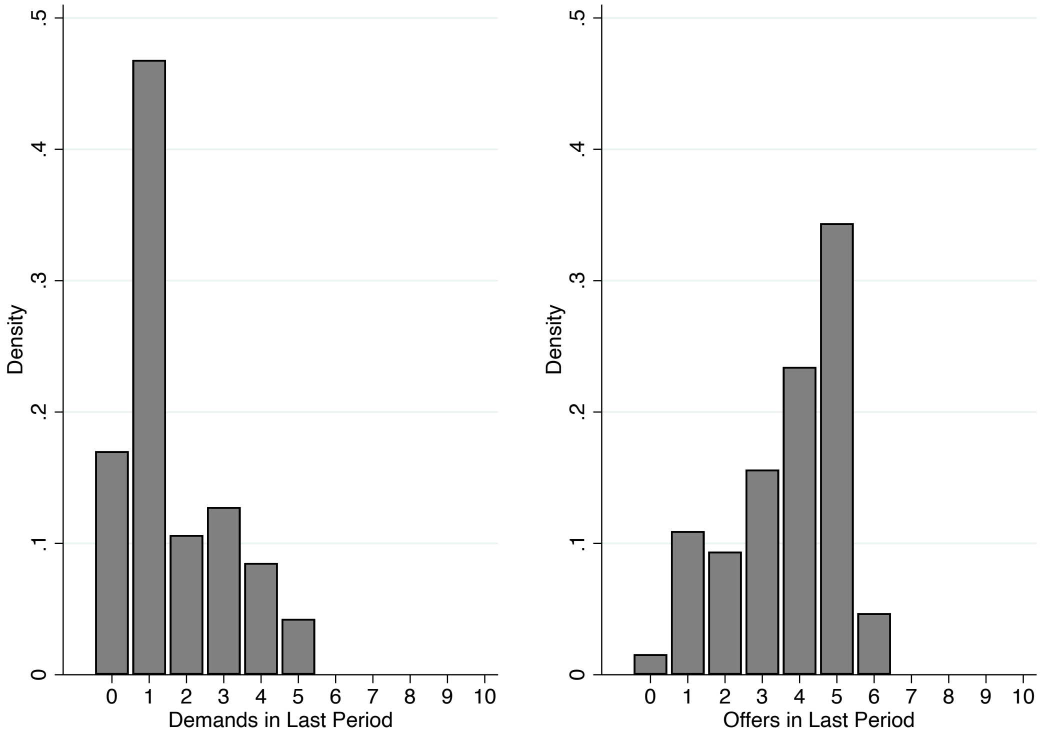

Figure 3: Demands and Offers in the Final Waves.

Summarizing our results so far, we find that the responders set lower demand as they interact with the computers rather than with the human proposers. However, the proposers do not become more aggressive as they interact with the more lenient computers than with the human responders. A likely explanation for these findings is that proposers have a hard time figuring out the computer responders' behavior. In the remaining results section, we explore the possible absence of such learning more carefully. We measure learning via the changes in offers and demands within the ten periods of a wave.

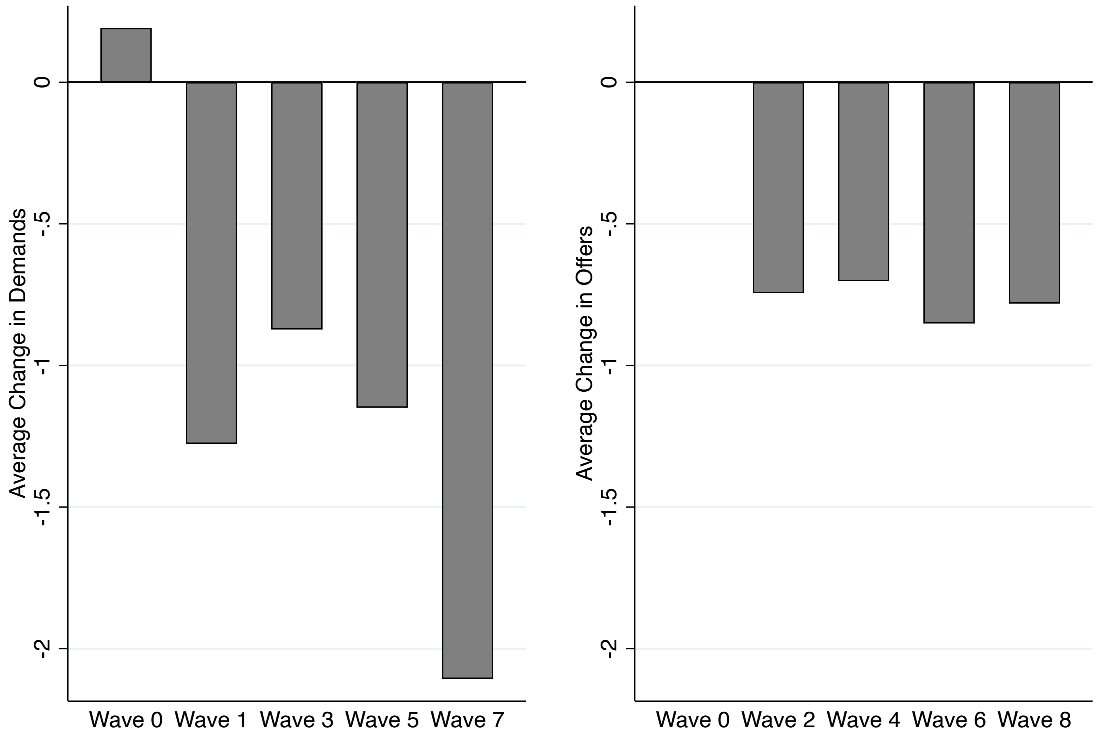

Figure 4: Average changes in demands and offers within waves.

Note: The changes in demands and offers are defined as the choice in the last period minus the choice in first period. The average change of proposers' offers in Wave 0 is exactly zero.

Figure 4 displays learning for all waves. In our social-preferences vacuum chamber, we observe that demands and offers drop and get closer to a trembling-hand perfect equilibrium with selfish preferences within each wave. However, learning is not very substantial, especially not for the proposers. Further, the changes in behavior are very similar across all waves in our social-preferences vacuum chamber. Kruskal-Wallis tests cannot reject that the distributions of offers are the same across Waves 2 to 8 and that the distributions of demands are the same across Waves 1 to 5 (p-values larger than 0.39). The only exception appears to be Wave 7, in which responders seem to lower their demands more strongly than in all other waves. A Kruskal-Wallis test indeed rejects that the distributions of demands are the same across Waves 1 to 7 (p-value smaller than 0.04).

Our last finding raises the question of why responders in Wave 7 reduce their demands more than in Waves 1 to 5, although all face essentially the same offers. As shown before, responders in Wave 7 make the same last-period demands as those in the previous waves. Therefore, their changes in demands must be driven by unusually high initial demands. Mann-Whitney rank-sum tests indeed show that initial demands in Wave 7 are significantly higher than those in Waves 3 and 5 (p-values weakly smaller than 0.036). The difference to the initial demands in Wave 1 is marginally insignificant. But there, the difference in the changes in demands is also less pronounced comparing these two waves (p-value of 0.137).

Finally, Figure 4 also suggests that social preferences and learning are interrelated. In Wave 0, we find that 47% of the responders and 57% of the proposers do not change their behavior at all, and 81% and 89% adjust their demands and offers by at most one. After Wave 0, we see more adjustments: 39% of the responders and 35% of the proposers do not change their behavior at all, and 59% and 70% of them change their demands and offers by at most one. This change in the stability in demands and offers suggests that when humans play against humans, they have a more fixed idea of what strategy they would like to pursue. Regression analyses reported in the appendix support these impressions. See Table 3 in the appendix for more details.

5 Conclusion

In this paper, we develop a social-preferences vacuum chamber to reduce social preferences' confounding effects while still allowing for indirect strategic interaction. We apply our vacuum chamber to study the empirical relevance of trembling-hand perfection in a normal-form version of the ultimatum game. We find social preferences to affect responder behavior profoundly. There is a sharp drop in the demands in the first wave with computer opponents compared to the initial wave where humans interact with humans. However, social preferences do not seem to play a large role for the proposers whose behavior is surprisingly stable across waves. Even towards the end of our experiment, behavior stays quite far from any trembling-hand perfect equilibrium with selfish preferences. We conclude that trembling-hand perfection is not empirically relevant in our setup. Future research could explore the conditions under which trembling-hand perfection better predicts human behavior.

We also find that the learning dynamics within a wave are different in an environment without social preferences. If humans interact with computers, both proposers and responders adjust their behavior in the direction of the trembling-hand perfect equilibrium. We do not find such an adjustment if humans play against humans. This interaction of learning and social preferences highlights the importance of considering social preferences as a potential confound when analyzing strategic interactions. Our social-preferences vacuum chamber is functional and awaits application to study other concepts from game theory.

References

Cooper, D. J. and Dutcher, G. (2011). The dynamics of responder behavior in ultimatum games: A meta study, *Experimental Economics* **14**: 519–546.

Cooper, D. J., Feltovich, N., Roth, A. R. and Zwick, R. (2003). Relative versus absolute speef of adjustment in strategic environments: Responder behavior in ultimatum games, *Experimental Economics* **6**: 181–207.

Cooper, R., DeJong, D. V., Forsythe, R. and Ross, T. W. (1996). Cooperation without reputation: Experimental evidence from prisoner's dilemma games, *Games and Economic Behavior* **12**(2): 187–218.

Ferraro, P. J., Rondeau, D. and Poe, G. L. (2003). Detecting other-regarding behavior with virtual players, *Journal of Economic Behavior and Organization* **51**: 99–109.

Ferraro, P. J. and Vossler, C. A. (2010). The source and significance of confusion in public goods experiments, *B.E. Journal of Economic Analysis and Policy* **10**(1): 1–40.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments, *Experimental Economics* **10(2)**: 171–178.

Gale, J., Binmore, K. G. and Samuelson, L. (1995). Learning to be imperfect: The ultimatum game, *Games and Economic Behavior* **8**: 56–90.

Gneezy, U. and Potters, J. (1997). An experiment on risk taking an evaluation periods, *Quarterly Journal of Economics* **112**(2): 631–645.

Greiner, B. (2015). Subject pool recruitment procedures: Organizing experiments with ORSEE, *Journal of the Economic Science Association* **1**(1): 114–125.

Güth, W., Schmittberger, R. and Tietz, R. (1990). Ultimatum bargaining behaviour - a survey and comparison of experimental results, *Journal of Economic Psychology* **11**: 417–449.

Harrison, G. W. and McCabe, K. A. (1996). Expectations and fairness in a simple bargaining experiment, *International Journal of Game Theory* **25**: 303–327.

Houser, D. and Kurzban, R. (2002). Revisiting kindness and confusion in public goods experiments, *American Economic Review* **92**(4): 1062–1069.

Johnson, E. J., Camerer, C., Sen, S. and Rymon, T. (2002). Detecting failures of backward induction: Monitoring information search in sequential bargaining, *Journal of Economic Theory* **104**: 16–47.

Kandori, M. (1992). Social norms and community enforcement, *Review of Economic Studies* **59**: 63–80.

Levine, D. K. (2012). *Is Behavioral Economics Doomed? The Ordinary versus the Extraordinary*, Open Books Publisher, Cambridge UK.

List, J. A. and Cherry, T. L. (2000). Learning to accept in ultimatum games: Evidence from an experimental design tha generates low offers, *Experimental Economics* **3**: 11–29.

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

Sanfey, A. G., Rilling, J. K., Aronson, J. A., Nystrom, L. E. and Cohen, J. D. (2003). The neural basis of economic decision-making in the ultimatum game, *Science* **300**(1755-1758).

Smith, V. L. (1976). Experimental economics: Induced value theory, *American Economic Review* **66**(2): 274–279.

van't Wout, M., Kahn, R. S., Sanfey, A. G. and Aleman, A. (2006). Affective state and decision-making in the ultimatum game, *Experimental Brain Research* **169**: 564–568.

Weibull, J. W. (2004). Testing game theory, *in* S. Huck (ed.), *Advances in Understanding Strategic Behaviour; Game Theory, Experiments and Bounded Rationality. Essays in Honour of Werner Güth*, Palgrave Macmillan, pp. 85–104.

Winter, E. and Zamir, S. (2005). An experiment with ultimatum bargaining in a changing environment, *Japanese Economic Review* **56**(3): 363–385.

Appendix A: Regression Results

	Demands		Offers	
	(1)	(2)	(3)	(4)
Wave 1	-1.57***	-1.63***		
	(0.25)	(0.25)		
Wave 2			-0.21	-0.14
			(0.25)	(0.25)
Female		-0.45		0.12
		(0.30)		(0.30)
Age		0.02		-0.01
		(0.04)		(0.03)
Risky Investment		0.08		0.25
		(0.22)		(0.25)
Constant	3.45***	3.20***	4.34***	4.18***
	(0.09)	(0.96)	(0.14)	(0.68)
Number of Observations	94	94	94	94
adjusted R^2	0.18	0.18	-0.00	-0.02

Table 1: Last-Period Demands and Offers with Human and Computer Opponents

Note: Linear regressions comparing last-period behavior in Wave 0 to last-period behavior in the first wave with computer players. The dependent variables are last-period demands in regressions (1) and (2) and last-period offers in regressions (3) and (4). The omitted reference category is Wave 0. Standard errors in parentheses, clustered at the session level for Wave 0 and at the individual participant level for Waves 1 and 2. Stars indicate statistical significance at the 10%, 5% and 1% level.

	Demands			Offers		
	(1)	(2)	(3)	(4)	(5)	(6)
Wave	-0.04	0.09	0.11	-0.08*	-0.24	-0.18
	(0.05)	(0.23)	(0.23)	(0.04)	(0.25)	(0.25)
Wave Squared		-0.02	-0.02		0.02	0.01
		(0.03)	(0.03)		(0.02)	(0.02)
Female			-0.16			0.30
			(0.24)			(0.21)
Age			-0.03			0.02
			(0.03)			(0.02)
Risky Investment			-0.09			-0.18
			(0.20)			(0.16)
Constant	1.99***	1.80***	2.63***	4.27***	4.59***	4.11***
	(0.23)	(0.38)	(0.83)	(0.24)	(0.56)	(0.73)
Number of Observations	188	188	188	205	205	205
adjusted R^2	-0.00	-0.00	-0.01	0.01	0.01	0.03

Table 2: Last-Period Demands and Offers in the Vacuum Chamber

Note: Linear regressions studying last-period behavior after the initial Wave 0. The dependent variables are last-period demands in regressions (1) to (3) and last-period offers in regressions (4) and (6). Standard errors in parentheses. Stars indicate statistical significance at the 10%, 5% and 1% level.

	Change in Demands		Change in Offers	
	(1)	(2)	(3)	(4)
Wave 1	-1.468***	-1.464***		
	(0.381)	(0.385)		
Wave 3	-1.064***	-1.065***		
	(0.309)	(0.319)		
Wave 5	-1.340***	-1.338***		
	(0.320)	(0.326)		
Wave 7	-2.298***	-2.322***		
	(0.365)	(0.370)		
Wave 2			-0.745***	-0.778***
			(0.286)	(0.289)
Wave 4			-0.702***	-0.744***
			(0.206)	(0.209)
Wave 6			-0.851**	-0.879**
			(0.358)	(0.364)
Wave 8			-0.781***	-0.829***
			(0.257)	(0.275)
Female		-0.0725		-0.123
		(0.238)		(0.251)
Age		-0.00503		0.0191
		(0.0326)		(0.0142)
Risky Investment		-0.165		-0.115
		(0.216)		(0.190)
Constant	0.191	0.546	0	-0.205
	(0.205)	(0.859)	(0.0527)	(0.471)
Number of Observations	235	235	252	252
adjusted R^2	0.141	0.144	0.027	0.032

Table 3: Changes in Demands and Offers with Human and Computer Opponents

Note: Results of linear regressions where the dependent variables are changes in demands in regressions (1) and (2) and changes in offers in regressions (3) and (4). Standard errors in parentheses, clustered at the session level for Wave 0 and at the individual participant level for the later waves. Stars indicate statistical significance at the 10%, 5% and 1% level.

Appendix B: Experimental Instructions

Please find attached a translation of the originally German instructions. The differences in the instructions for the two treatments are highlighted by different colors.

Instructions

Thank you very much for your participation in this experiment. The amount of money you can earn depends on your own decisions and possibly on the decisions of other participants. We will pay you the full amount you have made in cash at the end of the experiment. We assure you that your final earnings will remain confidential. Therefore, we will not pass on any information about your earnings to other participants in the experiment.

Please read the instructions for the experiment below. Please do not communicate with the other participants during the experiment. Please raise your hand if you have a question. One of the experimenters will then come to your table to answer your question in private.

Today's experiment consists of two completely independent experiments.

Experiment 1

In the first experiment, you do not interact with any other participant. In the beginning, you will receive 2 euros from us, of which you can invest any amount Z in steps of 20 cents in a risky project. With a probability of 50%, the project is successful, and you will get back 2.5 times the amount Z you invested. With the composite probability of 50%, the project fails, and you get nothing back. Regardless of the outcome of the coin toss, you receive the amount (2-Z) that you did not invest. You will find out your final payment and, thus, the project's success only after the second experiment.

Example 1: Suppose you do not invest anything in the project. You will then receive 2 euros from this part of the experiment regardless of the project's success.

Example 2: Assume you invest Z=1 euro in the project. If the project is successful, you will receive the 2.50 euros profit from the project, plus the one euro you did not invest. Therefore, we would pay you 3.50 euros for this part of the experiment towards the end of the experiment. If the project is not successful, you will only receive the one euro you did not invest at the end of the experiment.

Experiment 2 - Overview

There are two roles in this experiment: player A and player B. These players negotiate the division of ten euros. Player A first makes an offer to player B. Offer X can be between 0 euros and ten euros in steps of one euro. Player B can accept or reject player A's offer. If player B accepts the offer X, he receives X euros, and player A receives the remaining 10-X euros. If player B rejects the offer X, both players receive zero euros.

For experimental reasons, player B must decide which offers he is willing to accept before he learns about the actual offer from player A. For this purpose, player B indicates his minimum demand Y. This minimum requirement Y can be between 0 euros and ten euros in steps of one euro. When player A determines his offer, he does not know the minimum requirement of player B.

The offer X and the minimum requirement Y determine the payments for players A and B as follows. If player B demands equal or less than A offers (Y less than or equal to X), the division rule is accepted. Then player B receives the offer of X euros, while player A gets the remaining 10-X euros. If X is less than Y (Y greater than X), the division rule is rejected. Both players then receive 0 euros.

Example 1: Assume that player A makes an offer of X=7 euros to player B, while player B sets Y=5 as minimum demand. Because the minimum demand Y is smaller than the offer X, the offer X is accepted. Player B then receives the offer of 7 euros, while player A receives the remaining 3 euros.

Example 2: Assume that player A makes an offer of X=3 euros to player B, while player B sets Y=8 as the minimum requirement. Because the minimum demand Y is greater than the offer X, offer X is rejected. Both players then receive zero euros.

Experiment 2 - Implementation

Before Experiment 2 begins, we will ask you some comprehension questions to make sure you understand the game's flow. If you correctly answer one of these questions, you can increase your final payment by one euro. The experiment can only begin when all participants have correctly answered all the comprehension questions.

[Ultimatum Human] As said before, there are two roles in the experiment: player A and player B. At the beginning of the experiment, you will learn your role in the experiment. You will keep the same role throughout the experiment. Other participants in today's experiment will be assigned the other role. You will, thus, interact directly with other participants in the experiment. Therefore, your behavior directly determines your payment and the payment of other human participants in today's experiment.

[Ultimatum Computer] As said before, there are two roles in this experiment: player A and player B. At the beginning of the experiment, you will learn your role in the experiment. You will keep the same role throughout the experiment. A computer simulates the other role in the experiment. Therefore, you do not interact with another human participant in the experiment, but a computer simulates your opponents. This means that your behavior does influence only your payment, but does not directly affect the payment of other human participants. The behavior of the computer-simulated opponents corresponds to the behavior of human participants in earlier sessions of the experiment.

The experiment consists of ten identical rounds. In each round, you negotiate with another participant for the distribution of ten euros. The interaction between you and the other participant is as described above. After each round, you learn your payment and the payment for the [Ultimatum Human] human player with whom you have interacted. Therefore, as Player B, you will learn the offer from Player A, and as Player A, you will learn whether Player B has accepted your offer. You will not learn more.

[Ultimatum Human] We make sure that you will only interact with another participant as a fellow player once. Further, the matching algorithm ensures that you cannot influence the behavior of your future players by your behavior in any round. Basically, you can consider the ten rounds as completely isolated interactions.

[Ultimatum Computer] We make sure that you will only interact with another computer-simulated participant as a fellow player once. Further, the matching algorithm ensures that you cannot influence the behavior of your future computer-simulated opponents. Basically, you can consider the ten rounds as completely isolated interactions.

For your payment, we randomly select one of the ten rounds, each with the same probability. We will pay you the amounts earned in this round. In addition, each participant receives a fixed participation fee of 5 euros, regardless of his or her decisions.