

**Theoretical and empirical analysis of the evolution of
cooperation**

Dissertation zur Erlangung des mathematisch-naturwissenschaftlichen

Doktorgrades *Doctor rerum naturalium* der

Georg-August-Universität Göttingen

im Promotionsprogramm *Behavior and Cognition*

der Georg-August University School of Science (GAUSS)

vorgelegt von

Peter Bednarik

geboren am 11.02.1984 in Bratislava

Göttingen, August 2014

Betreuungsausschuss:

Dirk Semmann, Nachwuchsgruppe Evolution von Kooperation und prosozialem Verhalten, Courant Forschungszentrum Evolution des Sozialverhaltens, Georg-August-Universität Göttingen

Margarete Boos, Sozial- und Kommunikationspsychologie, Gregor-Elias-Müller-Institut für Psychologie, Georg-August-Universität Göttingen

Julia Fischer, Kognitive Ethologie, Deutsches Primatenzentrum und Georg-August-Universität Göttingen

Mitglieder der Prüfungskommission

Referent: *Dirk Semmann*,

Nachwuchsgruppe Evolution von Kooperation und prosozialem Verhalten, Courant Forschungszentrum Evolution des Sozialverhaltens, Georg-August-Universität Göttingen

Korreferentin: *Margarete Boos*,

Sozial- und Kommunikationspsychologie, Gregor-Elias-Müller-Institut für Psychologie, Georg-August-Universität Göttingen

Weitere Mitglieder der Prüfungskommission

Julia Fischer, Kognitive Ethologie, Deutsches Primatenzentrum und Georg-August-Universität Göttingen

Julia Ostner, Soziale Evolution von Primaten, Georg-August-Universität Göttingen

Oliver Schülke, Soziale Evolution von Primaten, Georg-August-Universität Göttingen

Bernhard Fink, Biologische Persönlichkeitspsychologie und Diagnostik, Georg-August-Universität Göttingen

Tag der mündlichen Prüfung: 10.09.2014

Contents

General Introduction.....	7
Chapter I: Costs for switching partners reduce network dynamics but not cooperative behaviour	18
Chapter II: Optimal Decision-Making in the Judge-Advisor System	48
General Discussion	76
Summary.....	79
Zusammenfassung	81
References.....	84
Acknowledgements.....	92
Curriculum Vitae	93
Erklärung über eigene Leistungen	95

General Introduction

Cooperative behavior is widespread among humans and the animal kingdom. It occurs in very simple organisms like social amoebae (Santorelli, Thompson et al. 2008), viruses (Turner and Chao 2003), eusocial insects (Nowak, Tarnita et al. 2010), fish (Bshary and Grutter 2006), birds (Ligon 1983), non-human primates (Cheney, Moscovice et al. 2010) and, of course, humans (Melis and Semmann 2010). However, in terms of evolution by natural selection, such cooperative behavior seems to be harder to explain than selfish behavior. The heart of the problem can be conveniently described with the language of game theory. Game theory is a mathematical theory to describe game-like situations. Originally, game theory was used to analyze classic games like chess or poker (Zermelo 1913). Soon, economists realized that it can be applied to a variety of economic interactions (e.g. Nash Jr 1950). Since the 1970's, evolutionary game theory became a big field of research by itself (Smith and Price 1973), extending its applications to biology, ecology, mathematics, physics and many more branches.

So, how does game theory help to understand the evolution of cooperation? One of the most prominent examples of a “game” in the sense of game theory is the so-called prisoner’s dilemma game. Here, we assume that two players must decide simultaneously between two options. They can either (i) cooperate or (ii) defect. Their earnings depend on both, their own action and the action of the other player. If player I cooperated, but player II defected, player I receives the so-called suckers payoff S and player II gets the temptation T . If both players cooperate, they receive the reward R and, finally if both defect, they receive the punishment P . The precise numerical values of S ,

T, R and P can vary but the order between them is characteristic for the prisoner's dilemma game:

$$T > R > P > S$$

Note that irrespective of the other player's action, defecting always yields the better result, because $T > R$ and $P > S$. However, if both players defect, they receive less than if both had cooperated. Thus, self-interested behavior leads to a poor outcome for both players. More generally, situations where group-interest and self-interest are in conflict are called social dilemmas. Another prominent example besides the prisoner's dilemma is the so-called public goods game, which describes a similar situation for more than two players.

Clearly, without any modification of this game, cooperative behavior cannot persist since defectors will always bring home the higher payoff. However, if embedded in a certain context, cooperation might be able to out-compete defection. Because such contexts work like a mechanism to promote cooperative behavior, they are often referred to as "mechanisms of cooperation". One well-established distinction between such mechanisms is due to Martin Nowak (Nowak 2006):

Mechanisms of Cooperation:

- kin selection
- direct reciprocity
- indirect reciprocity
- group (multi-level) selection
- network reciprocity

Kin selection (Hamilton 1964) refers to a situation where closely related individuals encounter a social dilemma. A gene that codes for cooperative behavior can be selected for, because helping close relatives implies helping bearers of this very gene. Famous examples are eusocial insects.

Direct reciprocity (Trivers 1971) is a straight-forward mechanism where we imagine that the prisoner's dilemma game is played repeatedly. As long as the chance to encounter an individual in the future is high enough, cooperative behavior can out-compete defective behavior. Since multiple rounds of prisoner's dilemma games are played, individuals can adopt potentially infinitely many strategies (i.e. combinations of cooperate and defect). The famous computer tournament by the political scientist Robert Axelrod (Axelrod and Hamilton 1981) aimed to detect best-performing strategies. It is remarkable that of all participating strategies, one of the simplest turned out to perform best: It played "cooperate" in the first round and then copied the opponent's move of the previous round; therefore it was called Tit-for-Tat.

Indirect reciprocity (Nowak and Sigmund 1998) works somewhat similar, except that the cooperative action is received by a third person. Cooperative behavior can succeed if individuals cooperate only with "good" individuals, but defect with "bad" individuals. Here, "good" individuals are those who helped others whereas "bad" individuals are those who did not help. Thus, reputation is the essence of indirect reciprocity.

Multi-level selection (e.g. Traulsen and Nowak 2006) refers to the fact that there might be different units of selection: Genes, groups of genes, individuals or groups of individuals. What might appear as cooperative behavior at one level of selection could

be something different at another level. If selection is faster on the latter level, cooperation on the former level could persist.

Network reciprocity is in the focus of the larger part of this thesis (chapter I) and thus deserves a more thorough introduction. In the 90's of the previous century, an evolutionary model (Nowak and May 1992) showed that cooperation could, in principle, be promoted by spatial structure. Here, individuals interact with neighbors, instead of with random individuals of the entire population. Thus, cooperators can assort into clusters and achieve higher payoffs than defectors. Interestingly, experimental attempts to find evidence of such a cooperation-enhancing mechanism in humans failed (Cassar 2007; Kirchkamp and Nagel 2007; Grujić, Fosco et al. 2010; Traulsen, Semmann et al. 2010; Gracia-Lázaro, Ferrer et al. 2012; Grujić, Rohl et al. 2012). To understand this, we must look at the assumptions of the evolutionary model. First, it uses the prisoner's dilemma game (see above) to model the interaction. Another model (Hauert and Doebeli 2004) showed that a very similar but slightly different game, the snowdrift game (Sugden 1986), does not lead to cooperation under the same circumstances. The only difference between the snowdrift game and the prisoner's dilemma game is that in the former, the order of payoffs is $T > R > S > P$. However, the above-mentioned experiments failed to show increased cooperation on spatial structure despite using the prisoner's dilemma game. The most likely reason for this is that humans do not update their behavior in any way that is assumed by evolutionary models. In general, evolutionary models consist of two major parts. First, the interaction between individuals is modeled, e.g. as a prisoner's dilemma game. Second, assumptions are made how individuals update their behavior in the game, e.g. whether they cooperate or defect. So-called imitation updating is one of the most frequently used

update mechanisms in evolutionary models. Therefore, we shall look at it in more detail. Out of the entire (possibly large) population, we look at four individuals (Fig.1, colored circles).




Figure 1. Segment of a network. Circles represent individuals, connected individuals can interact.

Lines between individuals mean that they can interact with each other (e.g. play a prisoner's dilemma game). Let us focus on the blue individual. The other individuals may also have interactions with other individuals that are not shown in Fig.1. The algorithm of a typical evolutionary model could look like this:

Step 1

Individuals play prisoner's dilemma games with all their neighbors (neighbors are indicated by a line between individuals). Their payoff from all of their prisoner's dilemma games is summed up and is called their payoff. An example of possible payoffs is shown in Fig.2.

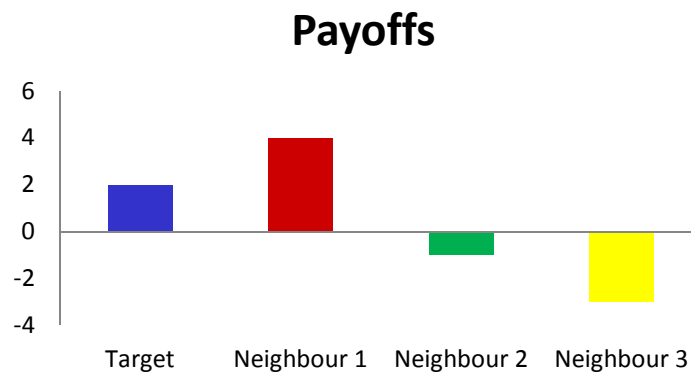


Figure 2. Examples of possible payoffs resulting from prisoner's dilemma games.

Step 2

To calculate the fitness, we use the function

$$fitness = 1 - w + w * payoff$$

where the selection coefficient w determines the strength of selection: $w = 1$ for strong selection and $w \rightarrow 0$ for weak selection. Thus, we get the fitness following values (Fig.3.)

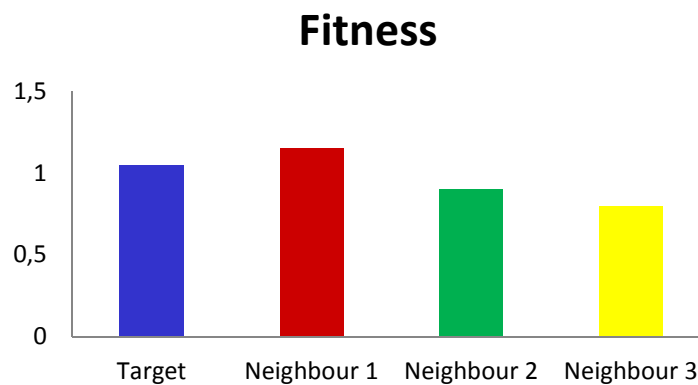


Figure 3. Fitness values resulting from payoffs.

Note that we can always choose a w such that the fitness of all possible payoffs is positive, which is important for calculating the relative fitness:

Step 3

$$\text{relative fitness} = \text{fitness} / (\text{sum of fitness of all agents})$$

Relative Fitness

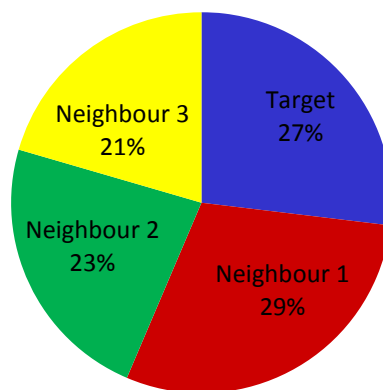


Figure 4. Relative fitness of individuals.

The relative fitness determines the probability that the target player will have the strategy of the respective agents in the next round. In this case, target node will keep its strategy with probability 27%, and will switch to the strategy of one of its neighbors with probabilities according to the relative fitness.

As we can see, evolutionary models like the one described assume rather specific and simplified learning behavior about the individuals. This may be a justifiable abstraction for very simple organisms, but it seems very unrealistic for humans. A recent study suggests that humans learning behavior resembles closely to so-called moody conditional cooperation (Grujić, Gracia-Lázaro et al. 2014). Here, individuals are more likely to cooperate if they, themselves cooperated in the previous round and are surrounded by enough other cooperators. However, as an updating mechanism, moody conditional cooperation does not seem to promote network reciprocity (Gracia-Lázaro, Cuesta et al. 2012). Therefore, static networks do not enhance cooperation among humans

However, the situation is different on dynamic networks. In contrast to the static networks that we discussed above, dynamic networks allow the individuals to change their neighbors over time. Recent models (e.g. Santos, Pacheco et al. 2006; Fu, Hauert et al. 2008; Pacheco, Traulsen et al. 2008; Fu, Wu et al. 2009; Wu, Zhou et al. 2010) have shown that dynamic networks allow for cooperation levels beyond the ones of static networks. Unlike with static networks, experiments with human participants did indeed show increased cooperation on dynamic networks (Fehl, van der Post et al. 2011; Rand, Arbesman et al. 2011; Wang, Suri et al. 2012). However, as I argue in chapter I, the main reason for this is different from the mechanisms suggested in the models. Although models use a variety of different assumptions, the cooperation levels always depend on how often individuals change their neighbours. In chapter I, I provide evidence where cooperation levels only depend on whether participants have the option to change their partners, but not on how often they actually do it.

While in evolution of cooperation, the state of empirical research is constantly lagging behind the development of new models, the situation seems to be rather opposite in behavioural psychology. In chapter II, I will develop a normative model of decision making. Although this is a rather different topic than the evolution of cooperation I will point out in the general discussion section of this thesis how the two topics could inspire each other. One important aspect of decision making is how to evaluate available information, e.g. advice from experts. A robust finding from various experiments (e.g. Harvey and Fischer 1997; Yaniv and Kleinberger 2000) is that participants do not use the information given by advisors in an optimal way, although advice generally increases the accuracy of judgements (e.g. Sniezek, Schrah et al. 2004; Gino and Schweitzer 2008; Minson, Liberman et al. 2011). A good framework to study how humans use information by peers is the so-called Judge-Advisor-System (Sniezek and Buckley 1995). Here, one person, “judge”, first estimates an unknown quantity. Then, the judge gets advice from another independent person, the “advisor”. Finally, the judge adjusts her or his estimate by taking into account the advisor’s estimate. The question is how much weight should be put on either of the two initial estimates. The most prominent model on the Judge-Advisor-System (Soll and Larrick 2009) focuses on two strategies: *averaging* and *choosing*. Averaging is a strategy that weights both initial estimates equally and in the choosing strategy one first attempts to find out the more accurate initial guess and neglect the other. To this end, the model (Soll and Larrick 2009) uses three parameters: the ability differences between judge and advisor (A), the probability that the judge can identify these differences (P) and the amount of systematic bias of both judge and advisor (R). It is called (R) because systematic bias of judge and advisor leads to redundant information. Because of these three parameters,

the model has become known as the PAR model. The main predictions from the PAR models are that judges should prefer the averaging strategy in most cases. Only if the ability differences are large and easy to detect, choosing outperforms averaging. Additionally, choosing performs well if the amount of systematic bias is very high. In this case, both estimates tend to either over- or underestimate the target quantity and therefore choosing outperforms averaging. Although the restriction to these two simple strategies seems to be backed up by their empirical data, where roughly 70% of the participants used one of the two, neither of them is optimal, except in rare examples. Therefore, I will present a model that compares the two strategies with a more accurate weighting strategy (which shall simply be called *weighting*). Weighting aims to identify the ability difference quantitatively (i.e. not only who is better, but also how much) and then assign weights accordingly. In chapter II, I derive a mathematical model to show that whenever the difference is large and moderately detectable, weighting is the best strategy.

Chapter I: Costs for switching partners reduce network dynamics but not cooperative behaviour

with Katrin Fehl¹ & Dirk Semmann¹,

Proceedings of the Royal Society B: Biological Sciences **281**: 20141661

¹Research Group Evolution of Cooperation and Prosocial Behaviour, Courant
Research Centre Evolution of Social Behaviour, University of Göttingen,
Kellnerweg 6, 37077 Göttingen, Germany

SUMMARY

Social networks represent the structuring of interactions between group members. Above all, many interactions are profoundly cooperative in humans and other animals. In accordance with this natural observation, theoretical work demonstrates that certain network structures favour the evolution of cooperation. Yet recent experimental evidence suggests that static networks do not enhance cooperative behaviour in humans. In contrast, dynamic networks do foster cooperation. However, costs associated with dynamism like time or resource investments in finding and establishing new partnerships have been neglected so far. Here, we show that human participants are much less likely to break links when costs arise for building new links. Especially, when costs were high the network was nearly static. Surprisingly, cooperation levels in prisoner's dilemma games were not affected by reduced dynamism in social networks. We conclude that the mere *potential* to quit collaborations is sufficient in humans to reach high levels of cooperative behaviour. Effects of self-structuring processes or assortment on the network played a minor role: participants simply adjusted their cooperative behaviour in response to the threats of losing a partner or of being expelled.

KEYWORDS cooperation, dynamic network, partner switching, prisoner's dilemma, evolutionary game theory

INTRODUCTION

One common way to deal with an unpleasant peer is to get out of her or his way and to become friends with someone else. In fact, the breaking of links serves as an effective mechanism to control cheating in social relationships and the evolution of cooperation among unrelated individuals (e.g. Santos, Pacheco et al. 2006; Fu, Hauert et al. 2008; Pacheco, Traulsen et al. 2008; Fu, Wu et al. 2009; Wu, Zhou et al. 2010). The process of finding new friends leads to assortment and populations with a social structure. First, this means that not all but only subsets of individuals interact with each other. Second, the social structure is dynamic. As such, dynamism in social networks has for instance been described in humans (Kossinets and Watts 2006; Saramäki, Leicht et al. 2014), bottlenose dolphins (Lusseau and Newman 2004), or chacma baboons (Henzi, Lusseau et al. 2009). Recent experiments (Fehl, van der Post et al. 2011; Rand, Arbesman et al. 2011; Wang, Suri et al. 2012) show that human cooperation is substantially higher in dynamic networks compared to static ones. For that matter, static network structures - though in theory impacting cooperation positively (Nowak and May 1992; but see Hauert and Doebeli 2004; Lieberman, Hauert et al. 2005; Ohtsuki, Hauert et al. 2006) - have trouble producing cooperative outcomes in experimental settings (Cassar 2007; Kirchkamp and Nagel 2007; Grujić, Fosco et al. 2010; Traulsen, Semmann et al. 2010; Gracia-Lázaro, Ferrer et al. 2012; Grujić, Rohl et al. 2012). It has to be noted that while most of the mentioned research assume that individuals make one decision for all partners, we allow them to choose freely for each of their partners, as in (Fehl, van der Post et al. 2011; Rand, Arbesman et al. 2011; Wang, Suri et al. 2012). This setup seems to be more realistic for human societies.

Conscious care of one's partnerships is beneficial for cooperation in humans. Both, evolutionary models and experiments, show that the directed breaking of social links to cheaters or so-called defectors in joint endeavours eventually leads to cooperative outcomes (Fu, Wu et al. 2009; Szolnoki and Perc 2009; Perc and Szolnoki 2010; Wu, Zhou et al. 2010; Fehl, van der Post et al. 2011; Rand, Arbesman et al. 2011; Wang, Suri et al. 2012; Shirado, Fu et al. 2013); further, long-term relationships occur between like-minded cooperators (Pacheco, Traulsen et al. 2006; Santos, Pacheco et al. 2006; Fu, Hauert et al. 2008; Wu, Zhou et al. 2010; Fehl, van der Post et al. 2011; Wang, Suri et al. 2012). However, not only humans make use of partner switching as control mechanism to achieve cooperative outcomes: Long-nosed parrotfishes stop interacting (at least temporarily) with their cleaner wrasses and choose a different cleaner station when being cheated on (Bshary and Schaffer 2002).

Real-world interactions offer a large variety of potential costs of partner switching in terms of resources like food, grooming, or money. Further, psychological, effort- and time-based costs may also occur. Theoretical evidence is sparse in regard to modelling link-related or partner switching-related costs and their effects on the evolution of cooperation. When costs in terms of distance - the further away a partner the larger associated (e.g. contacting or traveling) costs - arise for individuals, they are prone to choose those in close reach and hence with lower associated costs (Li, Min et al. 2013). Other approaches impacting the evolution of cooperation include the introduction of migration costs (Liu, Chen et al. 2012), participation costs (Masuda 2007; Szolnoki, Perc et al. 2008), and time-related costs (Poncela, G3mez-Gardees et al. 2011). Surprisingly, present experiments on cooperative behaviour in dynamic human networks (Fehl, van der Post et al. 2011; Rand, Arbesman et al. 2011; Wang, Suri et al.

2012; Shirado, Fu et al. 2013) assume that altering partners is cost-free. In sum, exploring the consequences of partner-switching costs on cooperation need to be investigated as the experimental evidence is lagging behind of insights of natural observations and the development of models.

While partner-switching costs might impact the positive effect of dynamic social networks on cooperative behaviour, they can also affect the dynamism per se. Conflicting work exists on the question *how much* partner switching in dynamic networks is optimal. On the one hand, higher rates of network dynamism should generally lead to more cooperation (Santos, Pacheco et al. 2006; Fu, Hauert et al. 2008; Rand, Arbesman et al. 2011; Wang, Suri et al. 2012). On the other hand, recent work suggest that the optimal rate of changing partners must be in medium ranges (Fu, Wu et al. 2009; Shirado, Fu et al. 2013) or the maximum number of partners has to be limited (Szolnoki, Perc et al. 2008; Szolnoki and Perc 2009). If defective individuals switch their partners too rapidly, they can exploit newly-linked partners and thus will out-compete cooperative individuals. According to (Fowler and Christakis 2010), another reason may be that defectors need to stay connected to cooperators in order to learn and adopt their strategy, yet recent experimental data suggest that humans imitate only selfish behaviour, but not cooperative behaviour (Jordan, Jordan et al. 2013). Former studies (Fu, Hauert et al. 2008; Fu, Wu et al. 2009; Wang, Suri et al. 2012; Shirado, Fu et al. 2013) approach this question by controlling for how often partners can be exchanged. Instead, we will keep the opportunity to switch partners at a maximum and constant throughout treatments, but expect that varying levels of costs influence partner-switching rates - with a so far unexplored impact on cooperative behaviour in humans.

Here, we examine the cooperative behaviour of human participants while either a static social structure or dynamic social structures define interaction partners. The switching of partners in dynamic networks occurs at (i) no, (ii) low, or (iii) high costs (i.e. breaking of social links is cost-free; however, costs occur for setting up a new link). Participants interact with three partners and play independent prisoner's dilemma games (PD; Rapoport and Chammah 1965; Axelrod and Hamilton 1981; see Electronic Supplementary Material [ESM] for a game description). Only if the game allows for repeated interactions then direct reciprocity (Trivers 1971; Nowak and Sigmund 1992; Nowak and Sigmund 1993) can produce cooperative outcomes. In general, as our experimental setup permits repeated interactions without a known endpoint, we expect cooperative behaviour of participants in both social structures. In particular, cooperation levels are predicted to be higher in the three dynamic networks (despite possible costs for partner switching). Especially, we will address the impact of costs of partner switching in dynamic networks: As costs for new partners increase, we expect participants' tendency to end partnerships will decrease and hence reduce the dynamism in the network. We therefore infer reduced cooperative behaviour under low and high costs compared to no costs for partner switching. Nonetheless, we predict to find directed link breaking in all dynamic network setup: mainly to unwanted, that is to defective, participants. In addition, we will examine the specific link-breaking behaviour that occurs along with no, low, or high costs of setting new partnerships.

MATERIAL AND METHODS

The participants

We conducted computerized experiments with 400 students, tested in fall 2009 and 2012. Students were recruited from a German University via the online recruitment system ORSEE (Greiner 2004) and came from a broad range of disciplines. They were composed of 49% females and were aged 22.07 ± 3.194 (mean \pm s.d.). Upon arrival participants were randomly seated in front of computers separated by opaque partitions. Participants were informed via written instructions about the game rules (available upon request) and came to know that their decisions were made anonymous towards other participants and the experimenters. They were forbidden to communicate except via computers. To allow for in-game identification while ensuring anonymity in regard to their real identity participants were given pseudonyms (randomly-assigned names of moons of the solar system). These pseudonyms were also used to ensure anonymous payment at the end (as described in Semmann, Krambeck et al. 2005; known by participants from written instructions). Sessions lasted approximately 90 minutes and participants earned $17.53\text{€} \pm 5.02$.

Static-network treatment and three dynamic-network treatments

We ran 10 sessions for each of our four treatments with 10 participants in each session, respectively: STATIC, DYNAMIC, DYNAMIC10 and DYNAMIC50 (see ESM Fig. S1-S5).

For STATIC interaction partners of participants were defined by the structure of the so-called Petersen graph (Holton and Sheehan 1993). This graph consists of 10 nodes (vertices; here players) and each node is linked to exactly three other nodes summing up

to a total of 15 interaction links. Importantly, players have identical starting positions in terms of network properties. Note, that participants did not receive any information regarding the network nor their position within. At the beginning of each session, participants were randomly assigned to one of these nodes. With each of their three linked partners, participants played independent PD games (i.e. different decisions could be made for different partners). Each of the three games lasted 30 rounds. To avoid end-round effects, participants were not informed about the duration. In each PD round, participants simultaneously decided whether to cooperate or to defect. However, instead of saying “to cooperate” and “to defect” we used the terms “ORANGE” and “BLUE” to avoid morally burdened language (cf. Fehl, van der Post et al. 2011). After all decisions in a PD round were made, participants were shown their own and their partners’ decisions and payoffs. For mutual cooperation (defection) the players received 0.25€ (0.00€) and if one player cooperated and the other defected the former received -0.10€ while the latter received 0.40€. Participants received only local information, that is, they were not informed about the outcomes and payoffs resulting from their partners’ interactions with others. Thereafter, the next round of PD games started. This treatment is called STATIC, because the network did not change throughout the experiment (no partner switching allowed).

Generally, in the three dynamic treatments the initial network in form of the Petersen graph could change over time due to link breaking (in all treatments, the maximum number of partners remained limited to three). DYNAMIC followed the setup of STATIC. Participants played three independent PD games but in addition they were given the option to quit any of their partnerships after learning the results of the current PD round. Independent link-breaking decisions were made. For a link to be broken, at

least one player had to decide to end the partnership. Thereafter, participants were informed about their and their partners' decisions. For each broken link, participants were randomly re-assigned to any other participant with less than three partners. Afterwards, the next round would start. Because of the random re-linking procedure, occasionally no partner could be found and participants had less than three partners (known by participants from written instructions; this situation generates the same payoff as mutual defection in partnerships, 0.00€). Further, depending on the total number of broken links in the network, there was a chance to get re-linked to the same partner (which could be noticed by pseudonyms).

In DYNAMIC10 and DYNAMIC50, participants made PD decisions followed by link-breaking decisions and, additionally, were asked - if they had less than three partners - whether they would like to buy a new link at the cost of 0.10€ / 0.50€, respectively, or not. Thus, there was no automatic re-linking to other available participants. The program randomly assigned partners to those willing to buy a new link; otherwise participants did not receive a new partner in the current round. If no new partner could be connected, the costs were not deducted. Note, that only in DYNAMIC10 participants could immediately compensate the costs of a new link either due to mutual cooperation (0.25€ payoff - 0.10€ costs) or due to exploitation of the new partner (0.40€ payoff - 0.10€ costs). However, DYNAMIC50 would require either the sum of payoffs of at least two partners or the sum of payoffs of at least two rounds with the new partner to compensate the initial costs for a new link.

Statistical analysis

We conducted group-level analyses as the behaviour of participants within each session was interdependent and thus required the unit of analysis to be groups rather than individuals. For statistical analysis R 2.15.2 (R Core Team 2013) was used. Probabilities are reported as two tailed at a 5%-significance level. Pairwise comparisons were corrected using the Bonferroni method (reported below are corrected probabilities for significant results).

RESULTS

Link-breaking behaviour in dynamic networks

Overall, the experiment showed that costs for setting new links affected the willingness of participants to break them. We found significant differences in the link-breaking rate between each of the three dynamic treatments: the higher the costs, the lower the link-breaking rate (Fig. 1a; Wilcoxon rank-sum test: DYNAMIC vs. DYNAMIC10: $W = 0$, $n_{1,2} = 10$, $p < 0.001$; DYNAMIC vs. DYNAMIC50: $W = 0$, $n_{1,2} = 10$, $p < 0.001$, DYNAMIC10 vs. DYNAMIC50: $W = 17$, $n_{1,2} = 10$, $p < 0.05$; see ESM Tab. S1).




Figure 1. Average rates of link breaking in the dynamic networks (a) and average cooperation levels in the prisoner's dilemma games (b). Interactions occurred either on a static structure (STATIC) or on dynamic structures. Here, social links could be broken after each round. Setting new links was either cost-free (DYNAMIC), or else low costs of 0.10€ (DYNAMIC10) or high costs of 0.50€ (DYNAMIC50) arose.

Additionally, we found that links were broken mostly during the first few rounds of the experiment (Fig. 2; comparing link-breaking rates of the average of rounds 1 to 10 against rounds 21 to 30, Wilcoxon signed-rank test: DYNAMIC: $V = 0, n = 10, p < 0.01$; DYNAMIC10: $V = 0, n = 10, p < 0.01$; DYNAMIC50: $V = 1, n = 10, p < 0.01$).

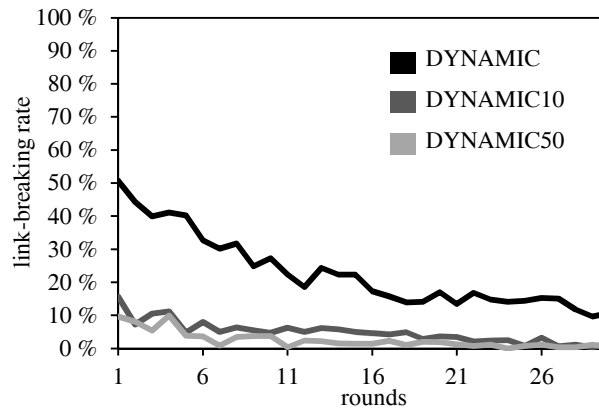


Figure 2. Average link-breaking rates over rounds. Participants played prisoner’s dilemma games with three partners and could break social links after each round. Setting new links was either cost-free (DYNAMIC), or else low costs of 0.10€ (DYNAMIC10) or high costs of 0.50€ (DYNAMIC50) arose. Average link-breaking rates differed significantly between treatments (p ’s < 0.05) and link-breaking rates significantly decreased over time (p ’s < 0.01).

Now, we asked which behaviour was likely to trigger link breaking. Generally, within all dynamic treatments significantly more links were broken to defectors than to cooperators (break rate to defectors: DYNAMIC: $82.9\% \pm 8.3$; DYNAMIC10 $89.5\% \pm 7.0$; DYNAMIC50: $91.3\% \pm 7.5$; Wilcoxon signed-rank test: $V = 0, n = 10, p < 0.01$ for all dynamic treatments). Here, we also found that the link-breaking rates to defectors marginally significantly differed between the dynamic treatments (Kruskal-Wallis test: $\chi^2 = 5.17, df = 2, p = 0.075$). Moreover, when taking the participants own behaviour into account we observed that if both players cooperated participants almost never broke

such a link. Out of a total of 4706 mutually-cooperative PD outcomes in DYNAMIC50 only one link was broken (DYNAMIC10: 1 out of 4828; DYNAMIC: 12 out of 5040).

Cooperative behaviour in the prisoner's dilemma games

Next, we examined the cooperative behaviour of participants as well as its relation to link-breaking behaviour. The overall cooperation levels between the four treatments were significantly different. Detailed analysis revealed that cooperation levels were significantly lower in STATIC compared to all dynamic treatments (Fig. 1b; Wilcoxon rank-sum test: STATIC vs. DYNAMIC: $W = 4$, $n_{1,2} = 10$, $p < 0.001$; STATIC vs. DYNAMIC10: $W = 13$, $n_{1,2} = 10$, $p < 0.05$; STATIC vs. DYNAMIC50: $W = 5$, $n_{1,2} = 10$, $p < 0.01$, for a round-by-round figure, see ESM Fig S2). However, we did not find significant differences between the dynamic treatments themselves (see ESM Tab. S2).

To understand the high levels of cooperation in the dynamic treatments despite very different dynamism in the networks (note, the network structure in DYNAMIC50 was nearly static), we first looked at the readiness to cooperate when a participant received a newly linked partner (i.e. participants who had no current partnership, but who might have met each other before; behaviour of the first round was excluded). We found significantly higher cooperation levels of participants with a new link in DYNAMIC50 than in DYNAMIC (Fig. 3; Wilcoxon rank-sum test: $W = 10$, $n_{1,2} = 10$, $p < 0.01$; see ESM Tab. S3). Second, we examined the willingness of participants to reciprocate defection, that is, to defect when a link to a defector remained in place (note, this required the approval of both players; level of defection in DYNAMIC $64.9\% \pm 8.1$;

DYNAMIC50: $72.2\% \pm 8.5$; DYNAMIC10: $74.8\% \pm 4.3$). We found significant higher levels of reciprocating a defection in DYNAMIC compared to DYNAMIC10 (Wilcoxon rank-sum test: $W = 11$, $n_{1,2} = 10$, $p < 0.05$; see ESM Tab. S4).




Figure 3. Average cooperation level in the prisoner's dilemma game of two newly linked participants. After each prisoner's dilemma round participants could break social links. Setting new links was either cost-free (DYNAMIC), or else low costs of 0.10€ (DYNAMIC10) or high costs of 0.50€ (DYNAMIC50) arose.

DISCUSSION

In accordance with a previous series of experiments (Fehl, van der Post et al. 2011; Rand, Arbesman et al. 2011; Wang, Suri et al. 2012; Shirado, Fu et al. 2013) we contribute evidence on the beneficial effects of dynamic networks on cooperative behaviour in humans. More specifically, we build upon a rather new approach introduced in (Fehl, van der Post et al. 2011) where participants can choose independent actions for each partner. Here, cooperation in dynamic networks went beyond the purely reciprocity-guided cooperation of static networks. We advance previous findings by introducing costs that arise when one is in search for new partners. The willingness to break social links was clearly reduced when low costs arose. When costs were so high that they could not be compensated within a single interaction, participants further reduced their link-breaking behaviour. Nonetheless, high levels of cooperation in prisoner's dilemma games were achieved (i.e. cost treatments did not differ from a treatment without costs). Interestingly, though participants of dynamic networks with no, low or high costs differed quantitatively in their link-breaking behaviour, they followed similar strategies. In all dynamic treatments, link-breaking rates decreased over time, participants mostly broke links to defective players and under mutual cooperation participants almost never broke links.

Contrary to our assumption we found no impact of reduced dynamism on cooperative behaviour. This notion supports the view that the mere option to end partnerships is enough to maintain cooperation. The option of link breaking provides (i) the possibility to get rid of defectors, (ii) the threat of losing a partnership, as well as (iii) the option to stay with like-minded cooperative partners. While avoiding link breaking when costs occurred, participants adjusted their behaviour in the PD games in

response: Participants under high costs were more willing to cooperate with a newly-linked partner compared to participants who did not have to pay for new links; and participants under low costs reacted themselves with higher levels of defection when a link to a defective partner remained in place compared to participants who did not have to pay for new links. These behavioural adjustments played an important role for maintaining high levels of cooperation.

Currently, there is an on-going debate of how much partner switching is necessary to achieve the highest possible level of cooperation (Fu, Hauert et al. 2008; Fu, Wu et al. 2009; Wang, Suri et al. 2012; Shirado, Fu et al. 2013). Previously, different rates of partner-switching were achieved by actively controlling them (e.g. partner-switching every round versus every fifth round of an evolutionary game). Contrasting, our dynamic-network treatments in principal allowed for the same rates of switching. However, rates varied greatly across treatments due to different incentives resulting from costs. In fact, in our high-cost dynamic networks participants broke only 2.5% of the existing links and still cooperated 20% more than in the static networks. Our study implies that it is important to investigate effects which constrain or lead to network dynamics naturally, in addition to the approach of setting different rates exogenously. Thus, our findings can inspire this debate by showing a different perspective of partner-switching.

In our experiment, assignment of new partners was random. As such, we have implemented active link breaking but not active link seeking. Both, models and experimental findings support the view that partner-choice based on behavioural (Wang, Suri et al. 2012; Shirado, Fu et al. 2013) or reputational information (Fu, Hauert et al. 2008; Du and Fu 2011; Wang, Wang et al. 2012) fosters cooperation. Future research

will have to address questions whether costly partner choice will increase individuals' willingness to break social links as they are given more control over future partners as well as whether this will make partner-switching mechanisms faster and more efficient in achieving cooperative outcomes.

Furthermore, our experiment included also a more implicit form of costs as we constrained the maximum number of partners to three (true for all treatments). Time might be such a factor limiting ones interactions; but theoretically have not been found to derogate cooperative incentives of individuals (Poncela, Gómez-Gardeñes et al. 2011). Moreover, while our study provides important first insights into the impact of costs as of seeking new interaction partners, further experiments should investigate the impact of different types of costs on cooperation. As discussed in the theoretical literature (Liu, Chen et al. 2012; Li, Min et al. 2013) such costs may, for instance, relate to geographical distance or migration.

The mechanism of costly link breaking demonstrated here can also be seen as punitive incentive. Related cheater-control mechanisms are ostracism (Cinyabuguma, Page et al. 2005; Maier-Rigaud, Martinsson et al. 2010) and costly punishment (Fehr and Gächter 2002) of defective individuals. An Advantage of link-breaking mechanisms lies in the fact that they are based on individual decisions that lead to assortment and self-organizing processes on the level of the network (Fehl, van der Post et al. 2011). Hence, unlike ostracism they achieve isolation or even final exclusion of defectors without the necessary cooperative and coordinated decision of all group members. Further, costly punishment often suffers from ongoing retaliations and antisocial punishment (e.g. Herrmann, Thöni et al. 2008; Fehl, Sommerfeld et al. 2012). Such detrimental behaviour is avoided in settings where links can simply be broken. Which

of the two - partner switching or costly punishment - is the more favourable behaviour most likely depends on the environmental or social ecology of individuals. The interspecies mutualism of the cleaner wrasse and their clients (i.e. various reef-fish species) provides an example: Clients switch to a new cleaner after defection by their current one if they have access to several cleaners (Bshary and Schäffer 2002). This forces the cleaner to be more cooperative (Bshary and Grutter 2005). Whereas clients with access to only one cleaner make use of punishment by chasing defective cleaners (Bshary and Grutter 2002).

To conclude, we emphasize that the mere option to break the social link to a partner promotes cooperation among humans. Despite being reluctant to pay costs for seeking new partners and reducing link-breaking behaviour, participants cooperated at very high levels. Due to the minuscule network dynamics there was barely any room left for participants to assort and to alter their social environment. Thus, the difference in cooperation levels between the “nearly static” and the static network is even more remarkable. Hence, besides assortment, making relationships costly, and therefore valuable, is an additional way to achieve cooperative outcomes in dynamic networks.

ACKNOWLEDGEMENTS

We thank the students at the University of Göttingen for their participation. Special thanks go to Johannes Pritz and Frederik Nowak for technical support and Xaver Franiel, Sonja Bednarik and Christine Wittge for help during data collection. We thank two anonymous reviewers for helpful comments. The research is funded by the German Initiative of Excellence of the German Science Foundation (DFG).

DATA ACCESSIBILITY

The data used for this article can be accessed online at <http://datadryad.org/>, doi:10.5061/dryad.tp582. We are happy to provide more information on request.

APPENDIX to Chapter I: Electronic Supplementary Material (ESM)

The Prisoner’s Dilemma Game

A common tool to analyse human cooperation is the so-called Prisoner’s Dilemma game. In this simple two-player game, both players are given the choice between cooperation (play C) or defection (play D). They make their decision simultaneously and without communicating with each other, that is decisions are made in private. If both cooperate, they each receive the reward payoff (R). If one defects and the other cooperates, the defector receives the temptation payoff (T) and the cooperator obtains the sucker’s payoff (S). However, if both defect, they each receive the punishment payoff (P). The assumptions $T > R > P > S$ must hold (and if the game is repeated $2R > T + S$):

$$\begin{array}{cc}
 & C & D \\
 \begin{array}{c} C \\ D \end{array} & \begin{pmatrix} R & S \\ T & P \end{pmatrix} &
 \end{array} \quad [1]$$

If the individuals cooperate, both do better than if they had both defected ($R > P$). For a single individual, it is always better to defect no matter what the partner does ($T > R$ and $P > S$). It is clear that under these simple rules, only defection can be evolutionarily stable. In other words, cooperative behaviour is vulnerable to exploitation in one-shot interactions. However, the dilemma can be resolved in repeated interactions by direct reciprocity: “If you help me, I will help you next time”. It has been shown that cooperation can evolve if the probability of another round is high enough and that humans use direct reciprocity to establish cooperation.

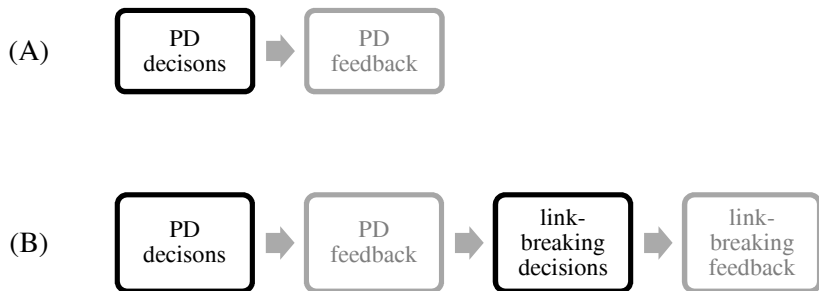




Figure S1. Decision and feedback diagram of one round. Participants played prisoner’s dilemma (PD) games with three partners (STATIC, panel A) or could break social links and receive new partners at random in either dynamic-network treatment. Setting new links was cost-free (DYNAMIC, panel B) or else low costs of 0.10€ (DYNAMIC10, panel C) or high costs of 0.50€ (DYNAMIC50, panel C) arose. For the PD decision, participants were asked “Do you want to play ‘orange’ or ‘blue’ ?” (orange-/blue-buttons, see methods for details). For a link-breaking decision they were asked “Do you want to keep playing with this partner in the next round?” (yes-/no-buttons). If there were open links they were asked for a link-making decision “Do you want to receive a random new partner for costs of 10 [50] cents?” (yes-/no-buttons).

Screenshots: decision making during the experiment

During the experiment participants were confronted with different decisions. In the STATIC treatment participants saw Fig. S3 and Fig. S4 (however, no decisions could be made in the latter case). Additionally, in the DYNAMIC treatment participants saw Fig. S5a and in DYNAMIC10 and DYNAMIC50 treatments they saw Fig S5b.

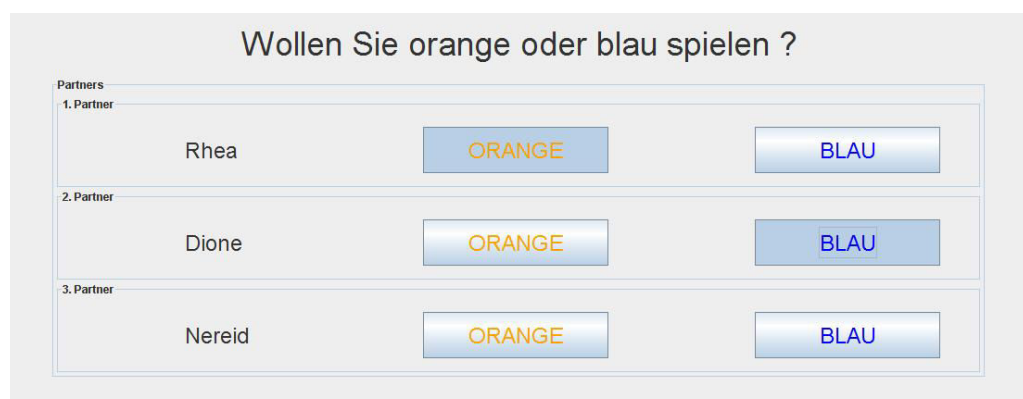


Figure S3. In the prisoner’s dilemma game, participants were asked whether to play “orange” (orange, in this particular case cooperation) or “blau” (blue, defection) and had to make one decision for every linked partner.

Mit welchen Partnern möchten Sie weiterspielen ?

Partners	Sie spielten ...	Ihr Gewinn bzw. Verlust.	Ihr Partner spielte ...	Gewinn bzw. Verlust Ihres Partners	Möchten Sie in der nächsten Runde weiter mit diesem Partner spielen?
1. Partner					
Rhea	orange	0,25	orange	0,25	<input type="button" value="ja"/> <input type="button" value="nein"/>
2. Partner					
Dione	blau	0,40	orange	-0,10	<input type="button" value="ja"/> <input type="button" value="nein"/>
3. Partner					
Nereid	orange	0,25	orange	0,25	<input type="button" value="ja"/> <input type="button" value="nein"/>

Figure S4. Participants were provided with the outcome of the prisoner’s dilemma decisions (the third column indicates the participant’s payoff and the fifth column the partner’s payoff). In the link-breaking stage of the dynamic treatments participants were asked whether to continue playing with a partner and could answer “ja” (yes) or “nein” (no). They had to make one decision for every linked partner.

Ihre Paarungen:

Partners	Sie wollen <u>nicht</u> mehr spielen mit:	Ihr Partner will <u>nicht</u> weiter mit Ihnen zusammenspielen:	
1. Partner			
Rhea			Sie spielen weiterhin zusammen.
2. Partner			
Dione		X	Sie erhalten einen neuen Partner.
3. Partner			
Nereid	X		Sie erhalten einen neuen Partner.

Figure S5a. In the DYNAMIC treatment participants were provided with a summary of the link-breaking decisions. Here, the participant continues to play with Rhea (pseudonym); Dione declined to keep playing with the participant; and in the case of Nereid the participant declined to continue the relationship. Thus, the participant would receive two new, randomly chosen partners.

Ihre Paarungen:

Partners	Sie wollen <u>nicht</u> mehr spielen mit:	Ihr Partner will <u>nicht</u> weiter mit Ihnen zusammenspielen:	Möchten Sie für 50 Cent einen zufälligen neuen Partner bekommen?
1. Partner Leda			Sie spielen weiterhin zusammen. Daher kein Anfordern möglich.
2. Partner Tethys	X		<input type="button" value="ja"/> <input type="button" value="nein"/>
3. Partner Metis		X	<input type="button" value="ja"/> <input type="button" value="nein"/>

Figure S5b. In the DYNAMIC50 (DYNAMIC10) treatment participants were also provided with a summary of the link-breaking decisions. In addition, for each broken link, participants must decide whether they wish to seek a new partner at a cost of 0.50€ (0.10€), i.e. answer “ja” (yes) or “nein” (no).

Table S1. Statistical test details of differences in the link-breaking rate between each of the three dynamic-network treatments where no (DYNAMIC), low (DYNAMIC10), or high costs (DYNAMIC50) arose for receiving new links.

	χ^2 / W	n	df	p
general test: Kruskal-Wallis test				
	21.17	10	2	< 0.001
multiple comparisons with Wilcoxon rank-sum test (3 tests)				
DYNAMIC vs. DYNAMIC10	0	10		< 0.001
DYNAMIC vs. DYNAMIC50	0	10		< 0.001
DYNAMIC10 vs. DYNAMIC50	17	10		< 0.05

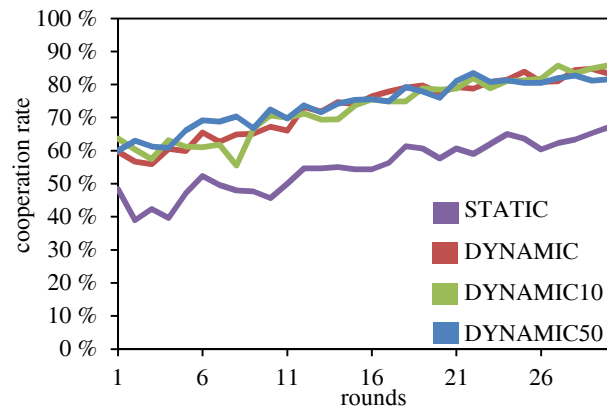


Figure S2. Average cooperation rates over rounds. Participants played prisoner’s dilemma games with three partners. Those partners were either fixed (STATIC) or participants could break social links after each round. In the latter case, setting new links was either cost-free (DYNAMIC), or else low costs of 0.10€ (DYNAMIC10) or high costs of 0.50€ (DYNAMIC50) arose.

Table S2. Statistical test details of differences in the average cooperation level in the prisoner’s dilemma games of the static-network treatment (STATIC) and the three dynamic-network treatments where no (DYNAMIC), low (DYNAMIC10), or high costs (DYNAMIC50) arose for receiving new links.

	χ^2 / W	n	df	p
general test: Kruskal-Wallis test				
	16.23	10	3	< 0.01
multiple comparisons with Wilcoxon rank-sum test (6 tests)				
STATIC vs. DYNAMIC	4	10		< 0.001
STATIC vs. DYNAMIC10	13	10		< 0.05
STATIC vs. DYNAMIC50	5	10		< 0.01
DYNAMIC vs. DYNAMIC10	42	10		0.58
DYNAMIC vs. DYNAMIC50	42	10		0.58
DYNAMIC10 vs. DYNAMIC50	47	10		0.85

Table S3. Statistical test details of differences in the readiness to cooperate when a participant received a newly linked partner in the dynamic-network treatments where no (DYNAMIC), low (DYNAMIC10), or high costs (DYNAMIC50) arose for receiving new links.

	χ^2 / W	n	Df	p
general test: Kruskal-Wallis test				
	12.42	10	2	< 0.01
multiple comparisons with Wilcoxon rank-sum test (3 tests)				
DYNAMIC vs. DYNAMIC10	25	10		0.06
DYNAMIC vs. DYNAMIC50	10	10		< 0.01
DYNAMIC10 vs. DYNAMIC50	31	10		0.16

Table S4. Statistical test details of differences in the willingness of participants to reciprocate defection in the three dynamic-network treatments where no (DYNAMIC), low (DYNAMIC10), or high costs (DYNAMIC50) arose for receiving new links.

	χ^2 / W	n	df	p
general test: Kruskal-Wallis test				
	12.95	10	3	< 0.01
multiple comparisons with Wilcoxon rank-sum test (3 tests)				
DYNAMIC vs. DYNAMIC10	11	10		< 0.05
DYNAMIC vs. DYNAMIC50	31	10		0.17
DYNAMIC10 vs. DYNAMIC50	64	10		0.32

Additional figures of conditional behaviour

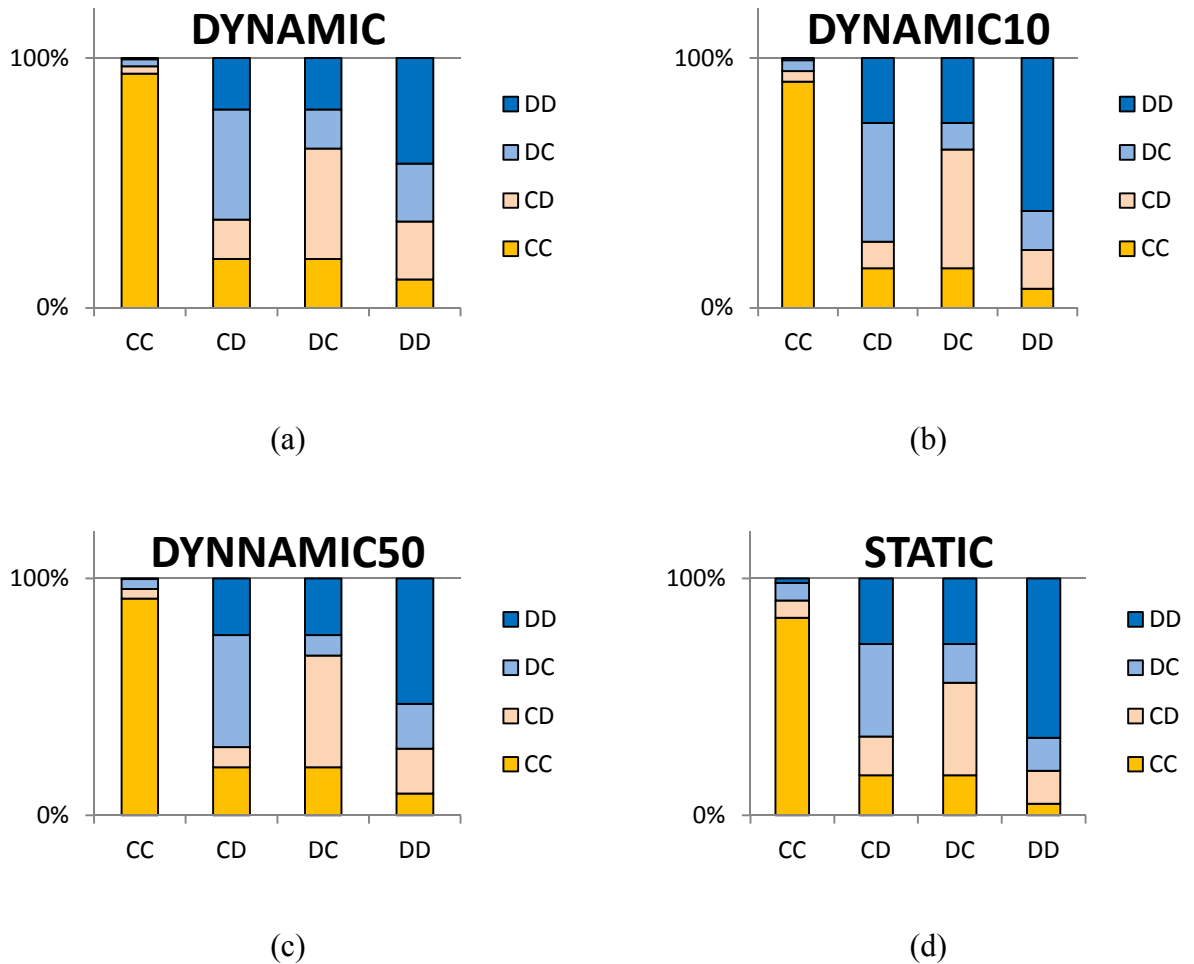
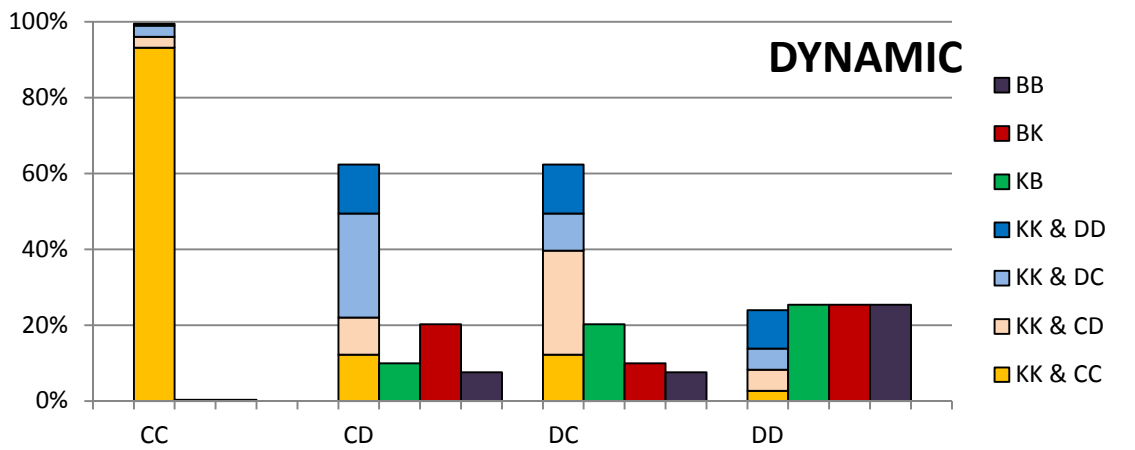




Figure S6. The x-axis shows behaviour of both partners in the previous round, i.e. CC means both players played C, CD means one player played C, the other player D, etc. In case the link between the two players was not broken, the plots show the frequency of actions in the following round. For example, in DYNAMIC treatment (a), if both players played C and kept the link, then in the following round they both played C in 93,65% (orange bar) of the cases. Similarly for DYNAMIC10 (b), DYNAMIC50 (c) and STATIC (d). This picture does not reveal how often links were kept and who of the players decided to break the link (cf. Fig. S7.)



(a)



(b)



(c)

Figure S7. The x-axis shows behaviour of both partners in the previous round. The plots show the frequency of all possible bilateral actions of the players in the following round: BB, if both players decided to break the link, BK or KB, if one player decided to break the link and the other player decided to keep it. If both players decided to keep the

link (KK), their actions in the following prisoner's dilemma game can be DD, DC, CD or CC. Because all links are automatically kept in the STATIC treatment, it is already covered in Figure S6d.

Chapter II: Optimal Decision-Making in the Judge-Advisor System

with Thomas Schultze¹

submitted to Journal of Judgement and Decision Making

¹Courant Research Centre Evolution of Social Behaviour, University of Göttingen,

Kellnerweg 6, 37077 Göttingen, Germany

Abstract

We investigate decision-making in the Judge-Advisor-System where one person, the “judge”, wants to estimate the number of a certain entity and is given advice by another person. The question is how to combine the judge’s initial estimate and that of the advisor in order to get the optimal expected outcome. A previous approach compared two frequently applied strategies, taking the average or choosing the better estimate. In most situations, averaging produced the better estimates. However, this approach neglected a third strategy that judges frequently use, namely a weighted mean of the judges’ initial estimate and the advice. We compare the performance of averaging and choosing to weighting in a theoretical analysis. We find that weighting outperforms both of these strategies when the judge can, without error, detect ability differences between judge and advisor. If we introduce errors in the assessment of the ability differences, the relative performance of weighting compared to averaging or choosing depends on the size of the actual ability differences as well as the magnitude of the error. However, for a wide range of ability differences and errors, weighting is preferably to averaging and more so to choosing. Our analysis expands previous research by showing that weighting is an appropriate advice taking strategy and under which circumstances judges benefit most from applying it.

Keywords: judgment; decision making; advice taking; modelling;

Optimal Decision Making under Uncertainty in the Judge-Advisor System

A famous saying holds that “two heads are better than one”. Accordingly, when making important judgments we rarely do so on our own. Instead, we consult others for advice in the hope that our advisor will provide us with additional insights, expert knowledge or an outside perspective – in short, an independent second opinion. Previous research on advice taking has consistently shown that heeding advice does, in fact, increase the accuracy of judgments (e.g. Sniezek, Schrah et al. 2004; Gino and Schweitzer 2008; Minson, Liberman et al. 2011). However, a commonly observed phenomenon is the suboptimal utilization of advice, that is, judges do not heed the advice as much as they should according to its quality (e.g. Harvey and Fischer 1997; Yaniv and Kleinberger 2000); for reviews see (Yaniv 2004; Bonaccio and Dalal 2006). As a consequence, the de facto improvement in judgment quality observed in many judge-advisor studies is inferior to the improvement that judges could have obtained if they had utilized the advice in the optimal way (Minson and Mueller 2012). The critical question, however, is what constitutes the optimal advice taking strategy. Our main goal is to provide an answer to this question that goes beyond previous research. To this end, we will first discuss the existing approach on the optimal utilization of advice and, then, build on it to arrive at a normative model of advice taking.

Our analysis will build on the logic of the framework commonly used for studying advice taking, the judge-advisor-system (JAS, Sniezek and Buckley 1995). In the JAS, one person (the “judge”) first makes an initial estimate regarding a certain unknown quantity and then receives advice in the form of the estimate another person

(the “advisor”), provided independently. The judge then makes a final, and possibly revised, estimate. Comparison of the initial and final estimates allows to determine the degree to which the judge utilized the advice, and advice utilization is usually expressed as the percent weight of the advice when making the final estimate (e.g. Harvey and Fischer 1997; Yaniv and Kleinberger 2000). The critical question is how strongly the judge should heed the advice in order to come up with the best possible final estimate. So far, our understanding of the optimal degree of advice utilization is rather limited. In situations in which judge and advisor are known to be equally competent or in which comparable expertise is the best assumption – for example when judge and advisor are drawn from the same population and there is not valid information on their relative expertise – the normatively correct strategy is to average the initial estimate and the advice (e.g. Harvey and Fischer 1997; Yaniv and Kleinberger 2000; Soll and Larrick 2009). However, for situations in which there are ability differences between judge and advisor, determining the optimal advice taking strategy is more difficult. To the best of our knowledge, there is only one formal model that addresses the question of optimal advice utilization in the face of ability differences between judge and advisor, the PAR model by Soll and Larrick (Soll and Larrick 2009).

The PAR model of advice taking

The PAR model makes statements about the effectiveness of advice taking strategies based on the three parameters of the JAS, ability differences between judge and advisor (A), the probability of the judge detecting these differences (P), and the degree to which the two judgments contain redundant information (R). Based on these parameters, the PAR model compares two very specific weighting strategies, namely equal weighting (i.e. averaging) and choosing the supposedly more accurate estimate. Averaging is a powerful strategy because it is a statistical truth that the arithmetic mean of the judges' initial estimate and the advice is, on average, equally or more accurate than the initial estimate (Soll and Larrick 2009). If the advisor's estimate is independent from the judge's initial estimate, averaging the initial estimate and the advice results in a reduction of unsystematic and - in some cases - systematic errors (Yaniv 2004; Soll and Larrick 2009).

The averaging strategy performs best if judge and advisor are equally competent. However, in real life, this might rarely be the case. For obvious reasons, averaging two judgments strongly differing in accuracy is unlikely to be optimal. The critical question, then, is how judges should utilize advice when they perceive it to be different from their own initial estimates with regards to accuracy. The PAR model offers an alternative to averaging in the form of the choosing strategy, that is, the judge either maintains the initial estimate or fully adopts the advice, depending on which of the two estimate he or she thinks is more accurate.

The theoretical analysis of the performance of the two advice taking strategies suggests that judges should average their initial estimate and the advice in most of the

cases. That is, even if judge and advisor differ regarding their ability to perform the judgment task, averaging often provides better results than choosing. The exception to this rule are situations in which there are strong and easily identifiable ability differences, and the advantage of choosing increases even more if judge and advisor share a systematic bias. In those cases, judges are usually better off simply choosing the supposedly more accurate estimate.

A possible downside of the PAR model is its focus on only two advice taking strategies. Soll and Larrick (Soll and Larrick 2009) provide strong arguments for this restriction, namely that these strategies are simple to use and that these strategies, averaging and choosing, account for about two thirds of the strategy choices in advice taking. They back up this argument with data from four experiments showing that judges used a choosing strategy in close to 50% of the cases and relied on averaging in about 20% of the cases. However, by implication this means that judges also may have adhered to a third strategy more than 30% of the time, namely weighting¹. In fact, while less frequent than choosing, judges seemed to prefer a weighting strategy to pure averaging based on the results of Soll and Larrick's experiments. A study by Soll and Mannes (Soll and Mannes 2011) showed a similar pattern; depending on the experimental conditions, judges utilized a weighting strategy in about 30 to 40% of the trials.

As previous studies (Soll and Larrick 2009; Soll and Mannes 2011) show, judges seems to engage in three rather than only two strategies when utilizing advice:

¹ In the following, we refer to "weighting" as taking a weighted mean of the two initial estimates, with the weights being based on perceived competence difference. Technically, one could argue that judges apply one strategy which is weighting, and that averaging and choosing are simply specific manifestations of this strategy. While this is certainly true, we will still refer to weighting, choosing, and averaging as three distinct strategies in order to maintain comparability to previous research.

choosing, averaging, and weighting. However, the PAR model only allows us to compare the effectiveness of choosing and averaging. In order to make claims about the appropriateness of weighting, we require a different model that informs us about the optimal weight of advice. Ideally, we want to know for any given constellation of a judge and an advisor who may differ with regards to their judgmental accuracy, how much weight the judge should assign to the advice in order to maximize the accuracy of the final estimates. Importantly, and comparable to the PAR model, these optimal weights need to be of normative character rather than being calculated post-hoc, that is, we need to state - a priori - which weighting scheme has the lowest expected judgmental error. In the following, we will describe a model that – similar to the PAR model – determines the effectiveness of weighted averaging based on ability differences between judge and advisor, as well as the ability of the judge to detect these differences. We will then compare the accuracy of the final estimates that would result from weighting to the expected accuracy of a pure averaging strategy as well as a choosing strategy and test under which conditions weighting is the more appropriate strategy.

2 The Model

2.1 A weighted mean minimizes the error of the final estimates

For the purpose of our model, and in accordance with the basic JAS, we assume that two people, a judge A and an advisor B , are tasked with estimating an unknown quantity (e.g. the distance between two cities). They first provide independent estimates, and then A wants to find the best possible final estimate after receiving B 's estimate as advice. Let us denote A 's a priori estimate by x_A and B 's a priori estimate by x_B . Let us further assume for simplicity that the judgments of both judge and advisor follow a Gaussian distribution centered on the true value x_T with variances σ_A^2 and σ_B^2 . First, we determine the most accurate final estimate in the case that A knows both his or her own error variance and that of the advisor. We compute the most likely estimate \tilde{x} with the Maximum-Likelihood method. Since x_A and x_B are drawn from independent distributions, their joint density function is given by

$$\begin{aligned} f_{AB}(\tilde{x}) &= f_A(\tilde{x})f_B(\tilde{x}) = \left(e^{-\frac{(x_A-\tilde{x})^2}{2\sigma_A^2}} \sqrt{\frac{1}{2\pi\sigma_A^2}} \right) \left(e^{-\frac{(x_B-\tilde{x})^2}{2\sigma_B^2}} \sqrt{\frac{1}{2\pi\sigma_B^2}} \right) \\ &= \frac{e^{-\frac{1}{2}\left(\frac{(x_A-\tilde{x})^2}{\sigma_A^2} + \frac{(x_B-\tilde{x})^2}{\sigma_B^2}\right)}}{2\pi\sigma_A\sigma_B} \end{aligned}$$

To optimize the expression with respect to \tilde{x} we get

$$\frac{d}{d\tilde{x}} (\log(f_{AB}(\tilde{x}))) = -\frac{1}{2} \frac{d}{d\tilde{x}} \left(\frac{(x_A - \tilde{x})^2}{\sigma_A^2} + \frac{(x_B - \tilde{x})^2}{\sigma_B^2} \right)$$

$$= \frac{1}{\sigma_A^2} (x_A - \hat{x})^2 + \frac{1}{\sigma_B^2} (x_B - \hat{x})^2 = 0$$

Solving with respect to \tilde{x} gives

$$\tilde{x} = \frac{x_A \sigma_B^2 + x_B \sigma_A^2}{\sigma_A^2 + \sigma_B^2}$$

which is a weighted average of x_A and x_B . In other words, our formal analysis shows that if the error variance of judge and advisor are known the expected error of the final estimate is minimal if that estimate is a weighted average of the judge's initial estimate and the advice. From the expression above, we can further derive that the optimal weight of the advice is

$$w_B = \frac{\sigma_A^2}{\sigma_A^2 + \sigma_B^2}$$

Under the assumptions stated above, the optimal weight of the advice is the ratio of the advisor's error variance and the sum of the combined error variances of judge and advisor. In other words, this weight represents the proportion of error variance in the combined estimate that is due to the judge's imprecision. Likewise, the weight of the judge's estimate is the ratio of the advisor's error variance and the combined error variances of the two judgments. On a conceptual level, the model indicates that the advice should be weighted more the higher the imprecision of the judge's initial estimate is. For example if the judge's error variance is 1 arbitrary unit and the advisor's error variance is 3 of those units, we can easily compute that the weight that should be placed on the advice is 25%, and if both error variances are equal, the optimal strategy is to weight the advice by 50%.

2.2 Comparing weighting to averaging and choosing

In this section we first compare the optimally weighted average with the arithmetic (non-weighted) average \bar{x} .

$$\bar{x} = \frac{1}{2}(x_A + x_B)$$

First, let us recall that for any random variable X and a real number a , we have

$$\text{Var}(aX) = a^2\text{Var}(X)$$

Further, if X and Y follow independent Gaussian distributions (μ_X, σ_X^2) and (μ_Y, σ_Y^2) , respectively, then their sum $X+Y$ also follows a Gaussian distribution with expected value $\mu_{X+Y} = \mu_X + \mu_Y$ and variance $\sigma_{X+Y}^2 = \sigma_X^2 + \sigma_Y^2$.

Now we look at the distributions of \tilde{x} and \bar{x} . Since they are both linear transformations of x_A and x_B we can directly apply the above two rules. Thus, \tilde{x} and \bar{x} follow a Gaussian distribution with expected value x_T and the variances

$$\sigma_w^2 = \frac{\sigma_A^2 \sigma_B^2}{\sigma_A^2 + \sigma_B^2}$$

$$\sigma_a^2 = \frac{1}{4}(\sigma_A^2 + \sigma_B^2)$$

where σ_w^2 is the error variance of the weighted mean and σ_a^2 is the error variance of the arithmetic mean. As stated above, weighting will perform equal or better than averaging because $\sigma_w^2 \leq \sigma_a^2$ with equality only if $\sigma_A^2 = \sigma_B^2$.

Further, and also in line with our reasoning stated above, we can see that $\sigma_w^2 \leq \min(\sigma_A^2, \sigma_B^2)$. This means that for any quality of advisors (that is, for any σ_B^2), the judge will necessarily provide a more accurate final estimate when taking the weighted mean \tilde{x} instead of choosing either the initial estimate x_A or the advice x_B with equality only if either σ_A^2 or σ_B^2 equals zero. The same does not hold for the arithmetic mean \bar{x} : for example, if the advisor is much worse than the judge (more precisely: if $3\sigma_A^2 \leq \sigma_B^2$), averaging will lead to less accurate final estimates than choosing the initial guess x_A . In other words, unlike pure averaging the optimal weighing strategy is always superior to choosing.

The downside of both weighting and choosing is, obviously, the need for additional information, namely knowledge of the ability difference between judge and advisor. We express these ability differences as the ratio of error variances.

$$m = \frac{\sigma_B^2}{\sigma_A^2}$$

The ratio m expresses how much the judge is more accurate than the advisor. If $m > 1$ the judge is more accurate than the advisor and if $m < 1$ the advisor is more accurate than the judge. We only require this single piece of additional information because we can express the weighted mean \tilde{x} as a function of m :

$$\hat{x} = \frac{m}{1+m} x_A + \frac{1}{1+m} x_B$$

Essentially, the model further yields two intuitive insights: first, as long as the error variance of both the judge and the advisor is nonzero and limited, their judgments

should never be completely ignored, that is, weighting is bound to yield more accurate judgments than choosing the more accurate judgment. Second, the expected error of the weighted average is always smaller or equal to that of the arithmetic mean (they are equal if the optimal weight is 0.5). On a theoretical level perfect weighting is therefore, per definition, superior to the PAR-models choosing and averaging strategies.

2.3 The effect of errors in assessing the ability differences

As we have shown in the previous section, perfect weighting is superior to choosing and equal to or better than pure averaging. However, perfect weighting requires that the ability difference between judge and advisor is known to the judge. Despite judges' ability to differentiate between good and bad advice beyond chance level (e.g. Harvey and Fischer 1997; Harvey, Harries et al. 2000; Yaniv and Kleinberger 2000; Yaniv 2004), exact knowledge of m is rather unlikely. Let us, accordingly, assume that m has to be estimated by the judge and is, therefore, subject to errors or biases. In essence, regardless of whether such a mistake is systematic or not, the judge can either under- or overestimate the true value of m , and we denote the degree to which the judge does so by the factor p . If p equals 1, the judge has a perfect representation of the ability differences. In contrast, values greater than 1 indicate that the judge's perception of the ability erroneously shift in his or her favor, whereas values smaller than 1 mean that the judge overestimates the ability of the advisor. Assuming the judge estimates the competence difference m with a mistake of p , the final result reads as

$$\hat{x} = x_A \frac{pm}{1 + pm} + x_B \frac{1}{1 + pm}$$

And the variance of $\hat{x}(p)$ is given by

$$\sigma_p^2 = \frac{m^2 p^2 \sigma_A^2 + \sigma_B^2}{(1 + pm)^2}$$

In this case, the final estimate obtained by weighting the two initial estimates differently might end up being worse than taking the simple average. This would happen if the competence difference is (i) not very large and (ii) poorly estimated. The

weighted mean might also end up being worse than choosing the better guess. This would happen if the ability difference is actually large, but is perceived as small, but in the correct order. To see the full picture we map the relative improvement

$$r = \frac{\text{error variance of final estimate}}{\text{error variance of initial estimate}}$$

of the judge against the real competence difference m . Values smaller than 1 indicate that the error variance of the final estimates is smaller than that of the initial estimate, that is, the final estimates are more accurate. In contrast, if the final estimates are less accurate than the initial estimates, r will assume values greater than 1. We determined the expected values of r for the three advice taking strategies as a function of the parameters m and p .

For the simple average, we get

$$r_{\text{averaging}}(m) = \frac{\sigma_a^2}{\sigma_A^2} = \frac{1}{4} \frac{\sigma_A^2 + \sigma_B^2}{\sigma_A^2} = \frac{1+m}{4}$$

For weighting, we get

$$r_{\text{weighting}}(m, p) = \frac{\sigma_p^2}{\sigma_A^2} = \frac{m^2 p^2 \sigma_A^2 + \sigma_B^2}{(1+pm)^2 \sigma_A^2} = \frac{m^2 p^2}{(1+pm)^2} + \frac{m}{(1+pm)^2} = \frac{m(1+p^2 m)}{(1+pm)^2}$$

Assuming that the better expert is identified correctly, the ratio r of the variances of initial and final estimate is given by

$$r_{choosing,correct}(m) = \begin{cases} 1, & \text{if } m > 1 \\ m, & \text{if } m < 1 \end{cases}$$

If the wrong expert is chosen, then 1 and m must be exchanged in the formula.

$$r_{choosing,wrong}(m) = \begin{cases} m, & \text{if } m > 1 \\ 1, & \text{if } m < 1 \end{cases}$$

Whether the correct expert is identified or not depends on m and p . If the judge has higher expertise and is overconfident ($m > 1, p > 1$) the choice is correct. Similarly, if the advisor has higher expertise and the judge overestimates the advisor's expertise ($m < 1, p < 1$), the choice will also be correct. In general, the choice is correct if and only if the perceived competence ratio pm and the real competence ratio m are either both greater than 1 or less than 1.

$$r_{choosing}(m, p) = \begin{cases} r_{choosing,correct}(m), & \text{if } (m > 1 \text{ and } pm > 1) \text{ or } (m < 1 \text{ and } pm < 1) \\ r_{choosing,wrong}(m), & \text{else} \end{cases}$$

In order to provide a simple comparison of the three advice taking strategies, we plotted their respective values of r against the ability differences m . We used a LogLog Plot² (Fig.1) to visualize the relevant part of the parameter space. In line with the reasoning above, the left panel of Figure 1 shows that in an ideal environment, in which the judge can correctly assess the ability differences, weighting is bound to outperform both averaging and choosing. However, as the right panel shows, the relative

² A brief explanation for readers unfamiliar with LogLog plots: Since the variables m and r that we wish to plot are relations, we need to scale the axes accordingly. A value of $m = 0.5$ means that the judge is twice as good as the advisor while $m = 2$ means that the advisor is twice as good as the judge. Similarly for $m = 0.1$ and $m = 10$. This means that we need to treat the two intervals $(0; 1)$ and $(1; \infty)$ equally. Further, we must center the plot around 1 instead of 0 because a value of $m = 1$ indicates equal accuracy of judge and advisor. Log(-arithmetic) scaling does both of these things. Double logarithmic scaling (i.e., LogLog Plots) scales both axes logarithmically.

performance of the three strategies changes when we introduce errors in estimating the ability differences. In the example we chose, the judge overestimates his/her ability relative to that of the advisor by 200% (i.e. $p = 3$). In this case, averaging outperforms weighting for small ability differences, and choosing outperforms weighting if the advisor is substantially more accurate than the judge.




Figure 1: Plots of relative improvement of accuracy (i.e., reduction of variance) before and after considering the advisor’s advice using three different methods: Choosing the better estimate (red), averaging both estimates equally (blue), and weighting the estimates according to competence difference (green). In (a) we assume that weighting and choosing are perfectly correct. In (b), the judge overestimates his/her ability relative to that of the advisor by 200% (i.e., $p = 3$), resulting in imperfect weighting and, for some values of m , choosing the wrong estimate. Both axes are in logarithmic scale.

This brief example shows that we need a more detailed analysis to shed light on the question which advice taking strategy performs better in a given situation. To this end, we need to compare the relative improvement in accuracy obtained by the different

strategies as a function of the model parameters p and m . Specifically, we are interested in the relative performance of weighting on one hand and either choosing or averaging on the other (for an in-depth comparison of choosing and averaging, see Soll and Larrick, 2009), which we denote as

$$k_{averaging} = \frac{r_{weighting}}{r_{averaging}} \text{ and } k_{choosing} = \frac{r_{weighting}}{r_{choosing}}, \text{ respectively.}$$

A value of $k = 1$ indicates that weighting and the comparison strategy (averaging or choosing) perform equally well whereas values of $k > 1$ indicate superior performance of weighting and values of $k < 1$ indicate that the respective comparison strategy performs better. The target value k is represented by the color and shade in the contour plot spanned by the parameters m and p (see Fig. 2). The bold line separating the blue and green areas is the iso-accuracy curve which indicates that the accuracy of the weighting strategy equals that of the comparison strategy (i.e. $k = 1$). For each subsequent line in the green area, k increases by 0.1, that is, the weighting-method performs 10% better than the comparison strategy, while in the blue area the opposite is true. In the same fashion we can compare weighting and choosing (Fig.2b).




Figure 2: Contour plot of the relative difference of averaging/weighting (a) and choosing/weighting (b). In the green area, weighting is the better option despite erroneous assessment of the ability differences between judge and advisor, while in the blue area simple averaging (a), or choosing (b), performs better. The contour lines denote increases or decreases in steps of 10%.

As can be seen in the left panel of Figure 2, if there are ability differences between judge and advisor and the judge has a rough representation of these differences, weighting is superior to pure averaging. In contrast, whenever the ability differences are small and/or difficult to detect, judges will benefit more from averaging. The accuracy differences between weighting and choosing are more pronounced (see the left panel of Figure 2). Obviously, the judge must make extreme errors when assessing m in order for choosing to be the better advice taking strategy. In addition, choosing can outperform weighting only if the more accurate person's estimates are chosen. This is the case above the white diagonal (Fig. 2b) for $m > 1$, and below the diagonal for $m < 1$. Note that the second prerequisite creates an asymmetry in the results. This asymmetry is rooted in the fact that choosing is heavily penalized if the judge erroneously chooses the

wrong estimate while weighting is much less prone to such extreme errors because it still assigns some weight to the more accurate judgment.

2.4 Combining Averaging and Choosing

Our analysis so far revealed that weighting is quite a powerful strategy when comparing it to either weighting or choosing. However, one rationale that we can derive from Soll and Larrick's (2009) PAR model is that judges should switch between averaging and choosing in order to maximize the accuracy of their final estimates. Specifically, they should average when ability differences are small and/or difficult to detect and choose when the opposite is true. An interesting vantage point, then, is to compare weighting to a combination of choosing and averaging. Let us, for simplicity, assume that judges know when they should switch from averaging to choosing based on their (potentially biased) perception of m . We can easily compute this threshold by equating $r_{choosing,correct}$ and $r_{averaging}$ which gives $pm = 3$ or $pm = 1/3$. In other words, a perfect application of the combined strategy implies that judges average their initial estimates and the advice until they perceive the initial estimates to be three times as accurate as the advice or vice versa; if this threshold is passed, they choose the more accurate estimate. If m is estimated without error (i.e., $p = 1$), dynamically switching between choosing and averaging is an extremely powerful strategy. However, we have to take into account that if $p \neq 1$, choosing will not always be correct, since the judge may erroneously choose the less accurate judgment. This problem may drastically reduce the performance of the combined strategy, because choosing the wrong expert has highly negative consequences.

In order to compare weighting to the combined strategy of choosing and averaging, we first determine the accuracy gains relative to the initial estimates that would result from a combination of choosing and averaging, $r_{combined}$. Figure 3 (left panel) compares the accuracy ratios of the combined strategy as well as that of

weighting as a function of m and assuming that the judge is strongly overestimating his or her own accuracy ($p = 3$). We next calculated the ratio of the accuracy gain obtained by weighting and that obtained by the combined strategy

$$k_{combined} = \frac{r_{weighting}}{r_{combined}}.$$

The right panel of Figure 3 shows $k_{combined}$ as a function of m and p . The white lines denote the threshold at which judges switch from averaging to choosing based on their perception of the relative accuracy of judge and advisor (i.e. when the product pm is greater than 3 or smaller than $1/3$). The bold lines, again, denote the iso-accuracy-curves. The analysis reveals some interesting findings. First, weighting is superior to the combined strategy in a wide range of situations. Second, the superiority of the weighting strategy is mostly due to the relatively weak performance of choosing. Specifically, application of the combined strategy leads judges to choose in situations in which averaging would outperform weighting but choosing does not, for example when ability differences are small but difficult to assess. Instances where the choosing part of the combined strategy performs better than the weighting strategy occur only for extreme competence differences outside of the parameter range of Figure 3.

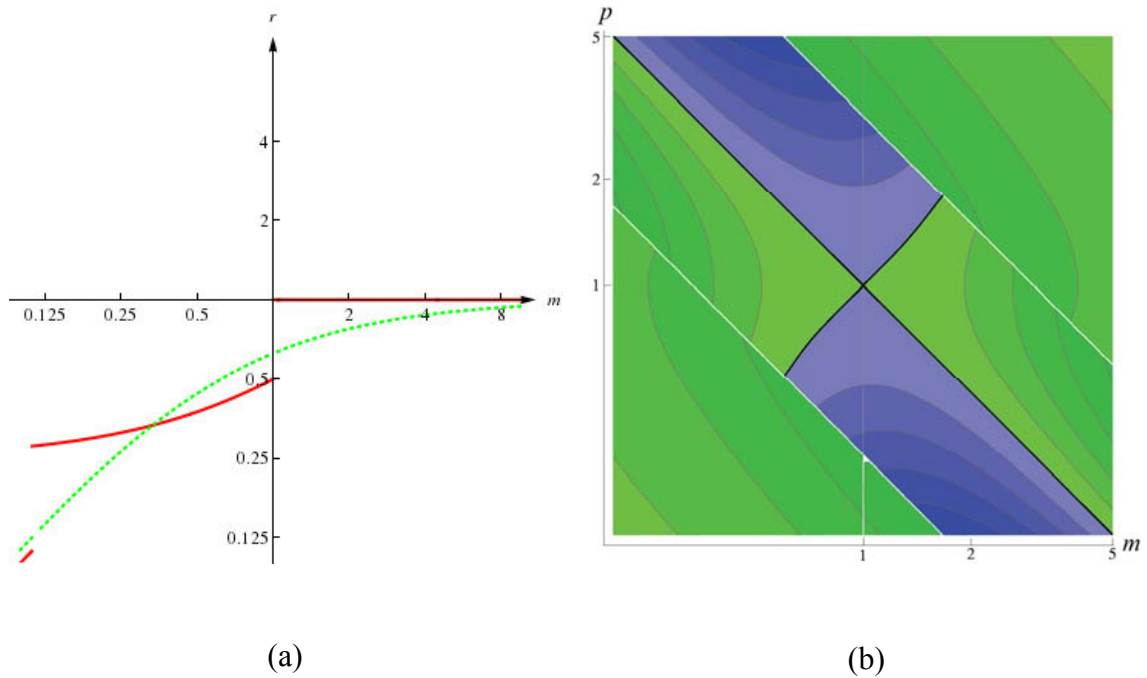


Figure 3: Comparing weighting to the combination of choosing and averaging. In (a) we plot the relative improvement of accuracy (as in Fig.1) of weighting (green) and the combined method (red), both for $p = 3$. Note that imperfect estimation of m leads to choosing the wrong judgment in a specific area. In (b) we generalize the picture by allowing for varying p (as in Fig. 2). In the green area, weighting is the better strategy, while in the blue area the combined method performs better. The contour lines denote increases or decreases in steps of 10%.

3 Discussion

The aim of our theoretical analysis was to answer the question which advice-taking strategy judges in a judge-advisor system should utilize in order to maximize the accuracy of their revised estimates. Previous research has suggested that judges should average their initial estimates and the advice unless the difference in accuracy between the two estimates is large and easily identifiable; in such cases they should simply choose the more accurate estimate (Soll & Larrick, 2009). It is a mathematical fact that averaging two independent and unbiased estimates leads to, on average, more accurate judgments (e.g. Yaniv 2004; Larrick and Soll 2006). However, if the error variance of the two judgments is unequal, there is an optimal weight of advice that produces combined estimates that are always equal or better than simple averaging with regards to accuracy. As a consequence, judges in a judge-advisor system would benefit the most from weighting the advice according to its accuracy relative to that of the judges' initial estimate (Budescu, Fiedler et al. 2006; Budescu and Yu 2006). Similar to choosing the better estimate, the potential superiority of the weighting strategy compared to pure averaging comes at the cost of additional information, namely knowledge of the ability difference between judge and advisor.

If this ability difference is known, a weighting strategy is bound to be superior to both, averaging and choosing. Yet, it is rather unlikely that judges will be able to correctly recognize differences between their own and their advisor's ability with perfect accuracy. Instead, previous research suggests that while judges have some ability to assess the relative quality of advice they frequently underestimate it (e.g. Harvey and Fischer 1997; Harvey, Harries et al. 2000; Yaniv and Kleinberger 2000). In other situations, for example, when judges perceive the task as very difficult (Gino and

Moore 2007) or when they are very anxious, they are prone to overestimate the quality of the advice relative to that of their own initial estimates (Gino, Brooks et al. 2012). If judges' assessment of the ability differences are subject to errors the resulting weighting strategy will result in less accurate judgments, and if these errors become too large, simple averaging turns out to be the better strategy. The fact that the averaging strategy can outperform weighting strategies that are based on erroneous weights has been previously documented in multi-cue judgments (Dawes 1979), and the advantage of averaging increases as the number of cues grows. Hence, the first question we aimed to answer was under which conditions imperfect weighting outperforms averaging. To this end, we compared the expected performance of both strategies as a function of ability differences between judge and advisor as well as the accuracy of the judge when estimating these differences.

Our analysis revealed that imperfect weighting outperforms averaging as long as there are at least moderate ability differences. This performance advantage of the weighting strategy is rather robust against moderate misperceptions of the ability differences. For example, if the judge's error was 50% larger than that of the advisor, weighting is superior to averaging even if the judge under- or overestimates the ability difference by 50%. Additionally, the larger the ability differences become the more robust the weighting strategy becomes against erroneous assessment of these differences. In other words, averaging is likely to produce better estimates than imperfect weighting only when ability differences are small and/or difficult to detect.

We also compared an imperfect weighting strategy to imperfect choosing, finding that the former outperformed the latter with very few exceptions. Specifically, choosing was superior to weighting only when there were large differences in accuracy

which the judge recognized but severely underestimated. The reason for this finding is that the choosing strategy is insensitive to the magnitude of the ability differences whereas the weighting strategy is not. Consider the case where the advisor is much more accurate than the judge but the judge erroneously perceives the advisor to be only slightly better than him- or herself. In this case the judge will still correctly identify the advisor as the expert, and because the actual difference in expertise is large, choosing the advice will produce a rather good result. In contrast, weighting will produce a final estimate that is not too different from (but slightly superior to) the one obtained by averaging because the difference in weights is bound to be small. Based on the misperception of the ability differences, the judge does not assign enough weight to the advice.

Finally, we compared imperfect weighting to a strategy that dynamically switches from averaging to choosing when the (potentially biased) perceived ability differences between judge and advisor become large (Soll and Larrick, 2009). Our analysis revealed that weighting is superior to the combined strategy in a wide range of situations. Interestingly, weighting is better than the combined strategy mainly because the application of the combined strategy leads judges to choose between estimates in situations where averaging would outperform weighting. These situations are characterized by the judge correctly recognizing whether the advisor is more competent than him- or herself or vice versa, but at the same time extremely overestimating the ability differences. The interesting thing about those situations is that simple averaging would have performed better than weighting, but since the ability differences are perceived as too high, the combined strategy must use choosing instead.

Implications and directions for future research

An important implication of our analysis is that weighting is a viable strategy in advice taking. The fact that participants in previous studies adhered to this strategy in a substantial number of trials (Soll and Larrick 2009; Soll and Mannes 2011) as well as its potential superiority to averaging highlight its importance when studying advice taking. Whereas the PAR model suggests that judges should engage in averaging in case of small or difficult to detect ability difference and rely on choosing otherwise, our analysis makes a partially different statement. In case of small and difficult to detect ability differences, averaging is still the best option. However, in case the ability differences become larger and easier to detect, judges should attempt to weight the two judgments by perceived accuracy instead of choosing between the two. Interestingly, weighting the two estimates by their perceived accuracy allows judges to mimic an aggregation strategy that has proven to be very effective if three or more judgments are involved, namely taking the median. Research on group judgment (Bonner & Baumann, 2008; Bonner, Gonzalez, & Sommer, 2004; Bonner, Sillito, & Baumann, 2007) suggests that the way in which groups or judges combine the individual estimates is best described by the median or similar models that discount outliers. The same is true when judges combine several independent judgments (Yaniv, 1997) or receive advice from multiple advisors (Yaniv & Milyavsky, 2007). Importantly, the median strategy outperforms the average because it discounts extreme judgments which are usually less accurate. Naturally, in the JAS with only one advisor, the median is per definition, equal to the mean, but assigning more weight to the more accurate judgment, even if the weight is not optimal due to misperceptions of the ability differences, also leads to discounting the less accurate judgments.

Our theoretical analysis does not only provide a normative framework to compare the expected performance of different advice taking strategies. It also allows to evaluate the effectiveness of judges' advice taking strategies. Similar to (Soll and Larrick 2009) empirical analysis, our model provides performance baselines against which to compare the de facto improvements in accuracy between judges' initial and final estimates. Soll and Larrick's analyses already showed that in the majority of the cases frequent averagers outperformed frequent choosers. An interesting question would, then, be whether or under which conditions frequent weighting can outperform frequent averaging.

Finally, a potential venue for further developing our model would be to include biased judgments. In our theoretical analysis, we made the simplifying assumption that there is no systematic bias in the judge's and advisor's estimates. Incorporating systematic biases of judge and advisor will necessarily make the model more complex, but it may be worthwhile if it allows us to draw conclusions about the relative performance of weighting, choosing and averaging in a wider range of decision situations.

Conclusion

Advice taking is not only an integral part of our daily social reality; it is also one of the most effective ways to increase the quality of our judgments and decisions. In order to make the best use of the wisdom of others, we need a thorough understanding of how well we utilize advice depending on its quality. An elegant way to provide answers to this question is provided by normative models of advice taking. We built on and extended the most prominent normative model of advice taking and, by doing so,

furthered our understanding of how effective different advice taking strategies are in different situations. More importantly, however, normative modelling allows us to detect and, ultimately intervene against, deviations from optimal strategies, that is, they can help us utilize the benefits of advice to its full effect.

Acknowledgements

Discussions with Dirk Semmann and Stefan Schulz-Hardt are gratefully acknowledged. The research is partly funded by the German Initiative of Excellence of the German Science Foundation (DFG).

General Discussion

The evolution of cooperation has become a widely-studied interdisciplinary field of research. Nevertheless, the processes related to cooperative behavior are still not fully understood, especially among humans. Empirical studies are consistently lagging behind models, because they are often difficult to design and time-consuming compared to modelling. In chapter I, I present a study which contributes to the literature on empirical studies of network reciprocity. The key feature of dynamic networks is that individuals can switch partners. Realistically, switching partners requires costs in terms of time or resources. Since such costs have been neglected in most of the present models and in all present experiments this work closes this gap. I showed that costs significantly reduce the dynamism on dynamic networks. But most importantly, cooperative behavior stayed at high levels even with such a reduced dynamism. Essentially, this sheds new light on the mechanism that drives cooperation on dynamic networks. The important difference for cooperation is whether participants have the option to switch partners or not.

Chapter II of the present thesis is concerned with advice taking. In the context of advice taking, I investigated how advice should be weighted optimally. I derived a mathematical model that compared three strategies, *choosing*, *averaging* and *weighting*. If the advisor's competence is sufficiently different from the judge or if the difference is easy to guess, weighting is the best strategy. Averaging turns out to be optimal if the difference is small or hard to guess. Choosing is best only in exceptional and unrealistic parameter regions. My results extend present models by highlighting the benefits of weighting and justify empirical findings where participants seemed to prefer weighting over other strategies.

Currently, there is only minimal scientific exchange between these two research areas. However, in the evolution of cooperation, information plays an important role, particularly in form of reputation. While reputation seems to be especially relevant for humans, other species are able to make use of reputation, too. A recent study shows that cleaner fish and their clients effectively use reputation to maintain cooperative behavior (Bshary and Grutter 2006). In advice taking literature, a robust finding is that participants are inclined to put too much weight on their own opinion (e.g. Harvey and Fischer 1997; Yaniv and Kleinberger 2000). Interestingly, the very opposite has been found in an experiment on the effects of gossip on human cooperation (Sommerfeld, Krambeck et al. 2007). Here, participants received full information on the cooperative behavior of their partners and, additionally, gossip about them. The gossip consisted of short statements of other players who had previously played with that person. Surprisingly, participants adjusted their behavior due to gossip even though it contained no additional information.

A possible explanation of the difference between over-using information from gossip in cooperation experiments and under-using information in advice taking is that humans are sensible to the specific kind of information. In advice taking, typical tasks include guessing difficult quantities, e.g. the average yearly water consumption of Japanese households. Future work could focus on identifying whether different types of information lead to different weighting patterns and, if so, why. On the other hand, experiments on human cooperation could be inspired by the advice taking literature as well. If participants received some information about the behavior of their partners by others and some by own observation, it could be analyzed how they distribute weights on the two sources of information.

Thus, the major outlook for future research is to integrate reputation in social networks, and also take into consideration how humans deal with advice by others. Clearly, future models must give specific and testable predictions. Empiricists, on the other hand should try to understand the underlying mechanisms of models. Understanding and good communication between empiricists and theoreticians are especially relevant in the evolution of cooperation and advice taking, because both topics rely on a balanced progress of theoretical and empirical knowledge.

Summary

Cooperative behaviour is widespread among humans and throughout the animal kingdom. Previous models suggest that the evolution of cooperation can be enhanced by network structure. However, recent experiments were not able to detect cooperation-enhancing capabilities in static networks. Only if the network is dynamic, experiments with humans report increased cooperative behaviour. Since dynamic networks imply the possibility of changing partners, an important aspect is how costs for changing partners affect behaviour. Since this aspect has been neglected so far, chapter I of this work is dedicated to close this gap and explore the effects of costs on dynamic networks. I showed that the willingness to break links is drastically reduced when links to new partners are costly. For very high costs, the rate of breaking links was so low that the network was nearly static. Interestingly, cooperative behaviour stayed at a high level nevertheless. This implies that cooperative behaviour depends, above all, on whether there is an option to switch partners or not. Even if costs are so high that this option is rarely used, cooperation levels are substantially higher than without the option.

Chapter II of this thesis is dedicated to the investigation of decision-making. In the so-called Judge-Advisor-System, one person, the *judge*, estimates an unknown quantity. Then, the judge receives advice from another person, the *advisor*. Importantly, the estimates by the judge and the advisor are made independently. The task is to find out how the judge should best use the information from the advisor. Existing approaches mainly focused on two methods, (i) taking the average, and (ii) choosing one of the initial estimates. This simplistic approach is mainly driven by empirical data, where it seems that in some experiments over 70% of participants used one of these methods.

However, other weights are also frequently assigned and a thorough theoretical investigation of optimal weights is necessary. Therefore, I derived a normative model that tells under which circumstances it is better to (i) take the average, to (ii) choose what you think is the better estimate or to (iii) try to assign proper weights. Which of the three is the best depends on the difference in expertise of judge and advisor as well as on the judge's likelihood to know this difference. If the judge has a good representation of this difference, assigning weights is always the best bet. The simple average is useful if the difference in expertise is small or difficult to guess. Finally, choosing performs well if the difference is large but its amount is difficult to guess. Motivated by previous approaches, I also explored the performance of a combination of choosing and averaging, i.e. a method that uses averaging for small difference in expertise and choosing for a large difference. Surprisingly, the performance of this combined method was very poor. The main reason is the uncertainty in guessing the difference of expertise. Therefore, assigning proper weights is almost always better than using the combined method. Since choosing the worse expert has performs so poorly, the combined method requires that the risk of choosing the wrong person is low. But this means that the difference in expertise is easy to guess and therefore weighting is the best method. Over all, I showed that weighting is a viable method for a wide range of parameters.

Zusammenfassung

Kooperatives Verhalten lässt sich in vielen Bereichen menschlichen Zusammenlebens sowie im gesamten Tierreich beobachten. In evolutionären Modellen wurde gezeigt, dass Netzwerkstrukturen die Kooperation erhöhen können. Empirische Studien versuchten vergeblich diesen Mechanismus auch bei Menschen nachzuweisen. Es scheint, als würden Netzwerke nur dann die Kooperation erhöhen, wenn die Strukturen nicht statisch sind, sondern dynamisch. Das heißt, dass die Individuen die Möglichkeit haben, ihre Partner zu wechseln. Eine wichtige – aber bislang unerforschte – Eigenschaft dynamischer Netzwerke ist jedoch, dass derartige Wechsel von Partnern in der Regel Kosten verursachen, ob in Form von Zeit oder Ressourcen. Kapitel I meiner Arbeit schließt diese Lücke, in dem es sich mit den Effekten von Kosten auf dynamischen Netzwerken befasst. Ich konnte nachweisen, dass Menschen seltener Interaktionen mit Partnern beendeten, wenn die Kontaktaufnahme mit einem neuen Partner mit Kosten verbunden war. Bei sehr hohen Kosten, wurden Partner so selten gewechselt, dass das Netzwerk fast statisch war. Interessanterweise blieb die Kooperation dennoch sehr hoch. Das bedeutet, dass für kooperatives Verhalten entscheidend ist, ob man die Möglichkeit hat, Partner zu wechseln. Im Gegensatz zu bisherigen Annahmen ist es daher nicht wichtig, wie oft tatsächlich Partner gewechselt werden, sondern lediglich ob es die Möglichkeit dazu gibt.

In Kapitel II beschäftige ich mich mit optimalem Entscheidungsverhalten. Im sogenannten Judge-Advisor-System geht es darum, dass eine Person, der *Judge*, eine unbekannte numerische Größe schätzen will. Dazu erhält der Judge eine zweite unabhängige Schätzung als Rat von einer zweiten Person, des *Advisor*. Schließlich ist die Frage, wie der Judge optimal den Rat verwerten kann um seine Anfangsschätzung

zu verbessern. Bisherige Forschung konzentrierte sich hauptsächlich auf zwei mögliche Methoden, (i) das Bilden des Mittelwerts und (ii) das Wählen der besseren Anfangsschätzung. Das Hauptargument für diese einfachen Methoden ist deren häufige Verwendung in bisherigen Experimenten. Allerdings wurden sehr wohl auch andere Gewichtungen beobachtet und daher ist eine gründliche Analyse der optimalen Gewichtung erforderlich. In der vorliegenden Arbeit leitete ich ein normatives Modell her, das beschreibt, unter welchen Bedingungen welche Methode das bestmögliche Ergebnis liefert. Es wurden drei Methoden verglichen: (i) das Bilden des Durchschnitts, (ii) das Wählen der besseren Anfangsschätzung, und (iii) das Bilden eines gewichtetet Mittelwerts, wobei das Gewicht vom Kompetenzunterschied abhängt. Welche Methode optimal ist, hängt davon ab, wie groß der Kompetenzunterschied ist und wie gut er vom Judge erkannt wird. Die Durchschnittsbildung ist immer dann vorteilhaft, wenn der Kompetenzunterschied nicht groß ist, oder nur schwer richtig eingeschätzt werden kann. Wenig überraschend lohnt sich das Wählen der besseren Anfangsschätzung, wenn der Kompetenzunterschied hinreichend groß ist, vorausgesetzt es wird tatsächlich die bessere Anfangsschätzung gewählt. Wenn der Kompetenzunterschied vom Judge gut eingeschätzt werden kann, ist eine Entsprechende Gewichtung immer die beste Methode, unabhängig vom tatsächlichen Unterschied. In Übereinstimmung mit bisheriger Forschung wurde auch die Kombination von Durchschnittsbildung und Wählen der besseren Anfangsschätzung untersucht. Diese Kombinationsmethode beruht darauf, bei als gering eingeschätztem Kompetenzunterschied den Durchschnitt zu bilden und ansonsten die bessere Anfangsschätzung zu wählen. Interessanterweise schneidet diese Kombinationsmethode sehr schlecht ab, was hauptsächlich daran liegt, dass zu oft

die falsche Anfangsschätzung genommen würde. Insgesamt ist das gewichtete Mittel also eine geeignete Methode für einen großen Parameterbereich.

References

- Axelrod, R. and W. D. Hamilton (1981). "The evolution of cooperation." Science **211**(4489): 1390-1396.
- Axelrod, R. and W. D. Hamilton (1981). "The evolution of cooperation." Science **211**: 1390-1396.
- Bonaccio, S. and R. S. Dalal (2006). "Advice taking and decision-making: An integrative literature review, and implications for the organizational sciences." Organizational Behavior and Human Decision Processes **101**(2): 127-151.
- Bshary, R. and A. S. Grutter (2002). "Asymmetric cheating opportunities and partner control in a cleaner fish mutualism." Animal Behaviour **63**: 547-555.
- Bshary, R. and A. S. Grutter (2005). "Punishment and partner switching cause cooperative behaviour in a cleaning mutualism." Biology Letters **1**: 396-399.
- Bshary, R. and A. S. Grutter (2006). "Image scoring and cooperation in a cleaner fish mutualism." Nature **441**(7096): 975-978.
- Bshary, R. and D. Schäffer (2002). "Choosy reef fish select cleaner fish that provide high-quality service." Animal Behaviour **63**: 557-564.
- Budescu, D. V., K. Fiedler, et al. (2006). "Confidence in aggregation of opinions from multiple sources." Information sampling and adaptive cognition: 327-354.
- Budescu, D. V. and H.-T. Yu (2006). "To Bayes or not to Bayes? A comparison of two classes of models of information aggregation." Decision analysis **3**(3): 145-162.
- Cassar, A. (2007). "Coordination and cooperation in local, random and small world networks: Experimental evidence." Games and Economic Behavior **58**: 209-230.
- Cheney, D. L., L. R. Moscovice, et al. (2010). "Contingent cooperation between wild female baboons." Proceedings of the National Academy of Sciences **107**(21): 9562-9566.
- Cinyabuguma, M., T. Page, et al. (2005). "Cooperation under the threat of expulsion in a public goods experiment." Journal of Public Economics **89**(8): 1421-1435.

- Dawes, R. M. (1979). "The robust beauty of improper linear models in decision making." American psychologist **34**(7): 571.
- Du, F. and F. Fu (2011). "Partner selection shapes the strategic and topological evolution of cooperation: The power of reputation transitivity." Dynamic Games and Applications **1**: 354-369.
- Fehl, K., R. D. Sommerfeld, et al. (2012). "I dare you to punish me - vendettas in games of cooperation." PLoS ONE **7**(9): e45093.
- Fehl, K., D. J. van der Post, et al. (2011). "Co-evolution of behaviour and social network structure promotes human cooperation." Ecology Letters **14**(6): 546-551.
- Fehr, E. and S. Gächter (2002). "Altruistic punishment in humans." Nature **415**(6868): 137-140.
- Fowler, J. H. and N. A. Christakis (2010). "Cooperative behavior cascades in human social networks." Proceedings of the National Academy of Sciences of the United States of America **107**(12): 5334-5338.
- Fu, F., C. Hauert, et al. (2008). "Reputation-based partner choice promotes cooperation in social networks." Physical Review E **78**(2): 026117.
- Fu, F., T. Wu, et al. (2009). "Partner switching stabilizes cooperation in coevolutionary prisoner's dilemma." Physical Review E **79**(3): 036101.
- Gino, F., A. W. Brooks, et al. (2012). "Anxiety, advice, and the ability to discern: feeling anxious motivates individuals to seek and use advice." Journal of personality and social psychology **102**(3): 497.
- Gino, F. and D. A. Moore (2007). "Effects of task difficulty on use of advice." Journal of Behavioral Decision Making **20**(1): 21-35.
- Gino, F. and M. E. Schweitzer (2008). "Blinded by anger or feeling the love: how emotions influence advice taking." Journal of Applied Psychology **93**(5): 1165.
- Gracia-Lázaro, C., J. A. Cuesta, et al. (2012). "Human behavior in Prisoner's Dilemma experiments suppresses network reciprocity." Scientific reports **2**.

- Gracia-Lázaro, C., A. Ferrer, et al. (2012). "Heterogeneous networks do not promote cooperation when humans play a prisoner's dilemma." Proceedings of the National Academy of Sciences of the United States of America **109**(32): 12922-12926.
- Greiner, B. (2004). "The online recruitment system ORSEE 2.0 - A guide for the organization of experiments in economics." *Working Paper Series in Economics*, 10. University of Cologne: Cologne, Germany.
- Grujić, J., C. Fosco, et al. (2010). "Social experiments in the mesoscale: Humans playing a spatial prisoner's dilemma." PLoS ONE **5**(11): e13749.
- Grujić, J., C. Gracia-Lázaro, et al. (2014). "A comparative analysis of spatial Prisoner's Dilemma experiments: Conditional cooperation and payoff irrelevance." Scientific reports **4**.
- Grujić, J., T. Rohl, et al. (2012). "Consistent strategy updating in spatial and non-spatial behavioral experiments does not promote cooperation in social networks." PLoS ONE **7**(11): e47718.
- Hamilton, W. D. (1964). "The genetical evolution of social behaviour. I." Journal of Theoretical Biology **7**(1): 1-16.
- Harvey, N. and I. Fischer (1997). "Taking advice: Accepting help, improving judgment, and sharing responsibility." Organizational Behavior and Human Decision Processes **70**(2): 117-133.
- Harvey, N., C. Harries, et al. (2000). "Using advice and assessing its quality." Organizational Behavior and Human Decision Processes **81**(2): 252-273.
- Hauert, C. and M. Doebeli (2004). "Spatial structure often inhibits the evolution of cooperation in the snowdrift game." Nature **428**: 643-646.
- Hauert, C. and M. Doebeli (2004). "Spatial structure often inhibits the evolution of cooperation in the snowdrift game." Nature **428**(6983): 643-646.
- Henzi, S. P., D. Lusseau, et al. (2009). "Cyclicity in the structure of female baboon social networks." Behavioral Ecology and Sociobiology **63**: 1015-1021.

- Herrmann, B., C. Thöni, et al. (2008). "Antisocial punishment across societies." Science **319**: 1362-1367.
- Holton, D. A. and J. Sheehan (1993). The Petersen Graph Cambridge, Cambridge University Press.
- Jordan, J., D. Jordan, et al. (2013). "Contagion of Cooperation in Static and Fluid Social Networks." PLoS ONE **8**(6): e66199.
- Kirchkamp, O. and R. Nagel (2007). "Naive learning and cooperation in network experiments." Games and Economic Behavior **58**: 269-292.
- Kossinets, G. and D. J. Watts (2006). "Empirical analysis of an evolving social network." Science **311**: 88-90.
- Larrick, R. P. and J. B. Soll (2006). "Intuitions about combining opinions: Misappreciation of the averaging principle." Management science **52**(1): 111-127.
- Li, Y., Y. Min, et al. (2013). "Partner switching promotes cooperation among myopic agents on a geographical plane." Physical Review E **87**(2): 022823.
- Lieberman, E., C. Hauert, et al. (2005). "Evolutionary dynamics on graphs." Nature **433**: 312-316.
- Ligon, J. D. (1983). "Cooperation and reciprocity in avian social systems." American Naturalist: 366-384.
- Liu, Y., X. Chen, et al. (2012). "Does migration cost influence cooperation among success-driven individuals?" Chaos, Solitons & Fractals **45**(11): 1301-1308.
- Lusseau, D. and M. E. J. Newman (2004). "Identifying the role that animals play in their social networks." Proceedings of the Royal Society B: Biological Sciences **271**: S477-S481.
- Maier-Rigaud, F. P., P. Martinsson, et al. (2010). "Ostracism and the provision of a public good: experimental evidence." Journal of Economic Behavior & Organization **73**(3): 387-395.

- Masuda, N. (2007). "Participation costs dismiss the advantage of heterogeneous networks in evolution of cooperation." Proceedings of the Royal Society B: Biological Sciences **274**: 1815-1821.
- Melis, A. P. and D. Semmann (2010). "How is human cooperation different?" Philosophical Transactions of the Royal Society B: Biological Sciences **365**(1553): 2663-2674.
- Minson, J. A., V. Liberman, et al. (2011). "Two to tango: Effects of collaboration and disagreement on dyadic judgment." Personality and Social Psychology Bulletin: 0146167211410436.
- Minson, J. A. and J. S. Mueller (2012). "The Cost of Collaboration Why Joint Decision Making Exacerbates Rejection of Outside Information." Psychological science **23**(3): 219-224.
- Nash Jr, J. F. (1950). "The bargaining problem." Econometrica: Journal of the Econometric Society: 155-162.
- Nowak, M. A. (2006). "Five rules for the evolution of cooperation." Science **314**(5805): 1560-1563.
- Nowak, M. A. and R. M. May (1992). "Evolutionary games and spatial chaos." Nature **359**: 826-829.
- Nowak, M. A. and K. Sigmund (1992). "Tit-for-tat in heterogeneous populations." Nature **355**: 250-253.
- Nowak, M. A. and K. Sigmund (1993). "A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game." Nature **364**: 56-58.
- Nowak, M. A. and K. Sigmund (1998). "Evolution of indirect reciprocity by image scoring." Nature **393**(6685): 573-577.
- Nowak, M. A., C. E. Tarnita, et al. (2010). "The evolution of eusociality." Nature **466**(7310): 1057-1062.
- Ohtsuki, H., C. Hauert, et al. (2006). "A simple rule for the evolution of cooperation on graphs and social networks." Nature **441**: 502-505.

- Pacheco, J. M., A. Traulsen, et al. (2006). "Active linking in evolutionary games." Journal of Theoretical Biology **243**(3): 437-443.
- Pacheco, J. M., A. Traulsen, et al. (2008). "Repeated games and direct reciprocity under active linking." Journal of Theoretical Biology **250**(4): 723-731.
- Perc, M. and A. Szolnoki (2010). "Coevolutionary games—a mini review." BioSystems **99**(2): 109-125.
- Poncela, J., J. Gómez-Gardeñes, et al. (2011). "Cooperation in scale-free networks with limited associative capacities." Physical Review E **83**(5): 057101.
- R Core Team (2013). R: A language and environment for statistical computing. Vienna, Austria, R Foundation for Statistical Computing.
- Rand, D. G., S. Arbesman, et al. (2011). "Dynamic social networks promote cooperation in experiments with humans." Proceedings of the National Academy of Sciences of the United States of America **108**(48): 19193-19198.
- Rapoport, A. and A. M. Chammah (1965). Prisoner's dilemma: A study in conflict and cooperation. Ann Arbor, University of Michigan Press.
- Santorelli, L. A., C. R. Thompson, et al. (2008). "Facultative cheater mutants reveal the genetic complexity of cooperation in social amoebae." Nature **451**(7182): 1107-1110.
- Santos, F. C., J. M. Pacheco, et al. (2006). "Cooperation prevails when individuals adjust their social ties." PLoS Computational Biology **2**: 1284-1291.
- Saramäki, J., E. A. Leicht, et al. (2014). "Persistence of social signatures in human communication." Proceedings of the National Academy of Sciences of the United States of America **111**(3): 942-947.
- Semmann, D., H. J. Krambeck, et al. (2005). "Reputation is valuable within and outside one's own social group." Behavioral Ecology and Sociobiology **57**(6): 611-616.
- Shirado, H., F. Fu, et al. (2013). "Quality versus quantity of social ties in experimental cooperative networks." Nature Communications **4**: 2814.

- Smith, J. M. and G. Price (1973). "The Logic of Animal Conflict." Nature **246**: 15.
- Sniezek, J. A. and T. Buckley (1995). "Cueing and cognitive conflict in judge-advisor decision making." Organizational Behavior and Human Decision Processes **62**(2): 159-174.
- Sniezek, J. A., G. E. Schrah, et al. (2004). "Improving judgement with prepaid expert advice." Journal of Behavioral Decision Making **17**(3): 173-190.
- Soll, J. B. and R. P. Larrick (2009). "Strategies for revising judgment: How (and how well) people use others' opinions." Journal of Experimental Psychology: Learning, Memory, and Cognition **35**(3): 780.
- Soll, J. B. and A. E. Mannes (2011). "Judgmental aggregation strategies depend on whether the self is involved." International Journal of Forecasting **27**(1): 81-102.
- Sommerfeld, R. D., H.-J. Krambeck, et al. (2007). "Gossip as an alternative for direct observation in games of indirect reciprocity." Proceedings of the National Academy of Sciences **104**(44): 17435-17440.
- Sugden, R. (1986). The economics of rights, co-operation and welfare, Blackwell Oxford.
- Szolnoki, A. and M. Perc (2009). "Resolving social dilemmas on evolving random networks." EPL (Europhysics Letters) **86**(3): 30007.
- Szolnoki, A., M. Perc, et al. (2008). "Making new connections towards cooperation in the prisoner's dilemma game." EPL (Europhysics Letters) **84**(5): 50007.
- Szolnoki, A., M. Perc, et al. (2008). "Towards effective payoffs in the prisoner's dilemma game on scale-free networks." Physica A **387**: 2075-2082.
- Traulsen, A. and M. A. Nowak (2006). "Evolution of cooperation by multilevel selection." Proceedings of the National Academy of Sciences **103**(29): 10952-10955.
- Traulsen, A., D. Semmann, et al. (2010). "Human strategy updating in evolutionary games." Proceedings of the National Academy of Sciences of the United States of America **107**(7): 2962-2966.

- Trivers, R. (1971). "The evolution of reciprocal altruism." Quarterly Review of Biology **46**: 35-57.
- Trivers, R. L. (1971). "The evolution of reciprocal altruism." Quarterly Review of Biology: 35-57.
- Turner, P. E. and L. Chao (2003). "Escape from prisoner's dilemma in RNA phage $\phi 6$." The American Naturalist **161**(3): 497-505.
- Wang, J., S. Suri, et al. (2012). "Cooperation and assortativity with dynamic partner updating." Proceedings of the National Academy of Sciences of the United States of America **109**(36): 14363-14368.
- Wang, Z., L. Wang, et al. (2012). "Inferring reputation promotes the evolution of cooperation in spatial social dilemma games." PLoS ONE **7**(7): e40218.
- Wu, B., D. Zhou, et al. (2010). "Evolution of cooperation on stochastic dynamical networks." PLoS ONE **5**: e11187.
- Yaniv, I. (2004). "The benefit of additional opinions." Current directions in psychological science **13**(2): 75-78.
- Yaniv, I. and E. Kleinberger (2000). "Advice taking in decision making: Egocentric discounting and reputation formation." Organizational Behavior and Human Decision Processes **83**(2): 260-281.
- Zermelo, E. (1913). Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels. Proceedings of the fifth international congress of mathematicians, II, Cambridge UP, Cambridge.

Acknowledgements

First I want to thank Dirk Semmann for excellent supervision. The atmosphere in Dirk's work group was always very good and everybody was invited to join the many constructive discussions. I also thank Katrin Fehl for patiently supporting me with her expertise. Open-minded discussions, easy communication and excellent collaboration with Thomas Schultze are gratefully acknowledged.

Further, I thank everybody who helped our group: Johannes Pritz, Christine Wittge, Frederic Nowak, Xaver Franiel and Rene Lange.

Special thanks go to Karl Sigmund who filled me with enthusiasm for the topic in the first place and supported me in getting this position.

Most of all, I want to thank my family, especially my wife Sonja and my sons Michael and Daniel, for their support and understanding.

Finally, I want to thank every member of the Courant Research Center Evolution of Social Behavior for interesting talks, fruitful discussions and a great experience that immensely broadened my horizon.

The funding was provided by the German Initiative of Excellence of the German Science Foundation (DFG).

Curriculum Vitae

Peter Bednarik

Born February 11th 1984 in Bratislava, Slovakia. Married, 2 children.

Education and Work

Since 08/2011 PhD-student, Courant Research Center Evolution of Social Behavior,
University of Göttingen

Thesis: “Theoretical and empirical analysis of the evolution of
cooperation”

2011 - 2012: Teaching assistant for Game Theory for Economists at
Microeconomic Department, University of Goettingen

07/2011 Diploma (Mag.rer.nat.) in Mathematics
University of Vienna

Thesis: “Discretized Best-Response dynamics for cyclic games”

2006-2011 Studies of Mathematics
University of Vienna, Austria and University of Leeds, UK

2010-2011 Teaching assistant for Mathematics for Biotechnologists at
University of Natural Resources and Life Sciences, Vienna

2002-2006 Studies of Electrical Engineering,
Vienna University of Technology, Austria

2002-2003 Shell Austria AG, Assistant in Corporate Quality Management

2002 Matura (A-level), Theresiansiche Akademie Wien, Austria

Grants/Scholarships

- 07/2013-09/2013 Research Interships in Science and Engineering (RISE),
funded by German Academic Exchange Service (DAAD) and the
Drexel University (USA), Funding amount: €2100
- 06/2012-08/2012 Research Interships in Science and Engineering (RISE),
funded by German Academic Exchange Service (DAAD),
funding amount: €1700
- 06/2008 Leistungsstipendium of the University of Vienna for outstanding
achievements (Performance Scholarship funded by private
institutions), Funding amount: €360

Publications

Bednarik, P., Fehl, K., Semmann, D. (2014). Costs for switching partners reduce network dynamics but not cooperative behavior. *Proc. R. Soc. B* **281**: 20141661

Zhang, B., Li, C., De Silva, H., Bednarik, P., & Sigmund, K. (2014). The evolution of sanctioning institutions: An experimental approach to the social contract. *Experimental Economics*, **17**(2), 285-303.

Erklärung über eigene Leistungen

Hiermit erkläre ich, dass ich diese Dissertation selbständig und ohne unzulässige Hilfe Dritter und ohne Benutzung anderer als der angegebenen Quellen und Hilfsmittel verfasst habe. Alle den benutzten Quellen wörtlich oder sinngemäß entnommenen Stellen sind als solche einzeln kenntlich gemacht.

Das gesamte Promotionsvorhaben wurde gemeinsam von Dirk Semmann und Peter Bednarik geplant. Die Experimente wurden von Peter Bednarik und Christine Wittge durchgeführt. Die Datenauswertung erfolgte durch Peter Bednarik. Das Kapitel I wurde von Peter Bednarik verfasst. Das Kapitel II wurde von Peter Bednarik entworfen und berechnet, sowie von Peter Bednarik und Thomas Schultze geschrieben. Alle Koautoren haben zu der finalen Version der Manuskripte beigetragen.

Des Weiteren erkläre ich, dass ich nicht anderweitig eine entsprechende Promotion beantragt habe und dass diese Arbeit nicht in gleicher oder ähnlicher Form für eine andere Prüfung vorgelegt wurde.

Göttingen, August 2014