Project Description – Project Proposals

Dr. Hannes Rau, Department of Economics, Heidelberg University
(DFG-Erstantrag)

Preference-based games and Nash equilibrium play

Project summary:

The Nash equilibrium is a widely used concept in applied game theory and experimental economics. However, in several situations it fails to accurately predict behavior. Sometimes players even violate the criterion of strict dominance. As this happens in very simple games as well, it seems unlikely that a lack of rationality is the driving force behind behavior. In many applications, the Nash prediction is based solely on players' own (material) payoffs. These are assumed to represent the utility of the decision makers. However, it is a well-established fact that for many players their utility does not only depend on their own payoff. Instead, they frequently take additional factors into account, such as the payoffs of others or the type of interaction ("social preferences"). As a result, the actual strategic situation may be different from the one implied by the material payoffs of the game. People even might play a completely different kind of game compared to what the analysts think they do.

In this project, we intend to explore the latter aspect by conducting a comprehensive analysis how the nature of a strategic situation changes when incorporating players' (social) preferences into the decision problem. We plan to first examine simple one-shot 2x2 games and subsequently extend the analysis to more complex categories of games. We plan to use experimental methods to elicit people's evaluations of the outcomes of games, which in our context usually correspond to payoff tuples. We first start with an ordinal preference ranking and later on extend it to a cardinal measure. These evaluations can then be used to infer the strategic properties of the game the agents are playing according to their preferences (the "preference-based game"). Empirical frequencies of different types of preferences yield information how often specific pairings of types occur. From that may determine the likelihood by which games transform into other classes of games when accounting for the social preferences of the involved players. Because the Nash-prediction that is based solely on players' own payoffs is biased towards selfish behavior, through this approach more cooperative outcomes can be justified as an equilibrium.

Furthermore, we are able to test if the equilibrium prediction can substantially be improved when taking players' preferences into account. This can be done by comparing empirical frequencies of Nash equilibrium play, once based on the predictions from the original payoff structure with those from the corresponding preference-based games. If the latter frequencies are significantly higher, then it is not the concept of the Nash equilibrium itself to be blamed for the imprecise prediction of behavior. Rather the material payoff structure of the game (the "game-protocol") does not represent the strategic situation adequately.

1 State of the art and preliminary work

One prominent example of the violation of strict dominance is behavior in Social Dilemma games, such as the Prisoner's Dilemma game (PD) or the Public Good game (PG). Cooperation in a PD game or any positive contribution in a Public Good game is a strictly dominated strategy when the number of rounds is finite. Hence, it can never be part of a (subgame-perfect) Nash equilibrium strategy. However, experimental studies show that average rates of cooperation in Social Dilemma Games are typically between 30% to 40%, depending on the specific characteristics of the game (see e.g. Sally, 1995 for a meta-analysis). Further evidence for players not acting according to the Nash-equilibrium prediction can also be found in sequential games: For example, one frequently observes people rejecting positive offers in Ultimatum Games and returning positive amounts in Trust Games.[1] Clearly, the behavior of the second-movers in these games does not maximize their own payoff and thus violates the Nash-criterion (if it is based on own monetary payoffs, as it is often the case when implementing these kind of games in experiments). In repeated interactions, one could plausibly explain this behavior in terms of reputation or reciprocal motives. However, when considering only one-shot interactions or last-round behavior of these games, this explanation cannot apply. Nevertheless, one often observes positive rates of cooperation/ contributions even there.[2] Putting this together, suggests that a substantial proportion of players does not only seek to maximize their own payoff.[3] As these decision situations are rather simple, it is reasonable to suspect that such choices are intentional. Hence, we may conclude that some additional factors must influence their decisions, which are not captured by the players' own material payoffs.

Typically, game and decision theorists assume that payoffs should be interpreted as *complete descriptions* of all possible factors that may motivate the agents. This approach is appropriate from a theoretical perspective. However, in an applied context it is only possible to observe the material outcomes of players, but not how they evaluate the outcomes in terms of utility. Nevertheless, these outcomes often are interpreted as corresponding to the players' "payoffs" of the game. This is not necessarily true, as several other relevant factors might not be reflected in these values. As a result, it may lead to wrong conclusions about the underlying motivations of the players and wrong predictions about behavior.
A similar challenge applies, when we want to implement a given strategic situation in a lab or field experiment. There, one has to decide on a specific *game-protocol* (a specific reward scheme), which is intended to induce the strategic properties of the situation of interest. According to Weibull (2004), the game protocol describes *"a game-form[4] with specified material consequences"*; while for obtaining the actual game one in addition needs to incorporate players' preferences over these material consequences. The latter is what we call a *preference-based game*. This distinction reflects the idea that it is not guaranteed that the material payoff structure of a game includes all underlying motivations of the players.

When experimental methods started gaining more attention in economics, researchers were more or less aware of these challenges. To make the incentive structure work properly in experiments, it has to satisfy three basic criteria (Plott, 1986): Participants need to prefer larger to smaller rewards, they have to value experimental outcomes exclusively based on those

[1] See the meta-studies by Oosterbeek et al. (2004) and Johnson & Mislin (2011) for an overview.
[2] See e.g. Ledyard (1995) for an early survey about behavior in Public Good games.
[3] We frequently use the term "payoff" as a short version for monetary/material payoff.
[4] In line with the standard interpretation of the literature, we use the term "game form" such that it represents an exact and full description of the physical rules of a given of a situation.

rewards and they have to behave rationally in a sense that they try to choose their utility maximizing strategy.
Experimental researchers have strived to reduce potential difficulties that might arise by an appropriate design.[5] However, these efforts cannot eliminate the problems completely. Evidence suggests monotonicity in the size of rewards[6] and the salience of one's "optimal" (payoff-maximizing) strategy can be reasonably well achieved, at least in the games we plan to examine. However, the second criterion, which demands that participants derive utility only from their own material payoffs, is frequently violated across various experimental designs. Hence, in several cases the game-protocol and the corresponding payoff matrix do not accurately describe the actual game that people are playing according to their preferences. Furthermore, the magnitude by which behavior is potentially influenced often is ambiguous and varies across contexts. These issues have already been discussed in the literature, as for example in Weibull (2004), Bardsley et al. (2010), Hausmann (2005) and in Veszteg & Funaki, (2018). Yet, until now only few researchers have tried to take concrete steps to incorporate individual preferences into the analysis of decision problems.[7] Instead, it is usually assumed that participants value only their own material payoffs and game theoretic analyses are based on this assumption. As a result, the corresponding analyses may be biased.

To account for this, <u>two main challenges</u> emerge:

1. **Identification of the preference-based game**
2. **Providing common knowledge about the preference-based game**

First, we need to identify the game people are playing according to their preferences over the outcomes of the "original" game. These preferences have to capture all motivational aspects in addition to the players' own material payoffs. This will be our main task to deal with in this research project.

There have been several attempts to bridge the gap between players' material payoffs and the utilities they derive from these payoffs. Many important approaches use the concept of ***social*** or ***other-regarding preferences***. Social preferences refer to a class of preferences where decision makers do not only care about their own material payoffs but also about those of other parties. There exist two important branches of models: *outcome-based* and *intention-based models*. The models in the first category typically assign a utility value for the decision maker, which depends on both the values of his own payoff and those of other involved parties. Most prominent representatives for this category are the models by Fehr & Schmidt (1999) and Bolton & Ockenfels (2000). Alternatively, the type of interaction between the decision maker and others may play a role in how people evaluate the material consequences. This aspect is taken into account in *intention-based models* (see for example the models by Rabin, 1993 and Charness & Rabin, 2002). As our analysis will initially focus on simultaneous one-shot interactions, outcome-based models are more appropriate in terms of representing players' preferences.

[5] Measures for example include increasing the understanding of the task, reducing the influence of other factors (such as the degree of anonymity or experimenter-demand effects) and replicating effects with stronger monetary incentives.
[6] There exist many studies investigating behavior in often-applied games, such as the Ultimatum Game, the Trust Game etc. with high stakes incentives. On average, in these studies behavior is more close to the standard Nash prediction based on monetary payoffs. Therefore, it seems the monotonicity criterion is satisfied rather well.
[7] For some examples, see the studies by Healy (2011), Veszteg & Funaki (2018) and Alempaki et al.(2019).

A second challenge is that players typically can only observe the material payoffs of others, but not how they evaluate these payoffs in terms of utility. Hence, it seems unlikely that there is common knowledge about the structure of the preference-based game.[8] This knowledge, however, is relevant for the players for figuring out their optimal strategy/best response. Both factors, correct description of the preference-based game and common knowledge about it, need to be satisfied to expect that rational players will choose a Nash equilibrium strategy.[9] To avoid any misunderstanding, people still play the game according to its original game-protocol. The preference-based game is only used to incorporate additional motivational aspects into the analysis and fine-tune predictions.

We have initiated this line of research in our previous work in Brunner, Kauffeldt and Rau (2020). There we focus on the impact of revealing players' preferences on Nash-equilibrium play. The experiment consists of two treatments, both having the same two stages. In the first stage, we elicit subjects' preferences over several payoff tuples (x, y). Subjects create an ordinal ranking about a given set of eight different payoff tuples. In the second stage, subjects play four one-shot simultaneous 2x2 games against different opponents. The monetary outcomes of the games correspond to the same payoff-pairs subjects ranked beforehand. In one treatment, participants additionally receive information about the elicited preferences of their opponents before choosing a strategy, wereas in the other treatment subjects play the games without this information. Our design allows us to first identify the ordinal structure of the game subjects are actually playing according to their reported rankings (the "preference-based game"). To have a meaningful comparison, our analysis is based on these preference-based games and focus on cases where players have a unique non-dominant equilibrium strategy. Especially in these situations, knowledge of each other's preferences is supposed to have an impact, because it is relevant in figuring out one's best response. We then compare the frequency of equilibrium play between treatments. We find that subjects indeed are significantly more likely to play an equilibrium strategy, when preferences are mutually known. However, the magnitude of this effect depends significantly on the underlying game-protocol.

Through analyzing this experimental data, the examination of the difference between the game based on material payoffs and the game based on preferences became apparently pertinent. In our previous experiment, we only have data about preferences and behavior from a very limited number of specific classes of games and game protocols. Therefore, one of the main goals of our future work is to examine the research question *how games change when incorporating players' social preferences* in a comprehensive and systematic way. In doing so, we will propose a method to normalize the material payoffs of individual games, while preserving their respective ordinal structure. This makes their payoff structure more comparable, both across different classes of games as well as different versions within the same class. Because of that more general inferences concerning *all classes* of 2x2 games are possible. As it is often not feasible that both criteria are satisfied at the same time, it would be useful to check if analyzing the preference-based game is adequate to substantially improve predictions, even if there exists no common knowledge about it. Therefore, one of our additional goals is to disentangle individual influences of both factors and assess their joint and isolated effect.

---

[8] We do not distinguish between these concepts in detail. Arguing from a more theoretical perspective, we will use the term "common knowledge" and in the context of experiments, we will rather use the term "mutual knowledge", as it is very hard to assure that participants have common knowledge about a specific piece of information.

[9] For further details, see the theorem by Aumann & Brandenburger (1995) concerning equilibrium play in one-shot 2x2 games. In addition, one needs the assumption that all players are rational and that this is common-knowledge.

1.1 Project-related publications

1.1.1 Articles published by outlets with scientific quality assurance, book publications, and works accepted for publication but not yet published.

Brunner, C., Kauffeldt, F. and Rau, H. (2020). Does mutual knowledge of preferences lead to more Nash equilibrium play? Experimental evidence
(currently revise and resubmit at *European Economic Review*)

2 Objectives and work programme

2.1 Anticipated total duration of the project

The anticipated duration of the project is 36 months. Funding is required for the entire period.

2.2 Objectives

Our main objective is to analyze how the characteristics of relevant strategic situations change when accounting for individual (social) preferences over outcomes of games. We mainly focus on dominance and equilibrium properties, as these are the most important factors for predicting behavior of the involved players.

We start with a comprehensive analysis of all simultaneous one-shot 2x2 games (subproject 1). By using experiments, we can obtain data about preferences and behavior in the situations of interest. This provides an empirical assessment which preference-based game the players are actually playing. Furthermore, the behavioral data from the games allows us to assess if the frequency of Nash-equilibrium play increases when the analysis is based on the preference-based games. In a first step, we make inferences concerning the ordinal game structure and pure equilibrium strategies. In a further step, we plan to develop a cardinal measure for the utility agents are deriving from the outcomes of games. With the help of this measure, we can extend our analysis to both mixed and Bayesian Nash equilibria (subproject 2). In doing so players' monetary payoffs are replaced with the elicited values of our utility measure. This results in a new payoff matrix of the preference-based game, which can be used for computing these kinds of equilibria. In a last and more explorative step, we examine more complex categories of games such as sequential games (subproject 3).
We envision that our methods of preference elicitation will allow for considerably more precise predictions of behavior, both on the individual level and on the population level, compared to what is the current state of the art in experimental economics. They might also be employed by other researchers in cases where it is suspected that there are big differences between own monetary payoffs and the corresponding utility valuations.

2.3 Work programme incl. proposed research methods

The whole project can be divided into three subprojects. We explain each subproject in detail in the next sections.

Subproject 1: Preference-based games and equilibrium play in 2x2 games

We first examine possible transformations of one-shot simultaneous 2x2 games. In this class of simple games, it is reasonable to assume that people understand the mechanisms of the underlying strategic situation quite well.[10] Furthermore, as discussed earlier, the influence of additional factors besides the monetary payoffs is smaller compared to other categories of games. For this reason and in line with outcome-based models of social preferences, we will assume that players' utility u(x, y) is solely composed of the value of their own monetary payoff x and the other player's monetary payoff y. This assumption often is denoted as *consequentialism* in the decision-theoretic literature.[11] It provides an advantage for the upcoming analyses, as it allows us to elicit preferences over payoff pairs (x, y) independently of the context of the game. We do not impose any further restrictions on how these payoffs are evaluated in relation to each other. Preference relations over such pairs can then be used to make inferences concerning all games in which they appear. This procedure has the huge practical advantage that subjects only have to report their rankings of the tuples once, but do not need to play the games every time. This makes it possible to provide a comprehensive analysis of all classes of 2x2 games, which would never be feasible, if one had to acquire data for each specific game separately. However, this assumption is not completely without drawbacks. In line with the argument of Weibull (2004) and other researchers, when measuring players' preferences one has to capture all motivational factors. Concerning some of the situations we examine, as e.g. the Prisoner's Dilemma, one could argue that players receive an additional utility from the cooperative action itself, besides the monetary payoff pairs.[12] We cannot rule out this possibility completely, but believe the overall effect is rather small in simultaneous 2x2 games (compared to sequential games where reciprocity is a more pronounced factor concerning the evaluation of payoff allocations). To further account for this issue, we plan to run an additional treatment as a robustness check to examine if there is a difference between reported rankings within or without the context of a game. If in some cases there exist major differences, we could still adjust our elicitation method for those critical cases.

To be more specific on how our method works, consider the generalized payoff matrix of a simultaneous 2x2 game: This matrix consists of **four monetary payoff pairs (x, y)** which correspond to the four outcomes of the game resulting from the four strategy combinations. According to our approach, players have an ordinal preference relation about these four payoff pairs. Applying the topology suggested by Rapoport (1966) and used by Robinson & Goforth (2005), there exist **78 distinct strictly ordinal 2x2 games**.[13] From these 78 games, roughly 5-10 are of major interest in the field of game theory and experimental economics, for example, the Prisoner's Dilemma game, the Chicken game and the Stag Hunt game.[14] We plan to examine these games in more detail.

---

[10] That is, people understand what impact their choices have on the material consequences for the involved parties, given the behavior of the other players.

[11] Usually, this axiom states that the decision maker evaluates the material consequences independently of the state of the world. In our context, the "states of the world" mainly correspond to the different contexts of games and in later subprojects additionally to different histories of actions.

[12] Some authors, such as Healy (2011), provide empirical evidence for this claim.

[13] The authors count "equivalence classes of games", which can be transposed into each other by relabeling/switching the set of strategies or players, as corresponding to the same class of game.

[14] Some additional games of interest, as e.g. the Battle of Sexes and the Matching Pennies, may contain the same payoff value multiple times (depending on its exact representation). Our approach can easily be applied also to those cases, as in these games the set of possible payoff pairs is a subset from the set we use.

To allow for a comprehensive analysis of all 2x2 games we will *normalize* individual monetary payoffs to the values 1-4.[15] In doing so, number 4 corresponds to a player's highest payoff and number 1 to her lowest one. Applying this systematization, there exist exactly 16 different payoff pairs, ranging from the values (1, 1), (1, 2) … to (4, 4). These pairs cover all possible pairs appearing in all games. This normalization is possible as for now we only consider ordinal relations between payoffs. By having the same ordinal relation as in the original games the strategic properties of the games remain constant (existence of dominant strategies and pure equilibria). To illustrate how the normalization works, consider an arbitrary example of a symmetric Prisoner's Dilemma game.[16] In figure 1a the original version of the game is presented and in figure 1b the version with normalized payoffs.

	L	**R**
U	**4, 4**	**8, 3**
D	**3, 8**	**7, 7**

Figure 1a: Prisoner's Dilemma game with original payoff values

→

	L	**R**
U	**2, 2**	**4, 1**
D	**1, 4**	**3, 3**

Figure 1b: Normalized Prisoner's Dilemma game

In a similar way, all other 2x2 games can be normalized by using the 16 payoff pairs. By keeping the same "distance" between payoff values, it makes games also more comparable.

Experiment 1 – Determining the preference-based games:

Now individual's social preferences come into play. In our first experiment, we will ask the participants to create an ordinal ranking over the 16 normalized payoff tuples (x, y). As explained before, one of the main advantages of this approach is, that participants have to report only one preference ranking, instead of many rankings for each specific game-protocol. From the subject rankings of these 16 pairs, one can construct several main classes of preferences, for example, selfish preferences and mildly or strongly prosocial preferences. We use these categories to pool together more specific rankings. There are several ways, to induce incentive compatibility for reporting one's true preferences. Therefore, this should not pose a problem.[17] From our experimental data on individual rankings, we obtain an empirical assessment of the frequencies of the specific types of preferences.[18]

Forming (hypothetical) pairings of these different types can then be used to determine the dominance and equilibrium properties of the respective preference-based games and provide information how often these occur. We call this procedure "hypothetical", as people do not necessarily need to play the games themselves. Remember, the original games with their respective monetary payoff structure are not changed directly. These "preference-based games" are rather meant to provide a better understanding of the implied strategic situation that players with a specific preference structure would be facing. These insights can help to improve the game-theoretic predictions or to choose an appropriate game-protocol for inducing an intended strategic situation.

[15] Concerning the implementation, one could think of using appropriate multiples of these values.

[16] As an example, we use the payoff matrix from our previous experiment in Brunner et al. (2020). Of course, one could also use any other payoff values that satisfy the conditions of a Prisoner's Dilemma game.

[17] One could e.g. randomly draw two payoff pairs and select the better-ranked pair for individual payment.

[18] For example, Brunner, Kauffeldt & Rau (2020) report 45-50% subjects as having selfish preferences, 40-45% as being prosocial and approx. 10% as having competitive preferences.

The general procedure can best be illustrated by our example from before, the normalized version of the Prisoner's Dilemma game:

	L	**R**
U	**2, 2**	**4, 1**
D	**1, 4**	**3, 3**

Figure 2: Normalized Prisoner's Dilemma game

Typically, in a Prisoner's dilemma game, preferences over the four outcomes fall into two main categories.[19] Players either have *selfish preferences*, corresponding to the following ranking of payoff tuples: $(4,1) > (3,3) > (2,2) > (1,4)$
Alternatively, players exhibit *conditional cooperator* preferences, which correspond to the following ranking: $(3,3) > (4,1) > (2,2) > (1,4)$

If both players have selfish preferences, they only value their own payoff and the strategic situation remains a typical Prisoner's Dilemma game with the unique Nash-equilibrium (U, L). If players have conditional cooperator preferences, they do not have a strictly dominant strategy. In contrast to before, the resulting strategic situation corresponds to a coordination game with multiple Nash-equilibria. In this example both strategy combinations (U, L) and (D, R) would be an equilibrium in the preference-based game, adding mutual cooperation as another possible equilibrium outcome. **Assuming for simplicity that both types are equally likely**[20], there exist four different possible pairings, each with an ex-ante probability of 25%. The following table summarizes results about how the original game would transform for each possible pairing:

Original class of game	**Category of pairing**	**Probability**	**Class of preference game**	**Pure Nash-equilibria of preference game**
Prisoner's Dilemma	selfish vs selfish	25%	Prisoner's Dilemma	(U, L)
Prisoner's Dilemma	selfish vs cond. cooperator	50%	**Dominance Solvable Game**	(U, L)
Prisoner's Dilemma	cond. cooperator vs cond. cooperator	25%	**Coordination Game**	(U, L), **(R, D)**

Table 1: Possible transformations of the Prisoner's dilemma game for a given set of preferences

Given the assumed frequencies of preference types, **in only 25% of the cases** the preference-based game would still correspond to a Prisoner's Dilemma game. Most of the times players would actually face a game where only one player has a strictly dominant strategy ("dominance-solvable game"). Though the equilibrium (U, L) is the same as in the standard Prisoner's Dilemma game, the strategic situation is quite different: Being a conditional cooperator, the player has no dominant strategy and hence no unique equilibrium strategy ex ante. This brings us back to the discussion above where we argued on the importance of

[19] Of course, there might exist additional types as for example an "altruistic" player who ranks the tuple (1, 4) as better than (2, 2). However, such preferences types are not very common in the existing experimental literature.
[20] This assumption is not too far-fetched given the experimental data in Brunner et al. (2020) - see footnote 18.

players having common knowledge about each other's preferences. Furthermore, this example illustrates that the change of the class of game depends both on the specific preferences of the players as well as on the exact pairing. In a similar way, we will examine possible transformations of all other classes of 2x2 games.

From a theoretical point of view, our approach bears some similarities with Bruns (2012), who inspired our research ideas. The author analyses how (single) payoff swaps in 2x2 games can change situations of conflict into win-win situations. Accounting for player's preferences over payoff-tuples can have similar strategic consequences as a swap of a monetary payoff value in the original game (under the assumption that players are selfish). Both may change a player's ordinal preference ranking over the outcomes of a game. Our approach widely extents this idea by allowing *any kind of transformation* (and not only a single swap of one payoff value). Additionally, we bring this concept into life by having empirical data about when and how often these changes actually occur.

Concerning these results, we have to be aware that they strongly depend on the frequencies of particular types of preferences. These again depend on the exact values of the monetary payoff pairs appearing in the games, which we normalized to the values of 1-4 in our approach. When transferring our results to other games with arbitrary payoffs, we would expect some variation concerning the size and direction of effects. If one is especially interested in results of a specific game (for example the Stag Hunt game), one could perform the same analysis for several payoff configurations of this particular game (instead of applying the normalization). This would yield a more comprehensive picture with respect to this specific game. Of course, one needed many more observations, as one would need to elicit preferences for each specific payoff configuration separately. Nevertheless, the general results of the normalized versions of the games can still be very useful in providing a rough estimate in which situations the preference-based game is likely to be different from the game with original payoffs.

Experiment 2 – Impact on Nash-equilibrium play:

We additionally want to examine if one can achieve better predictions concerning Nash equilibrium play when using the preference-based game for the analysis instead of the original payoff structure. To answer this question, we need data from both individual preferences and actual behavior in the games of interest. As already discussed, this design bears some similarities with the one in Brunner et al. (2020), while extending our previous work. To incorporate players' preferences as before participants first have to report their rankings of the outcomes of the games. Afterwards in a first treatment they play those games in the "usual" way, which means with the original monetary payoffs and without having knowledge about other players' preferences. Then, we can simply compare the frequencies of equilibrium play based on the predictions from the monetary payoff structure with the predictions based on reported preferences.[21]
Since we already examined the sole effect of revealing each other's preferences on equilibrium play in our previous project, it is reasonable to study the joint effect. In order to do so, we need another treatment where one proceeds in the same way as before. In addition, players'

[21] Looking at our example, for the latter case this corresponds to the set of equilibria stated in the last right column. It might seem as if this is not a fair comparison, as here the set of equilibria gets richer when accounting for players' preferences. However, this impression is misleading. Depending on the specific class of game and the exact pairing the change in the sets of equilibria can go in either direction.

preferences are revealed before they make their decisions. Again, the frequencies of equilibrium play based on the predictions from the monetary payoff structure are compared with the ones from the preference-based games. In contrast to before, the preference-based games are now common knowledge to the players. Naturally, we expect the latter treatment to have the largest impact on equilibrium play. However, from an implementation point of view this procedure is also the most complicated. When participants know that their reported preferences will be revealed later on, we need to ensure incentive compatibility for truthful reporting. Furthermore, it might be that players' own behavior is influenced by the information about others' preferences, which would complicate matters further. Therefore, if one is mainly interested in results on the population level, it would be enough to simply adjust the predictions by including average degrees of social preferences in the analysis. Taken together, one can use the previous insights to evaluate which experimental design is the most appropriate for one's purposes.

Up to this point, we have only discussed ordinal preference relations and the existence of pure equilibria. In the next subproject we want to extend our analysis to mixed and Bayesian equilibria.

Subproject 2: Cardinal measure of preferences and mixed/Bayesian equilibria

So far, our proposed analysis only allows us to examine the equilibrium structure with respect to pure strategies. This already offers useful insights, especially when there exist strictly dominant and unique equilibrium strategies. However, in several of the preference-based games more outcomes can potentially result in an equilibrium compared to the analyses based on own monetary payoffs. The reason is that before the set of equilibria was determined assuming that only the selfish type of preferences existed. Now several additional preferences types are taken into account as well. This increases the set of possible equilibria. When for a given pairing the preference-based game contains several pure equilibria, then there exist also mixed equilibria. In this kind of situation, the optimal choice for the decision maker depends also on their exact evaluations of the equilibrium outcomes. For this, we need a cardinal utility measure. To elicit subjects' cardinal utilities we propose the method of **monetary equivalents.** This concept is very similar to the one of certainty equivalents. The decision maker is asked to state the value of a monetary payoff m^*, which makes her indifferent between receiving this payoff m^* instead of the outcome she has to evaluate. As in our context the outcomes correspond to payoff pairs, an important question concerning the implementation is what will happen with the payment to the other player. As any reference point is somehow arbitrary, we normalize these payments to zero.[22]

There exist also several alternative methods to measure subjects' preferences over outcomes in a cardinal way. All approaches have their advantages and disadvantages: A simple procedure, which is often employed by psychologists, is to elicit self-reported satisfaction levels of participants. While it is undeniably simple and easy to understand, it lacks proper incentives for truthful reporting. This can be problematic concerning the classes of games we plan to examine, because they often involve a conflict of interests between own material payoffs and those of other parties. The mechanism used by Alempaki (2019) and Brunner et al. (2020), where two outcomes are randomly determined and the better-ranked is selected for payment,

[22] Setting the other player's payment to a value of zero seems to be a natural reference point, but one could also choose any other value.

only works for eliciting ordinal preference relations. Another commonly employed method in experimental economics is the *binary lottery procedure (BLP)*. This mechanism is supposed to induce risk-neutral preferences in experimental decision making (under certain technical assumptions concerning the underlying expected utility function). As the evaluation of outcomes itself does not contain any sort of risk per se, there is no direct necessity to use a mechanism like the binary lottery procedure. Furthermore, from an implementation point of view, it is also rather complex, while there is an ongoing debate on whether the method works well from a behavioral perspective (see e.g. the studies by Harrison et al., 2013 and Oechssler & Sofianos, 2019). However, if one includes actual behavior into the analysis, individual risk preferences might come into play when the players choose mixed strategies. It is not desirable to confound different influential factors of the decision-making process. Therefore, we separate the evaluations of outcomes from the strategy selection process. Concerning the equilibrium predictions, in a first step we will simply assume risk neutrality with respect to our cardinal measure. If necessary, individual risk-attitudes could be incorporated in an additional step.[23] However, we feel this would complicate both the experimental design and the game-theoretic analysis without too many additional insights. For these reasons, we will start with a simpler approach.

To illustrate how the mechanism of *monetary equivalents* works, suppose a participant is asked to evaluate the outcome of a game, where both players receive a monetary payoff of 5. That corresponds to the payoff tuple (5, 5). Choosing a monetary equivalent of $m^* = 7$ means the decision maker is indifferent between receiving 7 monetary units for herself (and the other player receiving nothing) or both players obtaining a payoff of 5 units. As proposed in the first subproject, we assume an outcome-based utility function for the decision makers. In this case the procedure corresponds to a mapping from the utility function u(x, y) to the function u(m, 0), with m being the monetary equivalent of the player. A selfish decision maker is then characterized by the condition $x = m^*$, while for a prosocial player it holds $x < m^*$. For obtaining the payoff structure of the preference game, the values of the monetary equivalents can be used to replace individual monetary payoffs of the original games. Afterwards, one can perform game-theoretic analyses in the usual way, using these numbers as a proxy for the utility values of the players.
Let us demonstrate this procedure with the help of our leading example, the (normalized) Prisoner's Dilemma: If the players have selfish preferences, neither the individual payoff values nor the game change. Now, suppose both players exhibit conditional-cooperator preferences. A specific example of an outcome-based utility function reflecting these preferences would be:

$$\boldsymbol{U(x,y) = x + ay}, \quad \text{with } a = 1 \text{ if } x \geq y \text{, and } a = 0 \text{ if } x < y$$

Using the payoffs pairs from our game, we can calculate the corresponding monetary equivalents. For example, for the payoff pair (4, 1) we obtain $U(4,1) = 5 = U(m^*, 0)$. This yields a value of $m^* = 5$ for the monetary equivalent concerning this specific payoff pair. The other values can be derived analogously. The results are summarized in table 2:[24]

[23] Experimentally, one can simply elicit individual risk preferences in another task and use these results to make the predictions even more precise. From a more theoretical perspective, one can add a specific utility function, which assigns a utility value for the monetary equivalents to incorporate the risk attitude of the decision-maker.
[24] For the sake of simplicity, we created an example where the monetary equivalents correspond to integers.

PAYOFF PAIR (X, Y)	MONETARY EQUIVALENT (m^*)
3, 3	**6**
4, 1	**5**
2, 2	**4**
1, 4	**1**

Table 2: Players' monetary equivalents for all four outcomes of the game

Now, we can use these elicited values to replace the original payoffs with the players' monetary equivalents. In doing so, the payoff matrix transforms to the one depicted in figure 3:

	L	**R**
U	**4, 4**	**5, 1**
D	**1, 5**	**6, 6**

Figure 3: Corresponding preference-based game of the normalized PD game

In this new payoff matrix of the preference-based game, individual values are supposed to incorporate players' social preferences over the payoff pairs of the original game. From now on, we can proceed in the familiar way and calculate equilibrium predictions using own players' payoffs (the monetary equivalents). At first, we assume players have common knowledge about the preference-based game. When also considering mixed strategies, we obtain the following **three Nash-equilibria**: **Two pure equilibria (U, L), (D, R)** and **one mixed equilibrium [(1/4, 3/4), (1/4, 3/4)]**.[25] The preference-based game corresponds to a special class of a coordination game, namely the Stag Hunt game. It has one Pareto-dominant and one risk-dominant equilibrium. Considering existing empirical evidence, this seems reasonable, because mutual cooperation is seen as the globally best outcome for conditional cooperators. However, as they do not know for sure that the other player will cooperate some people might defect in order to avoid their worst outcome. We can derive the payoff and equilibrium structure of the preference-based games for all other 2x2 games in the same way. This work refines our insights from subproject 1 by adding information about the intensity of preferences. In turn, this makes it possible to extend the analysis to mixed strategies and mixed equilibria. This is especially helpful when the preference games contain the same ordinal structure as the original games, but their outcomes differ in their cardinal evaluation. In this case, we get a more precise prediction about behavior. Our approach has the attractive property that it can easily be applied to other classes of games (e.g. sequential games). Relating to the following discussion in subproject 3, in sequential games we would elicit preferences/monetary equivalents directly in the context of the specific games.

As a last step of this subproject, we want to extend our analysis even further towards a Bayesian framework. This is the proper approach when players do not have common knowledge about the preference-based game. It is appropriate, when it is reasonable to assume that people have a rough estimate about the overall distribution of types. In this case, a player's optimal strategy depends both on their beliefs of the distribution of types and the evaluation of outcomes. To fit into a Bayesian framework, we can use own players' cardinal evaluations of outcomes (the monetary equivalents) in combination with the general

[25] The basis for these calculations is the commonly used von-Neumann-Morgenstern (1947) expected utility function. As the payoffs in the matrix do not directly correspond to utility values, but to monetary payoffs instead, the assumption of risk neutrality is of more importance here. If one wants to additionally account for that, one could add a further step where these monetary equivalents m* are evaluated by a specific utility function u(m*).

distribution of elicited preference types to calculate Bayesian equilibria in the games of interest. Of course, this would be a more complex endeavor, as strictly speaking each distinct cardinal ranking corresponds to a unique preference type. Therefore, we need to develop reasonable categories to pool together different types of similar kinds. Here we can build up on results from subproject 1. Together one would have all necessary ingredients to examine if subject behavior is in accordance with some type of Bayesian equilibrium.[26]
In a last and more explorative step, we want to apply our methods to more complex categories of games such as sequential games.

Subproject 3: Preference-based games and equilibrium play in sequential games

We aim to examine sequential games involving two players and focus on ordinal preference relations. The general procedure in this subproject is similar to the ones before, with one major difference: We can no longer reasonably sustain the assumption that the utility, which individuals derive from the outcomes of games only depends on the monetary payoff pairs (x, y). In sequential games, the history of interaction can also play an important role in how these payoff pairs are evaluated. For example, Blount (1995) shows that the acceptance rate of low offers in the Ultimatum Game is much higher if they are made by a neutral third party or computer instead of a human player. This result suggests that participants may evaluate the very same material payoffs quite differently, depending on the exact conditions of their realization. This makes it necessary to elicit preferences directly in the context of the games and the history of actions. We plan to give participants a full description of the rules of the game, such that it is clear which history of actions leads to which outcome. Unfortunately, this approach reduces the number of situations we are able to examine significantly, as participants have to report preferences for each game separately. For this reason, we plan to focus on some prominent games like the Ultimatum Game and the Trust Game. As in the previous subprojects, our analysis will mainly deal with the question of how important properties of these games change when accounting for player's preferences over outcomes.[27]
As an example, consider the Ultimatum Game (UG) where a proposer (player 1) has to decide on how to allocate an endowment of 10 monetary units. In a simplified version of the Ultimatum Game, the proposer can either offer a fair split of the pie (F) or an unequal split (U). The fair split corresponds to the payoff tuple (5, 5), the unequal split to the tuple (8, 2). The responder (player 2) can either accept (A) or reject (R) the offer. If the offer is rejected, both players receive a monetary payoff of zero. The extensive form of the game is shown below:

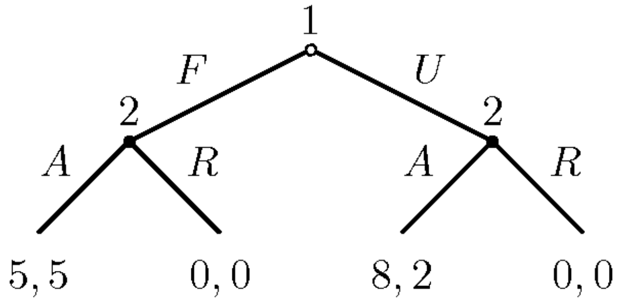

Figure 4: Extensive form of the Mini-Ultimatum Game

[26] Applying this idea to our leading example from before with the corresponding payoffs values and assuming there exist only two types, selfish players and cond. cooperators. The cond. cooperators do not know the type of their opponents, but have the belief that with a probability of α they are matched with the same type. Then, there exists a Bayesian-equilibrium where both players choose to cooperate for values of $\alpha \geq 0.75$.
[27] Again, we mainly focus on the change of the equilibrium structure.

Under the assumption that these payoffs fully represent the utility values of the players, the **unique subgame-perfect equilibrium** of the game is **(U, AA)**. However, experimental behavior is often not in line with this prediction. Assuming again, people are capable of understanding the nature of the decision problem, a plausible explanation is that players' utilities of the outcomes are different from the monetary payoffs. This is especially relevant regarding the responder's perspective as she has information about the choice made by the proposer and may want to reciprocate it. If the proposer has selfish preferences and the responder prefers the outcome (0, 0) to the outcome (2, 8) (e.g. due to fairness considerations or inequality aversion) then only the strategy combination **(F, AR)** would be a **subgame-perfect equilibrium** in the preference-based game.[28] Usually in the Ultimatum Game the equilibrium structure is mainly driven by the preferences of the responder, but changes can also be caused by the preferences of the proposer: If she has strong prosocial preferences such that she ranks the outcome (5, 5) as better than (8, 2), **two new equilibria emerge**: **(F, AA)** and **(F, AR).** In this case, the proposer has no incentive to deviate from the fair split, even if she knew the responder would accept the unequal offer. Note that we carefully distinguish between preferences and behavior. A selfish proposer might nevertheless offer the fair split because she fears the other offer would be rejected. In this example as well as in the Trust Game, the set of possible equilibria in the preference game will get richer when accounting for different types of preferences. Often it boils down to the question of how strong the second mover's social preferences are. This information is crucial in determining the optimal behavior of the first mover, even if she has selfish preferences. Again, the influence of common knowledge of preferences is of great importance and its effect might even be more pronounced than in simultaneous games.

Concerning the experimental design, an important question will be if participants should rank all outcomes of the games simultaneously ex ante (this is similar to what in the literature of experimental economics is known as "*strategy method*" or "*cold elicitation*"), or if they should make their decision when they are in the specific situation ("*hot elicitation*"). Concerning the hot or cold elicitation method, letting subjects evaluate all possible outcomes yields many more observations and offers more control.[29] On the other hand, the cold elicitation method is less able to capture all decision-relevant factors such as e.g. affective behavior and emotions to the same extent. The choice of the appropriate design is also linked to the question if one in addition needs to elicit data about actual behavior in these games. As we plan to do this as well, both options seem to be reasonable. Therefore, the decision concerning the specific experimental design will need careful consideration.

In a similar way as before, we plan to test if the frequency of Nash equilibrium play increases when incorporating players' social preferences into the equilibrium prediction (plus additionally making preferences commonly known). In doing so, one would need to elicit data about preferences and game-play at the same time. As the general procedure is very similar to the one in the first subproject, we do not describe it here in detail again.

Time schedule

The project is planned for a duration of 36 months. The following table shows the scheduling of the different tasks, which will partially be conducted in parallel.

[28] For completeness, one additionally has to assume that all players prefer the outcome (5, 5) to (0, 0).

[29] Otherwise, one is very much dependent on the exact pairings of preference types and likely has an unbalanced number of different histories of behavior.

Subproject	Task	Expected duration (in months)
1	Literature review: topology of 2x2 games and results	1-2
1	Finalizing of experimental designs	1
1	Programming & preparation of the studies	1-2
1	Data collection (all experiments)	3-4
1	Data analysis	1-2
1	Writing the paper & presentation of results	3
Duration for subproject 1		**∑ 12**
2	Literature review, theoretical analysis of preference elicitation mechanisms	2
2	Finalizing of experimental design	1-2
2	Preparing the study	1-2
2	Data collection	3
2	Data analysis	1
2	Extension to a Bayesian framework	2
2	Writing the paper & presentation of results	3
Duration for subproject 2		**∑ 14**
3	Literature review & selection of games	1
3	Finalizing of experimental design	1-2
3	Programming & preparation of the studies	1-2
3	Data collection	2-3
3	Data analysis	1
3	Writing the paper & presentation of results	2-3
Duration for subproject 3		**∑ 10**
Total duration of the project		**36 months**

Table 3: Estimated time for the different parts of the project.

2.4 Data handling

The experimental data will be made publicly available on the website of the host institution (Department of Economics, University of Heidelberg). The department has a special data repository for economic experiments. In case of publication, the data will be accessible on the website of the journal as well.

2.5 Other information

Covid-19 statement:

For the main parts of our work we plan to collect experimental data. Due to the current situation, at the moment (September 2020) it is not possible to perform experiments with participants showing up personally at the lab. The University of Heidelberg announced a statement that several proposals are being discussed to allow for such experimental work in the future (e.g. having smaller session sizes, using more rooms for conducting the studies etc.). We are optimistic to be able to pursue the data collection in the usual way by the beginning of the project (early-mid 2021). This would also be our preferred option. However, as this is not guaranteed we are developing alternative plans to collect the necessary data in case it will not be feasible by the means of a lab experiment. The parts involving preference elicitations can equally well be performed by an online experiment (e.g. using an appropriate software or commercial platforms) as they mainly involve individual decisions in a framework, which should be rather easy to understand. However, the parts involving strategic decision-making are harder to transfer to an online experiment, because there the settings are more complex and include interdependent decisions. Our idea for these tasks is to set up a live online experiment using the same subject pool registered for our lab experiments. There all participants would be active at the same time, but they do not need to be physically present. This would allow performing tasks where participants have to make choices simultaneously and the sessions can be monitored by the experimenters. In principle, the administration of the University has approved this practice. Only some minor issues concerning the data protection and payment mechanisms still need to be resolved.

3 Bibliography

List of references:

Alempaki, D., Colman, A., Kölle, F., Loomes, G., & Pulford, B. D. (2019). Investigating the failure to best respond in experimental games. *Working Paper*

Aumann, R., & Brandenburger, A. (1995). Epistemic conditions for Nash equilibrium. *Econometrica*, 63, 1161-1180.

Bardsley, N., Cubitt, R., Loomes, G., Moffat, P., Starmer, C., & Sugden, R. (2010). *Experimental Economics: Rethinking the rules*. Princeton University Press.

Blount, S. (1995). When Social Outcomes Aren't Fair: The Effect of Causal Attributions on Preferences. *Organizational Behavior and Human Decision Processes, 63(2)*, 131-144.

Bolton, G. E., & Ockenfels, A. (2000). ERC: A theory of equity, reciprocity, and competition. *American Economic Review*, *90(1)*, 166-193.

Bruns, B. (2012). Escaping Prisoner's Dilemmas: From Discord to Harmony in the Landscape of 2x2 Games. *Unpublished manuscript*, arXiv:1206.1880.

Charness, G., & Rabin, M. (2002). Understanding social preferences with simple tests. *The Quarterly Journal of Economics*, *117(3)*, 817-869.

Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, *114(3)*, 817-868.

Harrison, G. W., Martínez-Correa, J., & Swarthout, J. T. (2013). Inducing risk neutral preferences with binary lotteries: A reconsideration. *Journal of Economic Behavior & Organization*, *94*, 145-159.

Hausman, D. M. (2005). 'Testing' game theory. *Journal of Economic Methodology*, *12(2)*, 211-223.

Healy, P. J. (2011). Epistemic foundations for the failure of Nash equilibrium. *Unpublished manuscript*.

Johnson, N., & Mislin, A. (2011). Trust games: A meta-analysis. *Journal of Economic Psychology*, *32*(5), 865-889.

Ledyard, J. O., Kagel, J. H., & Roth, A. E. (1995). Public Goods: A Survey of Experimental Research. In *Handbook of Experimental Economics Vol. 1*, (pp. 111-194), Princeton University Press.

Oechssler, J., & Sofianos, A. (2019). The Binary Lottery Procedure does not induce risk neutrality in the Holt-Laury and Eckel-Grossman tasks. *Discussion Paper.*

Oosterbeek, H., Sloof, R., & Van De Kuilen, G. (2004). Cultural differences in ultimatum game experiments: Evidence from a meta-analysis. *Experimental Economics*, *7*(2), 171-188.

Plott, C. R. (1986). Rational choice in experimental markets. *Journal of Business, 59(4)*, 301-327.

Rabin, M. (1993). Incorporating fairness into game theory and economics. *American Economic Review*, *83(5)*, 1281-1302.

Rapoport, A. (1966). A taxonomy of 2× 2 games. *General systems*, *11*, 203-214.

Robinson, D., & Goforth, D. (2005). The topology of the 2x2 games: a new periodic table, *Routledge*, London and New York.

Sally, D. (1995). Conversation and cooperation in social dilemmas: A meta-analysis of experiments from 1958 to 1992. *Rationality and Society*, *7*(1), 58-92.

Veszteg, R. F., & Funaki, Y. (2018). Monetary payoffs and utility in laboratory experiments. *Journal of Economic Psychology*, *65*, 108-121.

Von Neumann, J., & Morgenstern, O. (1947). Theory of games and economic behavior, 2nd rev., *Princeton University Press*, New York.

Weibull, J. W. (2004). Testing game theory. In *Advances in Understanding Strategic Behaviour*, Palgrave Macmillan, London. 85-104.

4 Requested modules/funds

4.1 Basic Module

4.1.1 Funding for Staff

Conducting the experimental sessions of the project requires support from a student assistant. We plan to run separate experiments for each part of the project. The duration necessary to collect the experimental data is expected to be about 3 months per subproject. Altogether, we anticipate the data collection to last approx. 10 months in total. For that period, we need the help of the student assistant with a workload of 60h/month. Taking the standard gross payment of 13,60€/h, this amounts to **a total cost of approximately 8.200€**.

4.1.2 Direct Project Costs

N/A

4.1.2.1 Equipment up to Euro 10,000, Software and Consumables

We kindly ask for financial support of approx. **1.500€** for covering expenses related to the renewal of software licenses (Stata, Microsoft Office etc.) in case the host institution does not provide them.

4.1.2.2 Travel Expenses

As it is common, we plan to present the results of the project at international conferences. The results of each subproject should be presented at least at one major conference or workshop. Exact travel possibilities and costs are hard to estimate given the current situation. Furthermore, they depend a lot on the specific location. In the past, average costs were roughly 800€ per event (conference fees, travel expenses, hotel cost). Therefore, we ask for a total funding of **2.400€.**

4.1.2.1 Visiting Researchers (excluding Mercator Fellows)

To discuss the research questions and the results of the project in detail, an occasional face-to-face meeting with co-authors from the related previous project would be very helpful. We apply for two trips from the current location of the co-authors to Heidelberg, covering both travel expenses and hotel costs (approx. 300€ per trip, **600€ in total**).

4.1.2.2 Other Costs

As detailed in the work program, our project includes several laboratory experiments. In the first experiment from subproject 1, we plan to elicit preference rankings and pool them together into meaningful categories. To have reasonable numbers, we would like to have a minimum of 30 observations for each category. We can use our previous experiment (Brunner et al., 2020), as a benchmark. In that study, we have data from 164 subjects, with the lowest number of observations in a category being slightly below 20. Thus, for achieving a number of roughly 30 per category we expect a requirement of 250-300 subjects. In the second experiment for subproject 1, we compare frequencies of equilibrium play. Again, we can orientate on the effect size from our previous experiment. There, the treatment effect on the increase of equilibrium play was about 13%. As our research question is not exactly comparable, at least for one treatment we take a more conservative estimate of the increase being around 10%. Using typical parameters for a power calculation (significance level of alpha=5% and a power of 80%) yields a demand of approximately 400 observations. In addition, a robustness check might be necessary to test our assumption of consequentialism. For this around 150 subjects should be adequate. The numbers needed for the experiment in subproject 2 are harder to estimate, as we will employ a more detailed measure for preferences. Again, we will pool data into reasonable categories, but expect the type space to be larger. To account for this, we plan with the demand from the first experiment multiplied by a factor of 1.5. This would amount to approx. 450 subjects. In subproject 3, we will elicit data only from a limited number of games. In addition, the number of possible outcomes is smaller than before. On the contrary, we have to elicit data for different roles of players. These effects balance to some extent, in that we give a similar estimate like in our first experiment (300 subjects).

Taken together, this adds up to around **1.600 participants in total**. We expect the sessions to last between 45-60 minutes. The average payment for such experiments is about 12€/per hour (this includes the show-up fee and a variable component).[30] Given these calculations, **total expenses** for data collection would amount to approximately **19.200€.**

4.1.2.3 Project-related publication expenses

Submission fees for the journals we aim to publish in are currently about 120€. We kindly ask for covering the costs for submitting our papers five times to such journals. This corresponds to a **total cost** of **600€**.

4.2 Module Temporary Position for Principal Investigator

Funding for the position of the PI (Postdoctoral researcher, 100%) is requested for a period of 36 months. During the project, the PI will be located at the University of Heidelberg. From there he will pursue all project-related tasks. The environment provides excellent infrastructure concerning the execution of the project, both in terms of data collection as well as contact to other researchers and potential collaborators from the relevant field. Based on current remuneration schemes of the DFG's website, yearly salary costs are indicated with 74,100€. For the total duration of the project (36 months) this amounts to **costs of 222,300€.**

4.3 - 4.7

N/A

[30] These are the numbers for standard lab experiments. In case we need to conduct online experiments, the necessary payments will be significantly lower.

5 Project requirements

5.1 Employment status information

Dr. Hannes Rau, Postdoctoral Researcher, Department of Economics, University of Heidelberg (currently until March 31th, 2021).

5.2 First-time proposal data

Dr. Hannes Rau

5.3 Composition of the project group

It is planned to collaborate with several other renowned researchers (see below) for different parts of the project, but they are not directly part of the project group.

5.4 Cooperation with other researchers

5.4.1 Researchers with whom you have agreed to cooperate on this project

- **Prof. Jürgen Eichberger**, Head of the chair of Economic Theory 1 (emeritus), University of Heidelberg

Jürgen Eichberger has followed my work in the recent years and provided valuable comments and advice for this project proposal. He is an expert in all areas of the field decision-making under uncertainty and its applications to game-theory. He kindly agreed to share his expertise, especially to the parts of the project that are related to the theoretical aspects of game-theory and decision-making.

- **Prof. Jörg Oechssler,** Head of the chair of Economic Theory 2 and **Andis Sofianos** PhD (same chair), University of Heidelberg

In subproject 2, we intend to collaborate with Jörg Oechssler and Andis Sofianos concerning the extension to a Bayesian framework. Jörg Oechssler has excellent knowledge of game theory, both from the theoretical and applied perspective. Andis Sofianos is a specialist in experimental economics and his experience will be very valuable for elaborating the optimal design of the studies as well as for the implementation part.

- **Prof. Florian Kauffeldt**, Heilbronn, Professor of Quantitative Methods, Heilbronn University of Applied Sciences

Florian Kauffeldt was part of the research group of the previous project. He has expertise in both theoretical and in quantitative methods. Together we developed several ideas, some of which are part of this proposal. We plan to work mainly on the follow-up research questions from our previous project, where we examine the impact of incorporating players' preferences on the frequency of Nash-equilibrium play (subproject 1).

- **Prof. Stefan Trautmann**, Head of the chair of Behavioral Finance and Contract Theory, University of Heidelberg

Stefan Trautmann enjoys vast knowledge of the field of behavioral economics in general. In addition, he is a specialist in the area of social preferences and in methods of measuring beliefs and preferences in experiments. He promised to share his expertise about the latter by helping with the development of a cardinal utility measure and its concrete implementation within the experimental design.

5.4.2 Researchers with whom you have collaborated scientifically within the past three years

Christoph Brunner, Dietmar Fehr, Graciela Kuechle, Florian Kauffeldt, Jörg Oechssler, Alex Roomets, Christiane Schwieren, Stefan Trautmann, Yilong Xu

5.5 Scientific equipment

The University of Heidelberg provides a modern computer laboratory with more than 20 workstations for conducting economic experiments (AWI Lab). There is also a collaboration with the nearby experimental lab in Mannheim of similar size (mLab), which we are permitted to use for the data collection. Both labs offer access to a huge database of potential participants.

5.6 – 5.7

N/A

6 Additional information

N/A