

Essays on Teamwork

Dissertation zur Erlangung des Doktorgrades der
Wirtschaftswissenschaften (Dr. rer. pol.) an der Fakultät für
Wirtschaftswissenschaften der Universität Bielefeld

vorgelegt von
Philipp Külpmann

Juni 2016

Gutachter

Erster Gutachter: Prof. Christoph Kuzmics, PhD

Zweiter Gutachter: Prof. Dr. Frank Riedel

Adresse

Universität Bielefeld
Fakultät für Wirtschaftswissenschaften
Universitätsstraße 25
33615 Bielefeld
Germany

Printed on non-ageing paper – ISO 9706

Acknowledgements

Many people have supported me during my doctoral studies. So many that I won't be able to thank everyone but I'd like to thank a few of them.

First of all, I'm very grateful to Christoph Kuzmics for advising me, listening and helping me with all my smaller and bigger problems and always being infectiously enthusiastic and positive while doing so. Thank you for always being open to new ideas, for encouraging me to pursue them and for not holding back your opinion. I couldn't have wished for a better advisor.

Then, I'd like to thank Frank Riedel for agreeing to be my advisor, for always making time for me and helping to find a solution to my problems. You have given me so many insights over the years, not only into research but academia in general and you have opened up many unexpected opportunities.

Furthermore, I'd like to thank everyone at the IMW, especially, Anna, Bettina, Giorgio, Jan-Henrik, Oktay and Ulrike for always having an open door for me and for making me feel at home whenever I came back to Bielefeld. I'd also like to thank the BiGSEM, especially Herbert Dawid, Diana and Karin, the BGTS and Helga Radtke.

So many people have made me feel welcome in Paris and Warwick, I'd need another thesis to thank all of you, so let me just thank Agnieszka and her network group, Daniel for supervising and supporting me and Irina for all the great talks we had (Yay! Banana-cheesecake!).

Then, I'd like to thank all of my colleagues, especially Marieke, Oliver, Tobias, Yuanyuan, Davit, for being a great coauthor and person to talk to and for making my stay in Warwick so enjoyable, Johannes, for being a great coauthor and for always trying (and achieving) things that I thought to be impossible (I am still waiting for the drill-quadcopter though), Martin, for all the endless talks we had, for always being there to help and for always having a different view on almost everything, and Yuliia, for being the heart of your little group and for her incredible green thumb.

Finally, I'd like to thank Anne. I wouldn't have been able to start or complete this thesis without you, thank you!

Contents

1	Introduction	1
2	Higher-order Beliefs and Coordination Failure	4
2.1	Introduction	4
2.1.1	Related works	7
2.2	Example	10
2.3	The model	12
2.4	Hypotheses	13
2.5	Experimental design	15
2.5.1	Hypotheses	16
2.6	Results	18
2.7	Conclusion	23
2.8	Appendix	24
2.8.1	Belief hierarchies	24
2.9	Equilibrium selection and models of higher-order beliefs	25
2.9.1	Order effects	27
2.9.2	Data	29
2.9.3	Instructions	30
2.9.4	Quiz	33
3	Projects with Uncertain Requirements	38
3.1	Introduction	38
3.1.1	Literature	41
3.2	The model	43
3.3	The breakthrough effort distribution	45
3.4	Results	46
3.4.1	Non-cooperative solution	46
3.4.2	Welfare maximizing solution	48
3.4.3	Free-riding, encouragement effect and delay of effort	50
3.5	Discussion	52

3.5.1	Deadlines	53
3.5.2	Full information	54
3.5.3	Patient players	55
3.5.4	The effect of the number of players	55
3.5.5	Effect of costs and project quality on different projects	56
3.5.6	Concluding remarks	57
3.6	Appendix	59
3.6.1	Derivation of the expected utility	59
3.6.2	Theorem 1 (Optimal Effort)	60
3.6.3	Theorem 3 (Last-minute rush)	63
3.6.4	Lemma 1 (Asymmetric equilibria)	64
3.6.5	Theorem 2 (Social planner)	64
3.6.6	Constant hazard rate	64
4	Probabilistic Transitivity in Sports	66
4.1	Introduction	66
4.2	Setup	69
4.3	The optimization problem	71
4.4	Optimization under a known ranking	72
4.4.1	Transitivity without draws	72
4.4.2	Solution process for the case including draws	73
4.5	Ranking methods	75
4.5.1	The Linear Ordering Problem	75
4.5.2	Branch and Bound Algorithm	76
4.5.3	Tabu Search	78
4.5.4	Popular ranking methods	79
4.6	Comparing the explanatory power of rankings	80
4.7	Data	81
4.8	Empirical analysis	82
4.8.1	Soccer in Austria: Finding an optimal ranking	82
4.8.2	Ranking systems and maximum likelihood estimates	83
4.8.3	Hypothesis testing	88
4.9	Conclusion	91
4.10	Appendix	93
4.10.1	Proofs	93
4.10.2	Parameter space	94
4.10.3	Code	95

Chapter 1

Introduction

Over the past few decades, teamwork has gained more and more in importance and it is surely but steadily replacing individual work in the workplace. At the same time, there is evidence from psychology (e.g., Brennan and Enns (2015)) and management science (e.g., Harvard Business School Press (2004)) that working in teams is more efficient than working alone. Therefore, it is not surprising that teamwork is becoming more and more popular. This trend is especially visible in academia. Wuchty, Jones, and Uzzi (2007) have shown that in an increasing number of fields, the number of research done in teams dominates single-authored works not only in numbers but also in average citations. But teamwork is not just restricted to the workplace. Working in groups is also popular in classrooms all over the world (Hutchinson (2001)) and even in sports, team-based sports usually overshadow individual sports in popularity.

But working in teams creates problems, which do not occur when working alone. In this work, I will discuss three different problems inherent to teamwork: Coordination problems, moral-hazard problems and how to evaluate and compare the results of different teams.

Coordination problems, are problems in which the incentives of the team members are perfectly aligned but they have to coordinate on one of the available options. An example for this is agreeing on norms or technical standards, like a communications protocol for computers.

However, when working in a team, incentives are rarely perfectly aligned. When team members share a common goal but have to exert private effort, which cannot be observed, moral-hazard problems come into play. The most commonly known effect of moral hazard is free-riding: Having more team members can lead to everyone working at inefficiently low levels.

The last problem discussed in this work is how to evaluate and compare the results of

different teams or parts of teams. This is a problem occurring in many different settings, however, here we are going to focus on an application in sports: What ranking scheme should we use when teams only compete one-on-one?

In this thesis I am going to address each of the three aforementioned problems, using different methods, including game theory (Chapters 2 and 3), laboratory experiments (Chapter 2) and statistical modeling (Chapter 4).

In Chapter 2, based on a joint work with Davit Khantadze, we are analyzing coordination problems. A coordination problem might occur if you have lost your spouse in a department store and both of you are trying to find each other. In this example you have to try to guess the place your spouse will go to. At the same time she has to guess where you are going to be. This complicates the (seemingly simple) question of “Will she look for me at the coffee bar or at the exit?” which now depends not only on the answer to the question “Does she think I am looking for her at the coffee bar or at the exit?” but also on the answers to “Does she think that I think that she thinks that I am looking for her at the coffee bar or at the exit?” and on infinitely more levels of so-called beliefs.

We are using a game theoretic model to analyze if higher-order beliefs play an important role in coordination problems when the players are facing a pure coordination game, i.e., a game in which the players have perfectly aligned preferences. To do so, we are using a laboratory experiment in which we introduce cognitive types into a pure coordination game in which there is no common knowledge about the distribution of cognitive types. In our experiment, around 76% of the subjects managed to coordinate on the payoff-dominant equilibrium despite the absence of common knowledge. However, around 9% of the players had first-order beliefs that lead to coordination failure and another 9% exhibited coordination failure due to higher-order beliefs. Furthermore, we compare our results with predictions of different models of higher-order beliefs, commonly used in the literature.

In Chapter 3, I am analyzing a model in which there is not only unobservable effort choice but in addition uncertainty about the requirements to complete a project, i.e., the players don’t know how much they have to work to complete it.

In the model, I analyze a dynamic moral hazard problem in teams with imperfect monitoring in continuous time. In the model, players are working together to achieve a breakthrough in a project while facing a deadline. The effort needed to achieve such a breakthrough is unknown but players have a common prior about its distribution. Each player is only able to observe their own effort, not the effort of others. I characterize the optimal effort path for general distributions of breakthrough efforts and show that, in addition to free-riding, a delay of effort and an encouragement effect, similar to Bolton and Harris (1999) arises. In this model, the encouragement effect increase and

decrease the work players put into the project, depending on the type of uncertainty faced. Furthermore, the delay of effort is also a result of rational and even welfare-maximizing behavior.

In Chapter 4, a joint work with Johannes Tiwisina, we develop a statistical model to describe results in sports and discuss its implications on different ranking schemes. In most sports, we don't just have teams competing against each other but these teams are usually also organized in associations like the FIFA or UEFA for soccer or the NCAA for college sports in the United States. In these organizations there are frequent discussions about the way a tournament or league should be organized to ensure that, in the end, the best team wins. One example for this is the discussion during the recent (2016) European Championship in which the system was changed to accommodate more teams into the tournament. But also leagues change their system. In soccer, most countries changed from a 2-points-for-a-win to a 3-points-for-a-win system between 1980 and 2000 and in 2014 the NCAA made a widely discussed change to the scoring system of the first division of college football.

We seek to find the statistical model that most accurately describes empirically observed results in sports. The idea of transitive relations concerning the team strengths is implemented by imposing a set of constraints on the outcome probabilities. We theoretically investigate the resulting optimization problem and draw comparisons to similar problems from the literature. We propose a branch-and-bound-algorithm for an exact solution and a heuristic method for quickly finding a good solution. Finally we apply the described methods to panel data from soccer, American football and tennis and also use our framework to compare the performance of empirically applied ranking schemes.

Chapter 2

Higher-order Beliefs about Cognitive Skills Can Lead to Coordination Failure

2.1 Introduction

If you have lost your spouse in a department store and both of you are trying to find each other, the answer to the (seemingly simple) question of “Will she look for me at the coffee bar or at the exit?” depends not only on the answer to the question “Does she think I am looking for her at the coffee bar or at the exit?” (i.e., something we will call the first-order belief) but also on the answers to “Does she think that I think that she thinks that I am looking for her at the coffee bar or at the exit?” (i.e., the second-order belief or “What is her first-order belief?”) and on infinitely more levels of beliefs. This chapter addresses the question if people actually use beliefs of a higher order.

When modeling human behavior, we usually assume that players have common knowledge about the structure of the game, i.e., that all players know the structure, that all players know that everyone else knows the structure and so on ad infinitum. Furthermore, we assume that players do not only have common knowledge about publicly known properties of the game but also about the distributions of unknown factors of the game, like the other players’ types (for example if I’d rather wait at the coffee bar or the exit). As the absence of common knowledge leads to complex belief hierarchies, so called *higher-order beliefs*, common knowledge is usually assumed for tractability reasons. The first level of these beliefs, so called first-order beliefs, might be a belief over the other player’s type. A second-order belief would then be a belief over the be-

belief of the other player about your type (i.e., a belief over the other player's first-order belief) and so on. The question we are trying to answer in this chapter is, how large the influence of the assumption of common knowledge is and if people use higher-order beliefs in coordination games.

More important applications than the search for ones husband or wife in a department store are suggested by recent studies in sociology and development studies, for example by Bicchieri (2005). She claims that common knowledge plays a significant role in the fight against female genital mutilation.¹ Female genital mutilation is practiced in, predominately African, communities and is required in many of these communities to find a husband and to prevent social exclusion. Despite being very dangerous and unnecessary, it has a long standing tradition and is, in areas where it is still practiced, very common. It is estimated to effect up to 200 million women in 2016 (UNICEF (2016)). In game theoretic terms the problem is one of equilibrium selection: There is one equilibrium in which everyone accepts and uses female genital mutilation and one in which no one does. The latter equilibrium is, given enough knowledge about the subject, clearly better for everyone, but we still observe the former equilibrium in many communities.

An important tool in the fight against female genital mutilation is to inform people about the dangers and (lack of) benefits of it. However, studies like Bicchieri (2005) suggest that just educating might not be enough. She claims that common knowledge of this education plays an important role because negative beliefs about the opinion of the other members of a community might prevent a coordination on the better equilibrium (i.e., the one without female genital mutilation): Even if I am convinced that this practice should be abolished, I might still partake in it, to prevent my daughters from being excluded from the community, as the others might not be convinced (i.e., my first-order belief is that others have not been educated). I also might think that others will continue this practice because they think I wasn't educated (i.e., because of my second-order belief) and so on.

That means, that just educating a family (or, in game theoretic terms: changing their type) does not necessarily lead them to change their stance on female genital mutilation. But is there any evidence that families use beliefs? Mackie (1996) and Mackie and LeJeune (2009) have compared the old Chinese tradition of foot binding² and female genital mutilation and pointed out that both are similar: Both are required to find a

¹Most studies, however, don't use the terms "beliefs" or "common knowledge" but describe this concepts in their own words, frequently restricting their attention to first-order and therefore ignoring higher-order beliefs.

²Foot binding was a Chinese practice of bending and binding the four lower toes of young girls tightly under the foot.

husband, while being very painful and dangerous without having any known benefits. Furthermore, they have a long-standing tradition (both can be traced back more than 1000 years) and were widely spread in their respective cultures. However, around 1910, foot binding has dropped in certain parts of China from 99% to under 1% prevalence over the course of just 20 to 30 years, without any change in policy (Gamble (1943), Keck and Sikkink (1998)), whereas even a combined effort of the UN, several NGOs and governments over the last 40 years resulted only in a moderate decline from about 51% to 37% of women effected by female genital mutilation in certain countries (UNICEF (2016)). Mackie (1996) claims that the main difference is the method of information transmission: In China, societies have been founded in which members publicly pledged to not bind their daughters' feet and to prevent their sons from marrying women with bound feet, whereas the effort to prevent female genital mutilation was mainly focused on changing the laws and educating the people about the dangers and problems. The societies fighting foot binding made the education and position of the families common knowledge whereas most organizations fighting female genital mutilation focused on changing the opinion of the families without changing the higher-order beliefs.

But also between projects fighting female genital mutilation there have been differences. Tostan, a Senegal-based NGO, has, according to World Bank Group (2012) successfully reduced the number of female genital mutilation in some parts of Senegal significantly. So, why did Tostan succeed where others have failed? They claim that not only education but “[...] public declarations are critical in the process for total abandonment [of female genital cutting]“ (Tostan (2016)) and are supported by World Bank Group (2012) who emphasizes that education together with public discussion and public declaration was an important factor in Tostan's success.

These examples suggest that beliefs might play an important role, as the more successful campaigns against foot binding and female genital mutilation also addressed higher-order beliefs by introducing common knowledge whereas others who focused on pure education have been less successful. However, it is not clear that common knowledge is required to achieve coordination. It might be sufficient to explain that others have also been educated (i.e., to take care of the first-order beliefs), which would be much cheaper than providing common knowledge. Therefore, the question if people actually use higher-order beliefs is an important one.³

Unfortunately, the game theoretical literature is not able to model these higher-order beliefs consistently: In practice many different assumptions and models of higher-order beliefs exist and many of these lead to very different predictions even in simple games

³More examples in which common knowledge plays an important role can be found in Chwe (2013).

like the pure coordination game we are using in this chapter.⁴ The question, what kind of model of higher-order beliefs players actually use, seems to be an empirical question which we are trying to address in this chapter.

We are analyzing the effect of absence of common knowledge in an experimental setting, using a certain type of simple coordination games to ensure that the effect of strategic uncertainty is reduced to a minimum, i.e., players have no incentive to "outsmart" the other players. In these games, there is uncertainty about the type of the other player, but no common knowledge about the distribution of these types. We are building on the work of Blume and Gneezy (2010) who have shown that "beliefs matter", i.e., that some people use beliefs that cause coordination failure. Using and extending their design, we are trying to answer the following three questions:

- **Are players able to coordinate in the absence of common knowledge?**
- Can coordination fail because players underestimate the skill of the other players? Or, in other words, **do first-order beliefs matter?**
- Can coordination fail because players think "too much" about what others might think? Or, in other words, **do higher-order beliefs matter?**

Using Blume and Gneezy's (2000) 5-sector disc, we were able to find answers to all three questions: In the experiment, the majority of players had no problem coordinating on the Pareto-dominant equilibrium of the game. However, some players switch to the worse equilibrium because of first- and higher-order beliefs.

This chapter is organized as follows: In Section 2.1.1, we will give an overview of the relevant literature and how our work fits into it. Then we will explain an example of the game we use in Section 2.2. In Section 2.3 we will explain the model and briefly discuss predictions made by some commonly used models of higher-order beliefs for this game and formalize the three questions stated before in Section 2.4. This is followed by the experimental design in Section 2.5 and the results of the experiment in Section 2.6. Finally, we will conclude in Section 2.7.

2.1.1 Related works

There is a large theoretical literature, beginning with the seminal paper on the "email game" by Rubinstein (1989), showing that higher-order beliefs play a role in determining the outcome of a game. For instance, Carlsson and Van Damme (1993) use higher-order beliefs (in their model of global games) to identify the risk-dominant equilibrium as the unique rationalizable outcome of the coordination game. This uniqueness result

⁴A brief overview of some models of higher-order beliefs can be found in Section 2.1.1 and a more detailed discussion in Section 2.9.

spawned a large applied literature on, among other areas, bank runs and arms races, in e.g. Morris and Shin (1998), Morris and Shin (2004), Baliga and Sjöström (2004), Corsetti, Dasgupta, Morris, and Shin (2004), and Goldstein and Pauzner (2005). Weinstein and Yildiz (2007b), however, have shown that this uniqueness result, that this whole literature depends on, is fragile to the exact specification of the higher-order belief model. Other “nearby” higher-order belief models have very different “unique” predictions. In fact, they show that any rationalizable outcome of the original game, can be obtained as the unique rationalizable strategy profile of some higher-order belief model.

Weinstein and Yildiz (2007a) establish a condition, called “global stability under uncertainty”. This condition implies that, if the change in equilibrium actions is small in the change of k th-order beliefs and higher, equilibria can be approximated by the equilibrium with at most k th-order beliefs. Unfortunately, pure coordination games do not fulfill “global stability under uncertainty”.

Strzalecki (2014) and Kneeland (2016) develop different non-equilibrium approaches, inspired by the experimental literature discussed later, using bounded levels of reasoning to explain behavior in coordinated attack problems (e.g. Rubinstein’s (1989) email game).

A more in-depth discussion of models of higher-order beliefs and their predictions of the results of our experiment can be found in Section 2.9.

The experimental literature, however, has so far mostly focused on strategic uncertainty. The most prominent example for this is probably the literature on level- k thinking or cognitive hierarchy models, which was started by Nagel (1995) and Stahl and Wilson (1995). In recent years, there have been many studies conducted, using and analyzing level- k reasoning, for example Ho, Camerer, and Weigelt (1998), Costa-Gomes, Crawford, and Broseta (2001), Camerer, Ho, and Chong (2004) and Crawford, Gneezy, and Rottenstreich (2008). For a recent survey, see Crawford, Costa-Gomes, and Iriberry (2013).

But there also have been works which do not focus on strategic uncertainty. For example Heinemann, Nagel, and Ockenfels (2004), Cornand (2006), Cabrales, Nagel, and Armenter (2007) and Duffy and Ochs (2012) who directly test implications of the theory of global games, i.e. individuals play an incomplete information game as in Carlsson and Van Damme (1993). The results however, are mixed and range from full support to full rejection of the predictions made by global games.

Another, closely related work is Kneeland (2015), in which she explores the level of rationality, a requirement for higher-order beliefs, of players experimentally. She shows that, in her experiment, 94% of all players are rational with decreasing numbers for

second- (71%), third- (44%) and forth-order (22%) rationality.

We explore experimentally the “depth of reasoning” individuals employ when playing slightly difficult coordination games. In fact we want to abstract away from purely strategic concerns by only looking at coordination games in which the incentives of the players are perfectly aligned and a Pareto-dominant equilibrium exists. The fundamental uncertainty in the model will be one about the cognitive abilities of the opponents.

Differences in cognitive abilities have been studied before, for example by Gill and Prowse (2016), who have shown that more cognitively able subjects converge, in repeated p-beauty contests, more frequently to equilibrium play and earn more. Furthermore, Proto, Rustichini, and Sofianos (2014) have shown that intelligence affects the results of repeatedly played prisoner’s dilemmas, in which groups of higher intelligence tend to cooperate more frequently in later stages of the game. Agranov, Potamites, Schotter, and Tergiman (2012) have shown, by manipulating the perception of the cognitive levels of other players, that beliefs about the level of reasoning do play a significant role in the presence of strategic uncertainty. Alaoui and Penta (2015) establish a framework in which the depth of reasoning is endogenously determined by different cognitive costs of reasoning.

The way we model cognitive differences however, builds on another branch of literature. Motivated by Schelling’s (1960) discussion of focal points, a variety of authors have tried to formally capture his ideas, most notably Bacharach (1993) and Sugden (1995). The importance of focal points is supported by many experiments, for example by Mehta, Starmer, and Sugden (1994), who have replicated Schelling’s results and have shown that coordinating on a focal point is different from accidental coordination. Crawford, Gneezy, and Rottenstreich (2008) have shown that, in a pure coordination game with symmetric payoffs, salient labels lead to a high percentage of coordination whereas even slight asymmetries in payoffs might lead to a coordination failure. Isoni, Poulsen, Sugden, and Tsutsui (2013) extend the analysis to bargaining problems and show that payoff-irrelevant clues help to improve coordination, even if there is no efficient or equal division.

In the absence of clues however, the theory of focal points can not be applied. Formally, the absence of clues can be modeled as symmetries between strategies and players in a given game. In fact Nash (1951), has already discussed equilibrium under symmetry restrictions (and shown existence also of such symmetric (mixed) equilibria for finite games). Crawford and Haller (1990) have defined symmetries in games and used these definitions to see what focal points in highly symmetric repeated coordination games

would look like.⁵ Blume (2000) has further developed this symmetry concept to talk about play under the absence of a common language. Other notions of symmetries have been put forward and studied in Harsanyi and Selten (1988), Casajus (2000) and Casajus (2001). Alós-Ferrer and Kuzmics (2013) have then clarified the difference between different notions of symmetries and characterized all the possible ways a frame (the way a game is presented to players in the lab, for instance) could lead to different symmetry restrictions (and therefore to different focal points).

All these models of symmetries and restrictions are implicitly or explicitly investigated under the assumption of perfectly rational individuals. However, identifying all symmetries (and especially non-symmetries) in a game can be a difficult task. Bacharach (1993) has proposed his variable frame theory to allow for individual players with different states of mind or, as developed by Blume (2000) and employed by Blume and Gneezy (2000) and Blume and Gneezy (2010), with different cognitive abilities.

This finally brings us to the goal of our study. We want to take up the experimental results and setup of Blume and Gneezy (2010), in which there is an issue of cognitive difficulties, to analyze the effects of higher-order beliefs. Blume and Gneezy (2010) were able to show that participants form beliefs about the cognitive abilities of other participants and, if these beliefs are pessimistic, they hinder coordination between the players. However, they have not taken into account the effect of higher-order beliefs about cognitive abilities. Therefore, we modify their experimental setup in order to distinguish the effect of first-order beliefs players form about the cognitive ability of their opponents (i.e., if players trust in the cognitive ability of their partners) and higher-order beliefs.

2.2 Example

In this example, players only have access to two strategies l and h and are trying to coordinate on one of them; the payoffs are as depicted in the payoff matrix in Figure 2.1. As (h, h) has a higher equilibrium payment it would therefore be the focal point (and the risk- and payoff-dominant Nash equilibrium) of this particular game.⁶

However, if we introduce cognitive differences, i.e., if action h is only available to a high-cognition player and low-cognition players are forced to play l , beliefs about the

⁵Bhaskar (2000) and more comprehensively Kuzmics, Palfrey, and Rogers (2014), have studied theoretically and in the latter case also experimentally, what the possible focal points of the symmetric repeated battle-of-the-sexes and its generalizations could be.

⁶Or, in the words of Luce and Raiffa (1957) and Schelling (1960) a *solution in the strict sense*.

	l	h
l	1,1	0,0
h	0,0	3,3

Figure 2.1: Payoff matrix of a high-cognition player

other player’s type might lead to coordination failure,⁷ even if both players are high-cognition players. The driving force of this result is the absence of common knowledge about the players’ type or the fraction of high cognition players.

The following two examples show how beliefs could lead to coordination failure between two high-cognition players: First imagine that the first player (she) thinks that the other player (he) is a low-cognition player. Then she would play l , as he would have no other choice than playing l . This is what we will call coordination failure due to a first-order belief. The second example is that she thinks that his type is high, he thinks she is a high-type player but she thinks that he thinks her type is low. Again, she would play l as she thinks that he will play l . Here we have a coordination problem due to her second-order belief. Therefore, even if both players have the ability to coordinate on the best equilibrium, they might end up failing to coordinate on the better equilibrium (h, h) .

The existence of infinitely many levels of beliefs and that a “bad” belief at any level makes the player switch to the “bad” strategy l makes one wonder, if, even with a high fraction of high-cognition players, coordination on the good equilibrium (h, h) is possible.

Therefore, the first main question this chapter addresses is if coordination on the good equilibrium can be expected even in the absence of common knowledge. The second question is if systematic underestimation of other players’ skills can be a source of coordination failure, or if first-order beliefs matter. The third and last question is if higher-order beliefs, e.g. if she thinks that he thinks that she is a low type, are a possible cause for coordination failure or if these levels of reasoning are too complex and play no significant role in coordination games.

The concepts of coordination games and higher-order beliefs will be formalized in the following section.

⁷In this chapter, we follow the notion for coordination failure of Van Huyck, Battalio, and Beil (1990), i.e., the failure to coordinate on the best achievable outcome. That means, even if two high-cognition players coordinate on a Pareto-inferior equilibrium we will call it coordination failure.

2.3 The model

We begin by defining a pure coordination game for two players.

Definition 1 (Pure coordination game). A *pure coordination game* is a game with 2 players, who each have access to m different actions ($\{a_1, a_2, \dots, a_m\}$).

In this game payoffs of a player i are defined as

$$u_i(a_i, a_j) = \begin{cases} x_i & \forall i, j : i = j \\ 0 & \text{otherwise} \end{cases}$$

with $x_m > x_{m-1} > \dots > x_1$.

This means that each player can choose from the same set of actions and whenever they have picked the same action they get the same payoff and if they don't manage to coordinate their actions, both get nothing. Furthermore, there is a Pareto ordering of these equilibria. Figure 2.2 shows an example of a pure coordination game with three possible actions.

	a_1	a_2	a_3
a_1	1, 1	0,0	0,0
a_2	0,0	2, 2	0,0
a_3	0,0	0,0	4, 4

Figure 2.2: A pure coordination game

Let us now introduce cognitive differences into this pure coordination game. For the sake of simplicity, we are only introducing two cognitive types, a low-cognitive type and a high-cognitive type. The latter has access to a “better” strategy, which is not available to the low type. Furthermore, the low type is unaware of the existence of the high type, as proposed by Bacharach (1993).

Definition 2 (Pure coordination game with cognitive differences). A *pure coordination game with cognitive differences* is a game with 2 players. Each of the players has a type $t_i \in \{low, high\}$ and has access to different strategies, depending on his type t_i . The types of a player are her private information. Low cognition players have access to $\{a_1, a_2, \dots, a_{m-1}\}$ whereas high cognition players also have access to the action a_m , i.e., to $\{a_1, a_2, \dots, a_m\}$.

In this game payoffs of a player i are defined as

$$u_i(a_i, a_j) = \begin{cases} x_i & \forall i, j : i = j \\ 0 & \text{otherwise} \end{cases}$$

with $x_m > x_{m-1} \geq \dots \geq x_1$.

These cognitive differences can also be thought of as symmetry constraints on attainable strategies, as proposed by Crawford and Haller (1990) and further developed by Blume (2000) and Alós-Ferrer and Kuzmics (2013). Here, the high-cognition player has less symmetry constraints and has therefore more attainable strategies.

In the experiment we are using the notion of cognitive differences as proposed by Blume and Gneezy (2010) (a generalization of Bacharach’s variable frame theory, using different symmetry constraints on the attainable strategies as used in Blume (2000)).

For a formal description of the belief hierarchy of these games, we would like to refer to Section 2.8.1. However, we believe for understanding the results of this work, the idea conveyed in Section 2.2 should suffice.

2.4 Hypotheses

Before investigating the three original research questions, we will have a look at some preliminaries. First, we expect Nash equilibria and, if these equilibria can be Pareto-ranked, the Pareto-better equilibrium to be played in the pure coordination game. This is supported by the literature (e.g., Van Huyck, Battalio, and Beil (1990) or Cooper, DeJong, Forsythe, and Ross (1990)) and was also corroborated by the choice data of the experiment. For low-cognition player that means that the Pareto-dominant action a_{m-1} will be chosen over all other actions, as, for him, the game is a simple pure coordination game, because he does not know about the existence of the high type. The high-cognition player, however, has two valid options: a_m and a_{m-1} . As we have seen in the example, the answer to the question if the high-cognition player chooses the payoff-dominant strategy a_m or the second-best strategy a_{m-1} , depends on her beliefs. We have chosen the game in such a way that the fraction of high-cognition players is high enough (i.e., $> \frac{x_m}{x_{m-1}+x_m}$), so that playing a_m is the payoff-dominant strategy.

Unfortunately, neither the theoretical nor the experimental literature on higher-order beliefs can tell us which of the two will be chosen. Even small variations in the theoretical models of higher-order beliefs can generate both equilibria. Table 2.1 shows us the predictions of a few common models of higher-order beliefs for the game as described in Section 2.3. The derivation of these predictions and a more detailed discussion can be found in Section 2.9.

From the table we can see that even the question if there is coordination in this game depends very much on the model of higher-order beliefs.

Hypothesis 1 (Coordination is possible). *High-cognition players use the first-best strategy a_m which is not available to the low-cognition players, despite the absence*

Model	Coordination	First-order belief coordination problems	Higher-order belief coordination problems
Common knowledge	Full coordination	No	No
Common p-belief	Full coordination	No	No
Global games	No coordination	Yes	Yes
Almost common knowledge	No coordination	No	Yes.

Table 2.1: Models of higher-order beliefs

of common knowledge about the cognitive types.

The next two hypotheses extend on Blume and Gneezy’s (2010) hypothesis that “beliefs matter”: Hypothesis 2 formalizes the question “Does coordination fail because some high-cognition players underestimate the fraction of high-cognition players?”.

Hypothesis 2 (First-order beliefs matter). *There is coordination failure due to first-order beliefs.*

Most of the problems in models of higher-order beliefs stem from the fact that there are infinitely many levels of beliefs. However, evidence from the laboratory indicates that people are not able to use higher-order rationality,⁸ a requirement for coordination problems due to higher-order beliefs. Furthermore, even in studies of level-k reasoning, where players are framed and incentivized on using higher-order beliefs, players still rarely use high-levels of reasoning.⁹

Hypothesis 3 (Higher-order beliefs matter). *There is coordination failure due to higher-order beliefs.*

In the following section we are going to explain the experimental design to test the three hypothesis.

⁸Kneeland (2015) shows that only about 22% of all players use more than third-order rationality.

⁹In Arad and Rubinstein’s (2012) 11-20 game, 80 % of the players only use 3rd-order beliefs or lower despite the game being designed to facilitate higher-order reasoning.




Figure 2.3: 5-sector-disc

2.5 Experimental design

Measuring higher-order beliefs is very complicated, as there is an “uncertainty principle” (as already discussed by Blume and Gneezy (2010)) at work; i.e., it is hard to measure beliefs without introducing or changing them.¹⁰ Introducing absence of common knowledge is difficult. When told that they are given a random number, subjects usually assume that it is drawn from a uniform distribution. Explicitly stating that the distribution is unknown leads to a myriad of other problems. Subjects could for example assume a strategic selection of the distribution by the experimenter. Furthermore, we need to have some sort of control over the fraction of high-cognition players, so that the action only available to the high-cognition players is the one with the highest expected payoff (see Section 2.9).

We solve all three problems by utilizing Blume and Gneezy’s (2000) 5-sector disc. This is a disc with 5 equally large sectors on it, 2 black and 3 white, as depicted in Figure 2.3.¹¹ The disc has the same sectors on the front- and backside of the disc and can be flipped and rotated. As the disc can be flipped, the subjects face symmetry constraints and can therefore not distinguish all five sectors. These symmetries cannot be overcome and therefore not all Nash equilibria are possible given the particular frame. Only certain “attainable” equilibria are possible, as defined originally in Crawford and Haller (1990), and further developed by Blume (2000) and Alós-Ferrer and Kuzmics (2013).

The property of this disc which is most important for this chapter is that it has a single distinct white sector: The sector adjacent to both black sectors (Figure 2.3).¹²

¹⁰Either by making the subjects realize that there might be something like a higher-order belief or by them trying to be a good subject (Orne (1962)).

¹¹There is a second version of this disc, with a significantly harder to find distinct sector, with adjacent black sectors. However, for this disc, the fraction of players who were able to identify the distinct sector is too small.

¹²More about the properties of this disc can be found in Blume and Gneezy (2000).

For the subjects there are then, in principle, three distinguishable sets of sectors: the black sectors (B), the uniquely identifiable white sector (D), and the other white sectors (W').

The key assumption behind the experiment (and also behind Blume and Gneezy (2000) and Blume and Gneezy (2010) and very much supported by their findings), is that not all subjects realize that there is a uniquely identifiable sector, which leads to two different cognitive types, the high type, who can identify the distinct sector, and the low type, who cannot. The low type then faces an additional symmetry constraint and has only two distinguishable sectors to choose from: One of the two black sectors (B) or one of the three white sectors (W).

The subjects then played three treatments in a random order without feedback after hearing and reading the instructions and completing an extensive quiz:¹³

The **Self Treatment** in which the subject gets the disc twice, every time randomly turned and rotated, and gets £5 if she picks the same sector twice.

In the **Prediction Treatment** one subject (she) is told that another subject (he) plays the *Self Treatment* (with a possibly differently turned and rotated disc). She has to pick one sector and every time he picks the sector she picked, she gets £2.5.

Finally, the **Coordination Treatment**, in which two players pick simultaneously a sector on a (randomly turned and rotated) disc and, if both players pick the same sector, both receive £5.

2.5.1 Hypotheses

In the Self Treatment a high-cognition player has 9 possible choices: She can pick any of three actions (D,B,W') in the first stage and then pick any of the three actions in the second stage. This decision problem for the high-cognition player has a unique optimal solution: pick the distinct sector twice, giving her a probability to win of 1.

A low-cognition player is only aware of four possible choices: He can pick *B* or *W* in the first stage and then pick *B* or *W* in the second stage. The low-cognition player also has a unique optimal choice: pick *B* in both stages, giving him a probability to win of $\frac{1}{2}$.

Therefore, we would expect a high-cognition player to choose the distinct sector twice and a low-cognition player to pick a black sector twice.

In the Prediction Treatment, the action taken by a subject should only depend on her type and her first-order belief about the type of the other player. A low-cognition player will always choose B, whereas a risk-neutral, high-cognition player should pick D

¹³For the complete instructions and a description of the quiz see the Sections 2.9.3 and 2.9.4.

if his belief that the other player is also of the high type is at least $\frac{1}{3}$ and B otherwise.¹⁴

The coordination treatment is best depicted as a bi-matrix game with three (for the high-cognition player) and two (for the low-cognition player) pure strategies, with winning probabilities as depicted in Figure 2.4 and Figure 2.5. We expect a low-cognition player to play B, as it is the payoff- and risk-dominant equilibrium, whereas a high-cognition player's choice depends on her belief hierarchy: If anywhere in her complete hierarchy a belief lower than $\frac{1}{3}$ (or $\frac{1}{2}$ for very risk averse players) that the other player is a high-cognition player or that the other player thinks that she is a high-cognition player, ... (or, in short, that there is no common-p belief of $\frac{1}{3}$ or higher, that both players are high-cognition players), she will choose B, otherwise she will choose D.

	W'	B	D
W'	$\frac{1}{2}$	0	0
B	0	$\frac{1}{2}$	0
D	0	0	1

Figure 2.4: High-cognition player winning probabilities

	W	B
W	$\frac{1}{3}$	0
B	0	$\frac{1}{2}$

Figure 2.5: Low-cognition player winning probabilities

We are using a within-subject design to test the hypotheses as stated in Section 2.4. In the following we will use a shorthand for players' strategies: "W'W' B D" means that a player selected one of the two white sectors twice in the Self Treatment, one of the black sectors in the Prediction Treatment and the distinct sector in the Coordination Treatment.

Using our design, we can reformulate the hypotheses as stated in Section 2.3:

Hypothesis 1 (Coordination is possible). *High-cognition players choose in the Coordination Treatment D more often than any other choice.*

The idea is straight forward: Only high-cognition players can identify the best equilibrium, so we don't have to consider other types. We can identify these players with the help of the the Self Treatment. If high-cognition players, i.e., the ones who have been able to identify "D" in the Self Treatment, coordinate on D in the Coordination Treatment we know that coordination is possible, even in the absence of common knowledge.

The second question we want to answer is, if pessimistic beliefs about the other players'

¹⁴Allowing for risk-averse players, this fraction has to be between $\frac{1}{3}$ and $\frac{1}{2}$, depending on the degree of risk aversion.

types can lead to coordination failure.

Hypothesis 2 (First-order beliefs matter). *There are high-cognition subjects who choose a black sector in the Prediction Treatment and Cooperation Treatment, i.e., play “DD B B”.*

We already know that we can identify players’ types with the help of the Self Treatment. Furthermore, the Prediction Treatment identifies players who think that more than $\frac{1}{3}$ of the other players can not identify the distinct sector.

Hypothesis 3 (Higher-order beliefs matter). *There are high-cognition subjects who play the distinct sector in the Prediction Treatment and a black sector in the Cooperation Treatment, i.e., play “DD D B”.*

Our design allows for another robustness check: There is an attainable strategy which is very similar to the one we use to identify first- and higher-order beliefs: “DD B D”. This strategy will only be chosen if players believe that their partner is of the low type, but still plays “D” in the in the Coordination Treatment. This strategy can therefore not be explained using our model.

Hypothesis (Robustness check). *“DD B D” is played less often than “DD B B” and “DD D B”.*

2.6 Results

The experiment was conducted at the DR@W Laboratory at the University of Warwick using the experimental software "z-Tree" developed by Fischbacher (2007). 130 subjects were recruited and received payments between £3 and £18. Before showing the results, let us briefly discuss the preliminaries of the experiment design.

The first preliminary is the focality of the distinct and the two black sectors. From the choice data in Figure 2.6 we can see that more than 95% of all players have chosen one of these sectors in the Coordination Treatment. The second preliminary is that there are enough high-cognition players, so that playing the high-cognition exclusive action is a payoff-dominant equilibrium for the players. In Figure 2.7 we can see that 58% of all players have chosen the distinct sector and are therefore considered high-cognition players. Therefore, playing the distinct sector would maximize the expected utility of high-cognition players in a game with common knowledge about the type distribution, independently of the degree of risk aversion (see Section 2.9).

These results are in line with Blume and Gneezy’s (2010) results where around 52% (58% in our experiment) have been able to identify the distinct sector and around 23% (34%) have chosen the black sector. We contribute the significantly lower level of




Figure 2.6: Results of the Coordination Treatment

noise (8% vs 25%) to the extensive instructions and the quiz we conducted before the experiment.¹⁵

Due to the lower level of noise we are, unlike Blume and Gneezy (2010), able to use a within-subject design, in which each player has access to 625 possible strategies.¹⁶ Of these strategies we consider 96.32% as “noise”.¹⁷ As the number of strategies which support our hypothesis are very low (1, 4 and 2 out of 625), the probability that someone chooses them by mistake is very low. For a detailed overview of all possible strategies and how we categorize them see Table 2.2.

Given the preliminaries, we can test hypotheses 1 through 3.

Hypothesis 1 (Coordination is possible). *High-cognition players choose in the coordination treatment D more often than any other choice.*

The choice data from our experiment confirms this hypothesis. In Figure 2.8 we can see that 80% have chosen the strategy “DD D D”. As this strategy represents only 0.16% of all available strategies (or 4% when excluding the Self Treatment), we can reject the null hypothesis of this high level of coordination being a result of random

¹⁵For the instructions and an overview of the quiz see the Sections 2.9.3 and 2.9.4.

¹⁶We are here ignoring the order in which treatments are played.

¹⁷This noise includes not only players not understanding the experiment or behaving randomly but also “Eureka”-learning (which was a big problem in Blume and Gneezy (2010), see Section 2.9.1), making a mistake (e.g., picking a not distinct white sector instead of the distinct sector, a mistake, which both of the authors made multiple times while testing the experiment) and beliefs of low-cognition players.




Figure 2.7: Results of the Self Treatment

Description	Hypothesis	# of strategies	Proportion
DD D D	1: Coordination is possible	1	0.16%
DD B B	2: First-order beliefs matter	4	0.64%
DD D B	3: Higher-order beliefs matter	2	0.32%
BB B B	(Low-cognition players)	16	2.56%
“Noise“	-	602	96.32%
WW-W-W	(part of “Noise“)	80	12.80%

Table 2.2: Overview of the strategies

play ($p < 0.00001$).

Blume and Gneezy (2010) claim that “beliefs matter“ and we test in Hypothesis 2 if there are subjects whose pessimistic beliefs about the other players’ skills lead to coordination failure.

Hypothesis 2 (First-order beliefs matter). *There are high-cognition subjects who choose a black sector in the Prediction Treatment and Cooperation Treatment, i.e., play “DD B B”.*

Our data confirms this hypothesis. Figure 2.9 shows us the results of all players, Figure 2.10 of the high-cognition players. In these figures we can see that about 9% of the high-cognition players (or 5% of all players) have a first-order belief problem, leading to coordination failure. As the fraction of strategies leading to this conclusion




Figure 2.8: Results of the Coordination Treatment (high-cognition players)

is very small (0.64%) we can reject the null hypothesis that this result is due to chance ($p < 0.00001$).

But do players really use higher-order beliefs in this type of games? Hypothesis 3 tests for this question.

Hypothesis 3 (Higher-order beliefs matter). *There are high-cognition subjects who play the distinct sector in the Prediction Treatment and a black sector in the Cooperation Treatment, i.e., play “DD D B”.*

From Figure 2.9 and Figure 2.10 we can see that there are high-cognition players who think that their partner is with a high probability of the high type, they, however, still think there are coordination problems. Again, we can reject the null hypothesis at the 1% level ($p < 0.00001$).

Hypothesis (Robustness check). *“DD B D” is played less often than “DD B B” and “DD D B”.*

All these results are statistically significant at the 1% level, however, our design allows for another robustness check: There is a strategy which should not be played by rational players: “DD B D”, which is about as likely to be picked at random as “DD B B” and “DD D B” but can not be explained by our model. Figure 2.11, shows us that only 2 subjects have chosen this strategy.

We expected to have significant order effects, as in Blume and Gneezy (2010). However, it turns out, that the only robust order effect is a weak effect in the Self Treatment

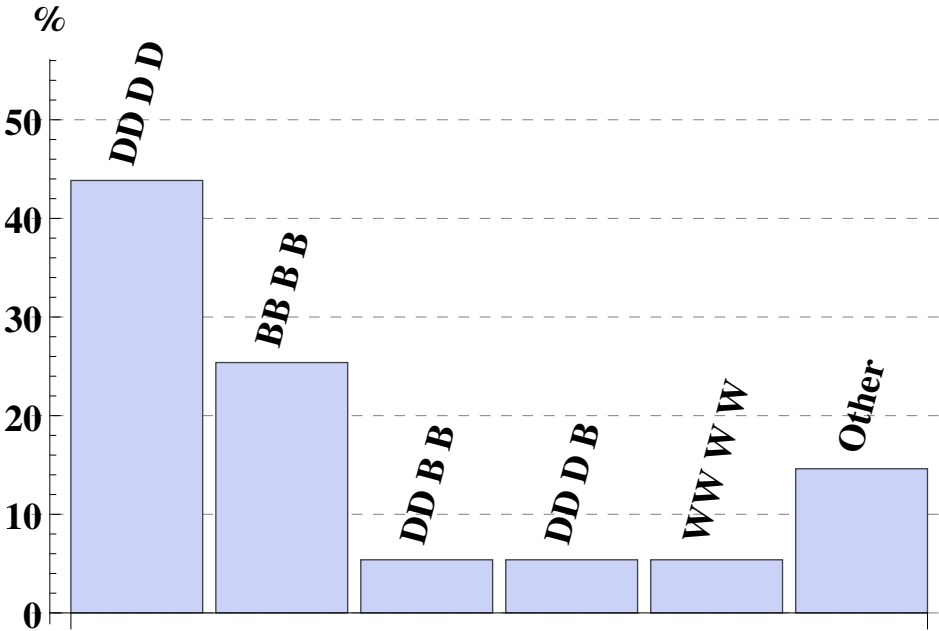


Figure 2.9: Used strategies




Figure 2.10: Used strategies (high-cognition players)




Figure 2.11: Robustness check

(i.e., more subjects have been able to choose the distinct sector twice later in the experiment).¹⁸ We attribute this to a small change in design. We have explained every treatment before the experiment started and we have conducted a quiz (Section 2.9.4), testing if the instructions have been understood. This probably led to "Eureka learning" before instead of during the experiment.

2.7 Conclusion

We have seen that, in this game, absence of common knowledge was not enough to prevent coordination on the Pareto-optimal equilibrium, as 76% of the high-cognition players have chosen the Pareto-optimal equilibrium. However, we still have a fraction of players who have beliefs that lead to coordination failure (around 18%) and of these only half could be attributed to first-order beliefs.

Of the models of higher-order beliefs discussed in Section 2.4 and Section 2.9, only "assuming common knowledge" or a common p -belief were able to explain coordination on the payoff-dominant equilibrium. However, these assumptions can not explain any coordination failure due to beliefs, as the beliefs are fixed by the model, whereas the models which can explain this type of coordination failure predict no coordination on the payoff-dominant equilibrium.

¹⁸For the full analysis of order effects see Section 2.9.1.

Therefore, as we have observed a coordination rate of about 76%, assuming common knowledge (or a common p-belief) might be the best tractable approximation available in coordination games without common knowledge, depending on the focus of the research.

But coming back to the example mentioned in the introduction, ignoring the higher-order beliefs can have severe negative consequences. Our results can explain why education without considering problems due to higher-order beliefs can have significant effect but they can also explain why NGOs like Tostan have significantly more success. Furthermore, these results give reason to believe that just explaining if others have been educated and are against female genital mutilation (i.e., changing the first-order beliefs) might not be sufficient and making this education common knowledge might be necessary to achieve all possible benefits from it.

However, the results from this experiment conducted with students at a European university should of course not be generalized to explain behavior in small rural communities without further research but just gives us reason to believe that higher-order beliefs do matter.

This opens up some questions for future research: Can these results be generalized to other populations? Are there certain parts of the populations who are more likely to exhibit first- or higher-order beliefs which lead to coordination failure? Are there other, maybe easier methods to make something common knowledge? Furthermore, it might be worthwhile to check more general structures of higher-order beliefs or if non-equilibrium models like Strzalecki (2014) or Kneeland (2016) can explain this phenomenon better.

2.8 Appendix

2.8.1 Belief hierarchies

Let $B_i^0 := T_j$ and $B_i^k = T_j \times \Delta(B_i^{k-1})$ with $\Delta(B)$ being the space of probability measures on B and $\Delta(X)$ being the space of probability measures on the Borel field of X , endowed with the weak topology. Using this notation, we can define a belief hierarchy as follows.

Definition 3 (Belief hierarchy). A *k-th order belief* is defined as

$$b_i^k \in \Delta(B_i^k)$$

with $B_i^0 = T_j$ and $B_j^k = T_j \times \Delta(B_j^{k-1})$. Furthermore, let us set $b_i^0 := t_i$. A belief hierarchy of a player i is then $b = \{b_i^0, b_i^1, \dots\}$

We therefore have a first order belief $b_i^1 \in \Delta(\{low, high\}) = [0, 1]$ and higher-order beliefs $b_i^k \in [0, 1]^k$.

Furthermore, we assume these beliefs to be coherent, i.e., that beliefs of different orders do not contradict one another,¹⁹ and that a low-cognition type does not know about higher cognitive types, i.e., $b_i^k = 0 \Rightarrow b_i^{k+1} = 0 \quad \forall k \geq 0$.

This excludes, on the one hand, that a low-cognition player thinks that the other player is a high-cognition player and, on the other hand, that a player has a first-order belief that the other player is of a the high type and a higher-order belief that the player is of the low type.

2.9 Equilibrium selection and models of higher-order beliefs

In this section we are going to discuss how different models of beliefs and frequently used assumptions on the structure of higher-order beliefs influence the specific game we analyze.

Using the results from the literature on focal points in coordination games (as discussed in Section 2.1.1), we know that we can restrict our attention on the two actions with the highest payoffs a_{m-1} and a_m . This simplifies the game to a Bayesian game with two types, a low type whose only attainable action is a_{m-1} and a high type, who has access to a_{m-1} and a_m , without common knowledge about the type distribution. Then, we can denote, with a small abuse of notation, the strategy of a player as the action she chooses if she is of the high-type, i.e., a_m or a_{m-1} , knowing that she will play a_{m-1} if she is of the low type.

Let us first start with the most common assumption, that the *distribution of types is common knowledge*. Then the expected utility of a (risk neutral) high-cognition player is as depicted in Table 2.3, given her and her partners strategies.²⁰ p denotes the probability of a player being of the high type. We can see that the prediction of the model then depends on p . If the probability of a player being of the high type p is too low ($p < \frac{x_{m-1}}{x_{m-1}+x_m}$), only (a_{m-1}, a_{m-1}) will be an equilibrium. In this chapter we are going to assume that $p \geq \frac{x_{m-1}}{x_{m-1}+x_m}$ which makes sure that the “better” equilibrium

¹⁹I.e., higher-order beliefs of a player mapped onto the space of beliefs of a lower order are the same.

²⁰In the analysis we restrict our attention to risk-neutral players. However, the analysis for the case of risk-averse players is analogous and the experimental results are valid for every possible degree of risk aversion.

always exists.²¹ For risk-averse players, it is required that $p \geq \frac{u(x_{m-1})}{u(x_{m-1})+u(x_m)}$, so we know that as long as $p \geq \frac{1}{2}$ the high-type equilibrium always exists, independently of the degree of risk aversion. Furthermore, if the equilibrium exists, it is payoff dominant.

	a_{m-1}	a_m
a_{m-1}	x_{m-1}, x_{m-1}	$(1-p)x_{m-1}, 0$
a_m	$0, (1-p)x_{m-1}$	px_m, px_m

Table 2.3: Expected utilities of two high-cognition players

Therefore, the prediction of assuming that the *distribution of types is common knowledge* is that, for a high-enough p , we should expect full cooperation.

Monderer and Samet’s (1989) *common p -belief* is a generalization of the concept of common knowledge and generates, in this model, the same predictions as assuming that the distribution of types is common knowledge, given a high-enough p .

The game we are analyzing is very close to the original description of a global game as introduced by Carlsson and Van Damme (1993). Written down as in Table 2.3 it is a very similar game as the main example used in Carlsson and Van Damme (1993). Therefore, we know that, given $\frac{x_{m-1}}{x_m} \leq p \leq \frac{2x_{m-1}}{x_m+x_{m-1}}$ (i.e., (a_m, a_m) is still a Nash equilibrium but (a_{m-1}, a_{m-1}) is risk dominant), (a_{m-1}, a_{m-1}) will be the only rationalizable solution to the global game. Furthermore, Hellwig (2002) shows that higher-order uncertainty about preferences leads to results similar to Carlsson and Van Damme’s (1993) higher-order uncertainty about payoffs, i.e., coordination on the "less risky" equilibrium.

Rubinstein (1989) shows that truncating common knowledge at any finite level is equivalent to the situation without any common knowledge at all and therefore suggests that players choose the save strategy a_{m-1} .

Weinstein and Yildiz (2007a) establish a condition, called “global stability under uncertainty” which implies that the change in equilibrium actions is small in the change of k th-order beliefs and higher. Therefore, under this condition, equilibria can be approximated by the equilibrium with lower-order beliefs. Unfortunately, pure coordination games do, in general, not fulfill the conditions for “global stability under uncertainty” as the best responses are very sensitive to every order of beliefs and even a small change in some higher-order belief might make a player change from a_m to a_{m-1} .

²¹In the experiment this assumption requires $p > \frac{1}{3}$. As the fraction of high-cognition players is 58%, this assumption is not problematic.

Model	Coordination	First-order belief coordination problems	Higher-order belief coordination problems
Common knowledge	Full coordination	No	No
Common p-belief	Full coordination	No	No
Global games	No coordination	Yes	Yes
Almost common knowledge	No coordination	No	Yes.

Table 2.4: Models of higher-order beliefs

2.9.1 Order effects

In the introduction we have briefly discussed an uncertainty principle, in which higher-order beliefs can not be measured without inducing them. This theory is related to the “good subjects hypothesis“ (Orne (1962)) according to which some subjects try to figure out the research question and then change their behavior to confirm said hypothesis. However, in this case the difference is more subtle: As soon as they realize that there is a higher-order belief problem, they might overestimate it.

Blume and Gneezy (2010) have encountered a different case of this uncertainty hypothesis. “Having a player play against himself may trigger an insight that switches a player from low to high cognition (“Eureka!” learning). There may be an uncertainty principle at work here in that we cannot measure a player’s cognition without altering it.“ (Blume and Gneezy (2010)) This suggests, that the order of treatments might be important. Therefore, we implemented a random order. However, it turns out that we have (almost) no order effect, as can be seen in Table 2.5. The only statistically significant effect is that, if the self treatment was the first treatment, there was a significantly higher number of “Other“ results than when it was the second ($p = 0.0062$) or third treatment ($p = 0.0139$). Furthermore, the distinct sector was played more often in the coordination treatment if it was the second than the first treatment ($p = 0.0277$), however, there were no significant effects when comparing the first and third and the second and third.²²The former has a intuitive explanation (i.e., practicing the task makes it less likely to make a mistake) whereas the later is considered to be a type II error by the authors.

²²Using the one-tailed Fisher’s exact test.

Treatment	Self			Prediction			Coordination		
Order	DD	BB	Other	D	B	W	D	B	W
1st	18	10	8	23	15	3	22	22	3
2nd	32	24	2	20	16	2	24	10	0
3rd	25	10	1	32	15	4	23	15	5

Table 2.5: Order effects of the different treatments

The question now is, why did Blume and Gneezy (2010) encounter strong "Eureka!"-learning effects whereas we had (almost) no significant effect. The authors attribute this to the fact that we used more extensive instructions and a quiz to make sure the instructions were understood. More importantly, the participants were instructed in all three treatments before they played the first game which most likely triggered the learning before the first decision, whereas in Blume and Gneezy (2010) the instructions for the second treatment were distributed after completion of the first treatment.

2.9.3 Instructions

Welcome to this experiment in economic decision making. It will take approximately 60 minutes. First of all, please check that the number on the card handed to you matches the number on the cubicle that you are seated in and that your mobile phones are turned off.

Before we start, we will explain the rules of this experiment. You will also find these rules on the paper provided, so you can read along and check again during the experiment. If you have any questions, please do not speak up but raise your hand and we will come to you and answer your question privately.


From now on, please do not talk, and listen carefully. In this experiment you will earn a minimum of £3, and potentially up to £18. How much money you earn will depend on your decisions and those of the other participants. Your reward will be paid out at the end of the experiment. None of the other participants will know how much money you made.

In this experiment you will be asked to make decisions related to a disc that has 5 sectors, similar to the disc provided to you. The disc has two identical sides. Your goal will be to pick the same sector twice (more on that later). During this experiment the disc will be flipped and/or rotated randomly.

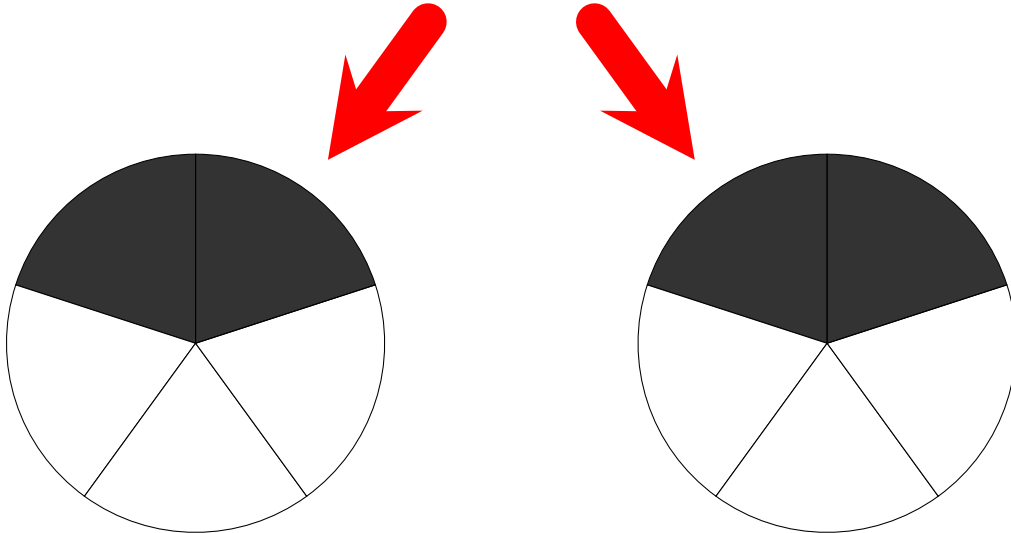
Pictures on page 2 illustrate rotation and flipping. Since you will not be told if the disc was flipped and/or rotated, it might even be the case that disc looks exactly the same though sectors have changed their positions.

The arrow tracks one specific sector that changes its position as the disc is rotated and/or flipped.

This is an example of rotating the disc by two sectors:



This is an example of flipping the disc:




In the experiment the disc will be surrounded by the letters A, B, C, D, and E. **These labels are not part of the disc!** They are only included to allow you to choose a sector.

In the experiment you will make decisions in the following environments (the order will be chosen randomly):

(Self Game) You will be asked to pick a sector twice; first you choose a sector; then the disc might be flipped and/or rotated. After this you are shown the same disc and have to choose a sector again. You will not observe the flipping/rotation of the disc. If you manage to guess the same sector twice, your payoff will be £5. Otherwise, you will receive 0. Therefore, to earn more money you want to maximise your chances to pick the same sector twice.

Here is an example of the choices made in a **Self Game**, using a simpler disc with only 2 instead of 5 sectors:




First you picked the black sector; then you picked the black sector again. Therefore, you pick the same sector twice and earn £5.

(Prediction Game) You are matched randomly with another person and you have to guess the choice of this person, while she plays the **Self Game**. First, you choose

a sector on the disc; each time the other person picks the sector you chose, you will receive £2.5. As the other player picks twice in the **Self Game**, you can earn £0, £2.5 or £5 in this situation, depending on your and the other person's choice. Therefore, to earn more money you want to guess what the other player is playing in the **Self Game** described above.


Here is an example of the choices made in a **Prediction Game**, again with the simpler disc:



First you picked the black sector. The other player then plays the **Self Game**. He first picks the black sector and therefore you earn £2.5. Then he picks the white sector and therefore you earn £0. Thus you earn £2.5 in total.

(**Coordination Game**) You are matched randomly with another person and both of you are asked to pick a sector on the disc simultaneously. Both of you know that you play the **Coordination Game**. You both see the same disc but possibly differently flipped and rotated. If both of you pick the same sector, then your payoff will be £5. Otherwise, you will receive £0. Therefore, to earn more money you want to guess the sector your partner is picking here, while he is trying to do the same.

Here is an example of the choices made in a **Coordination Game**, again with the simpler disc.



You picked the black sector. The other player picked the white sector. You therefore failed to coordinate and both of you earn £5 each.

The experiment consists of two periods. Each period consists of the three games as described above, using a 5-sector disc; the order of the games is random. At the end of the experiment one of the two periods will be randomly chosen. The earnings made in this period will be paid out in cash.

Again, please do not talk during this experiment! If you have questions just raise your

hand.

Before the experiment there will be a quiz to check your understanding. Read hints carefully if you get stuck during the quiz.

2.9.4 Quiz

In this appendix you can find screenshots of the quiz which was conducted before the experiment. Participants who made a mistake in some part of the quiz were given a small hint and then were asked to repeat this part of the quiz.

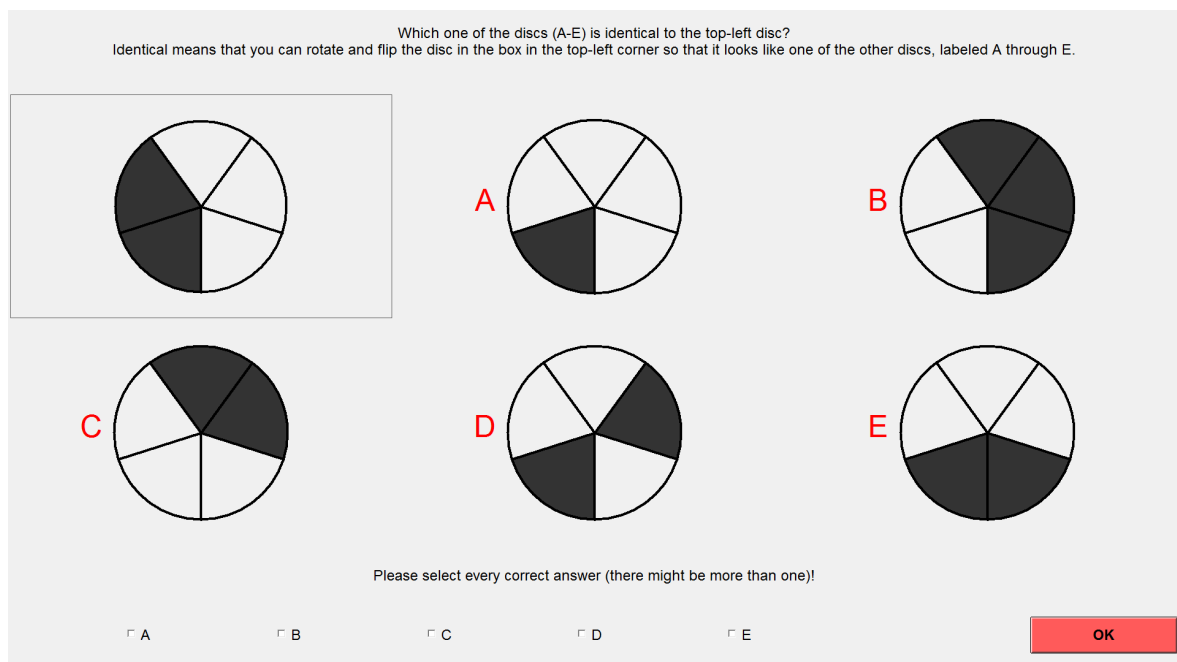


Figure 2.12: Quiz part 1

Which one of the discs (A-E) is identical to the top-left disc?
 Identical means that you can rotate and flip the disc in the box in the top-left corner so that it looks like one of the other discs, labeled A through E.
 A red arrow points to a sector you or your partner have chosen before and also indicates the same sector after the disc was rotated and/or flipped. This chosen sector is therefore also part of the disc and moves whenever the disc rotates and flips.

Please select every correct answer (there might be more than one)!

A B C D E

Figure 2.13: Quiz part 2

Assume you are playing the **Self Game** and you have chosen the same sector twice.

1) How much Pound will you get for this game?

Assume you and a random partner are playing the **Coordination Game** together and both of you have each picked a **different sector** of the same color.

2) How much Pound will you get for this game?

3) How much Pound will he get for this game?

Assume you are playing the **Guessing Game**. The participant you are observing is therefore playing the **Self Game**. He picked the same sector twice but you failed to pick the sector he picked.

4) How much Pound will you get for this game?

5) How much Pound will he get for this game?

Now imagine the same situation but this time you managed to pick the same sector he picked twice.

6) How much Pound will you get for this game?

7) How much Pound will he get for this game?

Figure 2.14: Quiz part 3

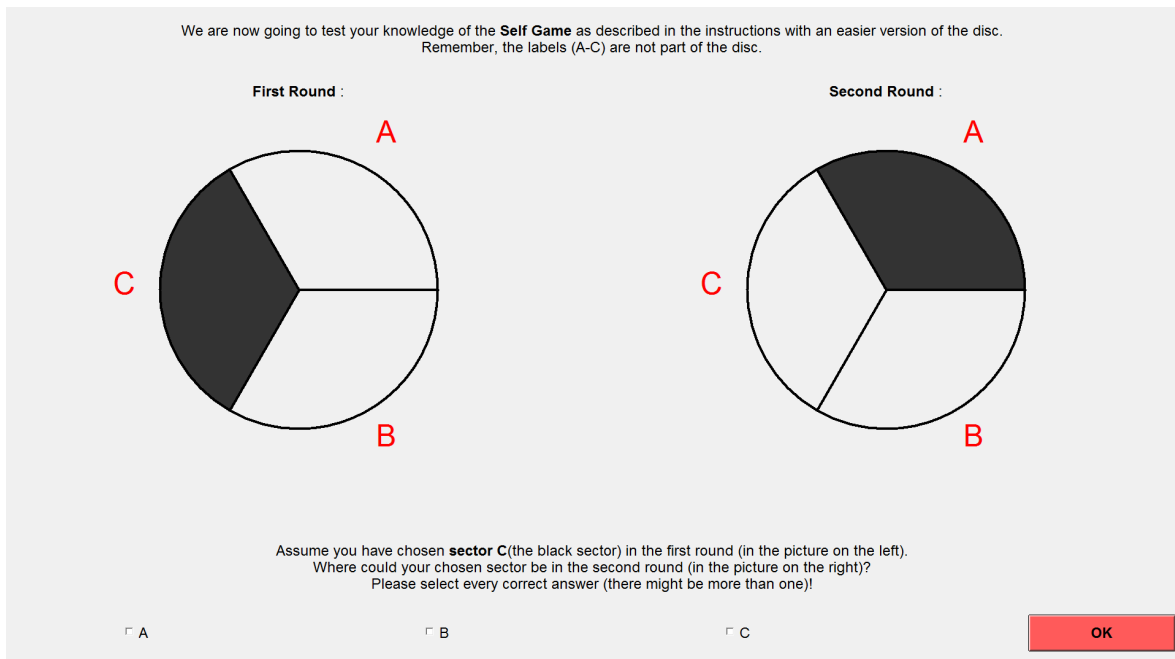


Figure 2.15: Quiz part 4


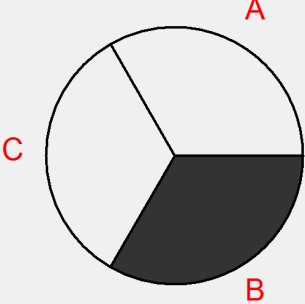


Figure 2.16: Quiz part 5

We are now going to test your knowledge of the **Guessing Game** as described in the instructions with an easier version of the disc. Assume you have observed the choices of the other player (which will not be possible in the experiment) and he has chosen a black sector twice).

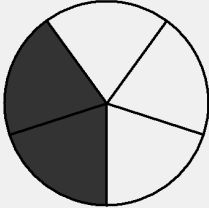


Now you are presented this disc and you are playing the **Guessing Game**.

<p>Assume you have chosen sector A (a white sector). What are your possible payoffs for this round?</p> <ul style="list-style-type: none"> <input type="radio"/> 0 Pounds <input type="radio"/> 2.5 Pounds <input type="radio"/> 5 Pounds 	<p>Now assume you have chosen sector B (the black sector). What are your possible payoffs for this round?!</p> <ul style="list-style-type: none"> <input type="radio"/> 0 Pounds <input type="radio"/> 2.5 Pounds <input type="radio"/> 5 Pounds
--	---

Figure 2.17: Quiz part 6

We are now going to test your knowledge of the **Coordination Game** as described in the instructions. Assume you are playing this game with the five sector disc below.

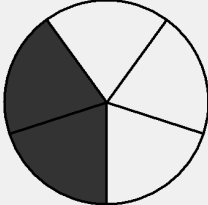


Assume you know that your partner has chosen a **black sector**.

<p>What is the probability of winning if you choose a black sector?</p> <ul style="list-style-type: none"> <input type="radio"/> 0% (no chance) <input type="radio"/> 20% (a 1 in 5 chance) <input type="radio"/> 33% (a 1 in 3 chance) <input type="radio"/> 50% (a 1 in 2 chance) <input type="radio"/> 100% 	<p>What is the probability of winning if you choose a white sector?</p> <ul style="list-style-type: none"> <input type="radio"/> 0% (no chance) <input type="radio"/> 20% (a 1 in 5 chance) <input type="radio"/> 33% (a 1 in 3 chance) <input type="radio"/> 50% (a 1 in 2 chance) <input type="radio"/> 100%
---	---

Figure 2.18: Quiz part 7

We are now going to test your knowledge of the **Coordination Game** as described in the instructions. Assume you are playing this game with the five sector disc below.



Assume you know that your partner has chosen a **white sector**.

What is the probability of winning if you choose a black sector?

- 0% (no chance)
- 20% (a 1 in 5 chance)
- 33% (a 1 in 3 chance)
- 50% (a 1 in 2 chance)
- 100%

What is the probability of winning if you choose a white sector?

- 0% (no chance)
- 20% (a 1 in 5 chance)
- 33% (a 1 in 3 chance)
- 50% (a 1 in 2 chance)
- 100%

OK

Figure 2.19: Quiz part 8

Chapter 3

Rational Delay of Effort in Projects with Uncertain Requirements¹

3.1 Introduction

You start doing it in school with your homework, continue while writing a term paper in college and are probably still doing it when you have to do your taxes: You postpone working on it until the very last minute, despite having a deadline. This phenomenon is not restricted to work you conduct on your own, sometimes it is even stronger when you work in a team.

Naturally this causes problems, not only for you, but for the whole team. In this chapter I will focus on project work, i.e., working together towards a fixed goal after which your team will be terminated. Project work is generally said to be more efficient and is frequently used in the workplace (Harvard Business School Press (2004)) and the classroom (Hutchinson (2001)).

Another important part of managing projects, apart from teamwork, is requirements management. According to a survey by Taylor (2000) unclear objectives and requirements are the most common cause for failure of IT projects. In this chapter, I am trying to establish a connection between uncertainty in the requirements of a projects and the often observed last-minute rush, in which workers delay much of the required work until the very end of the project.

This work proposes a continuous-time model of working in projects, which explains delaying effort not only in teams but also when working alone, not as a result of inefficiency or time-inconsistency but as an efficient, team-value maximizing consequence

¹Parts of this chapter can be found in Külpmann (2015).

of a deadline. The range of applications is quite broad, from large scale multi-national research projects to a single worker trying to write a report.

The main features of this model are:

- **public benefits**, which are realized upon completion of a project in the form of a lump-sum payment,
- **private costs**, which are assumed to be quadratic,
- an **unknown threshold for success** or uncertain requirements, with a commonly known distribution,²
- **unobservable efforts**, so only the player's own effort is known,
- and a **deadline**, after which the project cannot be completed anymore.

In the model, the players exert effort over time until either the deadline is reached or the project is successful. While doing so, they only know that they have not been successful yet. Projects in this model are described by the assumed distribution of the breakthrough effort. This *breakthrough-effort distribution* can cover many different projects, e.g., projects in which only the current effort influences the probability of success or projects during which players learn about the quality of the project while trying to complete it. One simple example for a breakthrough-effort distribution is the uniform distribution on $[\underline{e}, \bar{e}]$. This means that the players think that the project needs effort between \underline{e} and \bar{e} to be completed. For examples of different types of projects and the corresponding breakthrough-effort distributions see Section 3.3.

I find that in the equilibrium there are three different effects at work: **free-riding**, which reduces the overall effort the more players are working on the project. The second effect is **encouragement**, which depends on the threshold distribution: Given a decreasing hazard rate my own effort encourages the other players to work less, while given an increasing hazard rate, my work encourages my coworkers to work more in the future. The last effect is **delay of effort**, which causes players to work later rather than earlier, even with the presence of a discount rate which lets players want to have a breakthrough as soon as possible.³

Free-riding is a common effect in moral hazard problems and already well understood. Encouragement also occurs frequently in the literature, however usually either only as a positive encouragement effect (for example in Georgiadis (2014)) or only as a negative

²Which we will call *breakthrough-effort distribution*.

³This effect is more than just a consequence of discounting, as the discount rate does not only effect the costs but also the benefits. As the benefits are, by design, later than the costs and higher than the expected costs a discount rate lets players work earlier.

encouragement effect (as in many bandit models).⁴ The positive encouragement effect is very similar to strategic complementarity and the negative encouragement effect to strategic substitutability. My model, however, can incorporate both effects and the occurrence of positive or negative encouragement depends on the type of the project (i.e. the breakthrough-effort distribution).

Delay of effort as a result of rational players, has, to the best of my knowledge, not been analyzed before in this context. In this model it is caused by convex costs, which make it optimal for the players to spread their effort as evenly as possible, a deadline and uncertainty about the effort required for a breakthrough. When they start working on a project, the players have a belief about the threshold that includes very low effort levels and they are trying to find an optimal effort level given this belief. If they do not succeed at first, they realize that the threshold effort level is not that low and that they have to update their beliefs about the threshold. Therefore, they also have to increase their effort level to reflect the updated beliefs. Hence, we can expect some delay even with a rational social planner trying to maximize social welfare. Close to the deadline, the effect even outweighs any other effect, including encouragement and even strong discounting. Therefore, we can observe a *last-minute rush*.

These results have implications on the evaluation of projects: Not only should managers avoid the negative encouragement effect but, using a positive encouragement effect, they might be able to counteract the ever-present incentives to free-ride. As this encouragement effect depends on the type of requirements uncertainty it might be possible to switch from a negative effect to a positive effect by resolving some of these uncertainties or sometimes even by introducing new uncertainty into the project's objectives.

Another point to take away from the results of this chapter is the occurrence of a last-minute rush in the welfare-maximizing solution. This might have further implications on the evaluation of the often observed increased workload around deadlines. This effect is usually dismissed as a result of (irrational) procrastination. This chapter, however, shows that it might not only be a rational consequence of unclear objectives but possibly even welfare maximizing.

This chapter is organized as follows: In the next section, I will give an overview of the relevant literature and how this work fits into it. Then, I will explain the model in Section 3.2, followed by an explanation of the breakthrough-effort distribution and how it translates to different projects in Section 3.3. In Section 3.4, I derive the optimal effort for the non-cooperative and the welfare maximizing case. Additionally, I show that there is a last-minute rush, i.e., delay of effort which leads to a peak of effort towards

⁴One example is Bonatti and Hörner (2011) in which the negative encouragement effect is the unnamed effect leading to procrastination in their model.

the end in both cases. Then I will briefly discuss the effect of deadlines (Section 3.5.1), the model without discounting (Section 3.5.3) and how different types of projects react to changes in costs and changes in the quality of the project (Section 3.5.5). I will give some concluding remarks in Section 3.5.6.

3.1.1 Literature

This chapter is related to different fields of the literature: Holmstrom (1982) started the game theoretic literature on *moral hazard in teams*, which was then expanded by Ma, Moore, and Turnbull (1988), Legros and Matthews (1993) and Winter (2004), to mention only a few important contributions. A common theme is the focus on free-rider problems due to shared rewards but costly private effort. My work adds to this literature as it analyzes a dynamic moral hazard problem, in which players have very restricted information about the actions of others, which leads to free-riding and a delay of effort.

In parts, this model is related to the literature on *strategic experimentation* as it models the behavior of players who optimize their decisions while gathering information at the same time. In these games every player has to divide her time between a “safe” and a “risky” action (as in the arms of a two-armed bandit) with unknown but common payoffs. Bolton and Harris (1999) analyze a two-armed bandit problem with many players in which the arms yield payoffs which behave like a Brownian Motion, with different drifts for the safe and the risky arm. They characterize the unique symmetric Markov Perfect equilibrium and are able to identify free-rider and encouragement effects. In Keller, Rady, and Cripps’s (2005) model of strategic experimentation, in which the risky arm yields a lump-sum with a certain intensity if the the risky arm is good and nothing if the risky arm is bad, new information arrive as a Poisson process, as in most of the recent literature on bandit problems. Two examples for this literature are Klein and Rady (2011), where the risky arms are negatively correlated, and Klein (2013) who extended the model to three armed bandits.

My work is very closely related (and was inspired by) Bonatti and Hörner (2011).⁵ They analyze a bandit model, similar to Keller, Rady, and Cripps (2005), in which efforts are private information and only outcomes are observable. After a success the game ends and payoffs are realized.

The model presented in this chapter is a very particular model of strategic experimentation: Not only is the information a player gathers about the actions of the other players very restricted, but furthermore, players’ payoffs are perfectly correlated. However, models of strategic experimentation usually assume the news arrival to be

⁵In fact, their benchmark model is a special case of my model with an incomplete exponential distribution and linear instead of quadratic costs. See Example 2 for more information.

a Poisson process, whereas my model has hardly any restriction on this news arrival process.

Bonatti and Hörner (2011) also exemplifies another strand of literature which this chapter is related to: *dynamic contribution games*. Already suggested by Schelling (1960), these models analyze the dynamic contributions to public goods. Admati and Perry (1991) and Lockwood and Thomas (2002) are examples for games in which contributions are observable. In Georgiadis (2014), there is no uncertainty about the valuation of the public good but about how effort affects the provision of the public good. He assumes that effort affects the drift of a standard Brownian motion towards a (commonly known) threshold and is able to not only identify free-riding and encouragement effects, but also to show that the optimal contract only compensates on success. Although in this chapter the uncertainty is about the threshold and not about the effect of effort, these two models are closely related when hazard rates are increasing as shown in Example 3. My work introduces uncertainty about the effort needed to provide the public good. Therefore, players also have to incorporate information gathering into their decision process. Furthermore, I show that, due to the presence of a deadline, delaying effort is optimal.

There is a huge literature on *procrastination in economics and psychology*. However, these usually attribute procrastination to self-control problems (O’Donoghue and Rabin (2001)) or time-inconsistencies like hyperbolic discounting (Laibson (1997)). Another explanation for procrastination is given by Akerlof (1991): According to him, procrastination is a consequence of “repeated errors of judgment due to unwarranted salience of some costs and benefits relative to others“ (Akerlof, 1991, p. 3).

The literature on procrastination in psychology is much more prominent than in economics but, like the economic literature, it almost exclusively focuses on some form of cognitive biased decisions (e.g., Wolters (2003) or Klingsieck (2015)).

This chapter adds to this literature, as it models not only decision processes of a single person but also delayed effort in teams, i.e. in a game-theoretic model. Furthermore, it provides an explanation for observed procrastination as rational and even welfare-maximizing behavior and therefore gives rise to completely different measures that should (or should not) be taken.

Bergemann and Hege (2005) use a very similar information structure to the one presented in this chapter, but analyze a problem in discrete time with linear costs and a memoryless investment. Second, Khan and Stinchcombe (2015) analyze decision problems in which changes can occur at random times and require a costly reaction. They have identified situations in which delayed reaction is optimal, depending on the form of the hazard rate of the underlying changing probability distributions. The relation-

ship of this model to the latter paper is mostly in the use of the hazard rate as a description of the projects players are working on.

To summarize, this chapter contributes to the literature in two different ways: On the one hand, it provides a tractable model to analyze a very general class of dynamic contribution games in continuous time with many players and incomplete information about effort contribution. The model deals with very different types of projects: Projects in which the success probability decreases in effort already spent,⁶ e.g. through learning about the quality of the project (which is very common in bandit models), investment projects similar to Georgiadis (2014) where the past effort increases the chance of success now and even projects in which past effort increases chance of success on some intervals and decreases on others.

Furthermore, this model can explain situations in which delaying effort is not only rational (which was also observed in Bonatti and Hörner (2011))⁷ but even welfare maximizing and arises without the assumption of time-inconsistencies or cognitive biased players. In addition to this, I was able to identify a strategic encouragement effect which can be beneficial or harmful to the projects success, depending on the type of uncertainties the players are facing.

3.2 The model

Consider n risk neutral players working together on a project in continuous time $t \in [0, T]$. Players can only observe their own past effort and whether the project was successful. After a success the players get a lump sum payment, normalized to 1, and the game ends. Every player i chooses at every point in time t whether to exert a costly effort $u_i : [0, T] \rightarrow \mathbb{R}_+$ with quadratic instantaneous costs at t : $cu_i(t)^2$. If the project was not successful at time T , the deadline of this project, the game ends and the project can therefore never be completed.

The utility function of player i is, given a breakthrough at time \bar{t} , therefore given by

$$\tilde{V}_i(u_i, \bar{t}) = re^{-r\bar{t}} - r \int_0^{\bar{t}} e^{-rt} cu_i(t)^2 dt$$

with r being the common discount rate. We can see the two parts of the utility function here: the first part is the lump sum payment, which occurs only once at time \bar{t} and is therefore discounted by $re^{-r\bar{t}}$. The second part is the cost $cu_i(t)^2$ which occurs at

⁶Covered by decreasing hazard rates of the breakthrough effort distribution.

⁷However, what they call procrastination, is, in the terms of my model, a result of the (negative) encouragement effect.

every point in time (and depends on the effort $u_i(t)$) up until \bar{t} .

Therefore, we have to figure out the point in time \bar{t} at which the project is successful. Definition 4 specifies this time \bar{t} as the time at which the players have accumulated enough effort:

Definition 4 (Effort Threshold). The project is successful in \bar{t} if the players have exerted enough effort, i.e. if

$$\bar{x} \leq \int_0^{\bar{t}} \sum_{\forall i} u_i(t) dt$$

Remark. *This definition implies the assumptions of symmetric, additively separable and linear effects of efforts and non-depreciation of effort.*

This threshold \bar{x} is drawn before the game and is unknown to all players. They have a common prior about its probability density function f and hence, about its cumulative distribution function F . This breakthrough effort distribution can be interpreted as the type of task or project (see Section 3.3).

Due to Definition 4 we can define $x(t)$ as the overall effort already spent up until t : $x(t) := \int_0^t \sum_{\forall i} u_i(s) ds$ and $u_{-i}(t) = \sum_{\forall j \neq i} u_j(t)$ as the effort of all players except i at a certain time t without loss of information.

From the definition of the game above, we can derive the expected utility for player i , given effort profile $\{u_i, u_{-i}\}$:

$$V_i(u_i(t), u_{-i}(t), x(t)) = r \int_0^T e^{-rt} (1 - F(x(t))) \left[\frac{f(x(t)) \left(\sum_{\forall j} u_j(t) \right)}{1 - F(x(t))} - cu_i(t)^2 \right] dt \quad (3.1)$$

To find the expected utility, as stated in Equation (3.1), one has to take the expectations with respect to \bar{t} , using Definition 4. For a detailed derivation please refer to Section 3.6.1.

The expected utility has an intuitive interpretation: The factor in front of the squared brackets $1 - F(x(t))$ gives the probability that we had no success before time t or, in other words, that we reach time t . The two terms in squared brackets give us the updated belief of the player about having success at time t minus the costs they have to bear.

3.3 The breakthrough effort distribution

The most important characteristic of a project in this model is the “breakthrough-effort distribution” or, in other words, how much effort one has to spend for a certain chance of success, given the effort that has been spent by the team in the past. Therefore, this distribution describes how likely every possible effort threshold is at the present stage of the project. If the player thinks finding a cure for a disease costs around 100 billion man hours of research, she could assume for example some normal distribution around 100 billion. If I am certain I lost my keys in my apartment (again), but have no idea where they could be, assuming a uniform distribution over every place in my apartment seems reasonable.

In this section I am going to give examples of three basic classes of distributions and how they can be interpreted in the context of the model. The distributions will be denoted by their hazard rates $h(x(t)) := \frac{f(x(t))}{1-F(x(t))}$, which basically describes the effect of past effort on the effectiveness of current effort.

Example 1 (Constant hazard rate). *The first type of distribution has a constant hazard rate, i.e., the exponential distribution ($F(x) = 1 - e^{-\lambda x}$ with a rate parameter $\lambda > 0$). This distribution conveys the idea that the chance of success only depends on the current effort and past effort does not matter at all, for example if you are trying to push a boulder out of your way or trying to force a door open.*

Example 2 (Decreasing hazard rate). *A variation of the exponential distribution is an incomplete exponential distribution⁸ (i.e., an exponential distribution with a probability mass at infinity). Although technically a distribution with a decreasing hazard rate, the intuition is similar to the example of the memoryless distribution: The probability distribution itself is memoryless, however there is a chance of failure. As time proceeds, the expected probability of failure is updated and therefore increases in the effort already spent. A popular example for a decreasing hazard rate is a search model, similar to Keller, Rady, and Cripps (2005), where you search at the most likely places first or investments into R&D: the more you invest without success, the higher is your belief that there is no solution to the problem.*

Example 3 (Increasing hazard rate). *The last example is the class of increasing hazard rates (e.g., when the breakthrough effort is distributed uniformly on some interval). Possible applications are projects with a strong learning-by-doing effect and projects where the success in a certain period depends on the cumulative effort, not on current effort.⁹ A simple example for this class is moving something heavy from A to B.*

⁸Using this distribution in my model yields us a model very similar to the so-called good news bandit models. One example is Bonatti and Hörner’s (2011) benchmark model, the only difference being that I use quadratic instead of linear costs.

⁹One example is Georgiadis (2014). In his model the uncertainty is about the effect of effort and

For more examples I would like to refer to Section 2 in Khan and Stinchcombe (2015), who provide an overview about the meaning of success probability distributions, their hazard rates and their relations to different projects.

Although all examples in this chapter will be from one of the three classes, the results also hold for general distributions.

3.4 Results

3.4.1 Non-cooperative solution

The best response of player i to the strategies of the other players $u_{-i}(t)$ can be stated as the following optimal control problem (omitting the time index t from $x(t)$ and $u_i(t)$):

$$\max_{u_i, x} V_i = r \int_0^T e^{-rt} (1 - F(x)) \left(\frac{f(x)(u_i + u_{-i})}{1 - F(x)} - cu_i^2 \right) dt \quad (3.2)$$

with boundary conditions $x(0) = 0$ for the cumulative effort at time 0.

The following technical assumption restricts our attention to distributions for which there is neither a certain success nor a certain failure.

Assumption 1. The hazard rate $h(x) := \frac{f(x)}{1-F(x)} > 0$ is continuous in x and bounded above for every finite x .

Note that this assumption allows for a probability mass point at ∞ . Given Assumption 1 on the hazard rate of the breakthrough-effort distribution we can find the symmetric equilibrium path.

Theorem 1 (Equilibrium Effort). There exists a unique symmetric Nash equilibrium in which, on the equilibrium path, u (i.e., the individual effort of a player) evolves according to

$$\dot{u} = \frac{2n-1}{2} h(x) u^2 + ru - \frac{r}{2c} h(x)$$

and reaches $u_T = \frac{1}{2c} h(x_T)$ at the deadline T .

To find this equilibrium effort path, I use the Pontryagin maximum principle to solve the optimal control problem given by Equation (3.2) and then use symmetry to find

not the threshold, but this is just a different way to model uncertainty about the relationship between effort spent and success. One can therefore generate a very similar model in the framework presented by choosing the appropriate breakthrough effort distribution, which would have an increasing hazard rate.

a candidate for the equilibrium effort. I then verify existence and uniqueness of this symmetric equilibrium and show sufficiency via a convexity argument. For the complete proof see Section 3.6.2.

Remark (Nash equilibria). *All equilibria in this chapter are symmetric Nash equilibria in pure strategies. To see that these are Nash equilibria, it helps to check the possible histories of the players: At every point in time they only know the time t , their past effort and that they were not successful so far. Therefore, every information set is just a point in time t . It follows from Assumption 1 that each information set (i.e., each time t) is reached with a positive probability. Hence, every information set is part of the Nash equilibrium. From this reasoning we also know that the symmetric Nash equilibrium in pure strategies is also a perfect Bayesian Nash equilibrium, given correct beliefs: $\hat{u}_j = u_j, \forall j \neq i$ for every player i . However, the off-equilibrium behavior is still discussed later in this section.*

Remark (Asymmetric equilibria). *In this chapter I am not analyzing asymmetric equilibria, as, given the symmetric setting, restricting our attention to symmetric equilibria seems natural. In addition, it is clear that every asymmetric equilibrium is, in terms of welfare and as a consequence of the convex cost structure, inferior to the symmetric equilibrium, as can be seen in Lemma 1 in the following section.*

Example 1 (continuing from p. 45). *For the exponential distribution with rate parameter λ , the solution from Theorem 1 reduces to:*

$$\begin{aligned} \dot{u} &= \frac{2n-1}{2} \lambda u^2 + ru - \frac{r}{2c} \lambda \\ u_T &= \frac{1}{2c} \lambda \end{aligned} \tag{3.3}$$

which has an explicit solution given in Section 3.6.6. Given this solution, some observations about this class of distributions can already be made: Independent of the number of players (and the discount rate), the individual effort right before the deadline is always the same. Furthermore, we can see that the individual effort decreases in the number of players, despite the fact that the reward for completion for each player is independent of the number of players. Figure 3.1 shows an example of the equilibrium effort path for one and three players.

Remark (Off-equilibrium behavior). *In the following, I am going to focus my attention on the equilibrium behavior. However, let me briefly discuss the off-equilibrium behavior that arises if one player deviates from the equilibrium path. If she exerts less effort at time t , her continuation strategy after t depends on the hazard rate of the underlying breakthrough effort distribution:*

- *Given a **constant hazard rate**, nothing changes for her. In this special case, behavior is independent of the past, hence she will immediately revert to the equilibrium effort.*




Figure 3.1: Equilibrium effort (constant hazard rate)




Figure 3.2: Equilibrium effort (decreasing hazard rate)

- Given an **increasing hazard rate**, her belief about the probability of success is now lower than that of her collaborators. Therefore, she will also exert less effort in the future. Given a high enough slope of the hazard rate, this leads to divergence of her belief (and therefore effort) and the beliefs of the other players.
- Given a **decreasing hazard rate**, her belief about the probability of success is now higher than that of her collaborators. This leads to a higher effort until her belief coincides again with the belief of the other players, as soon as she has made up the effort she previously failed to exert. So, given enough time, in this case the player will revert to the symmetric equilibrium.

3.4.2 Welfare maximizing solution

To solve the problem of the social planner, we have to solve a problem similar to Equation (3.2). However, now we maximize the combined utility and therefore:

$$\max_{u_i} V_i = r \int_0^T e^{-rt} (1 - F(x)) \left(n \frac{f(x)(\sum_{\forall i} u_i)}{1 - F(x)} - \sum_{\forall i} c u_i^2 \right) dt \quad (3.4)$$

We can focus on the symmetric problem in which every player exerts $\bar{u}(t)$ without loss of generality, as the following Lemma shows us:

Lemma 1. Every welfare-maximizing effort path has to be symmetric.

The intuition for Lemma 1 is as follows: Due to the assumptions of symmetric and additive-separable effects of efforts (Definition 4) and convex costs an equal distribution of the efforts exerted at every point in time results in the same probability of success but a lower sum of costs. For a short proof see Section 3.6.4.




Figure 3.3: Nash equilibrium effort




Figure 3.4: Welfare maximizing effort

Therefore we get the social planners optimization problem

$$\max_{\bar{u}} V_i = r \int_0^T e^{-rt} (1 - F(x)) \left(\frac{f(x)(n^2 \bar{u})}{1 - F(x)} - nc\bar{u} \right) dt \quad (3.5)$$

And its solution which is derived in a similar fashion to Theorem 1 in Section 3.6.5.

Theorem 2 (Welfare Maximizing Effort). The unique effort u that every player has to exert that maximizes the social planners problem (Equation (3.4)) evolves according to

$$\dot{u} = \frac{1}{2}h(x)u^2 + ru - \frac{nr}{2c}h(x)$$

and reaches $u_T = \frac{n}{2c}h(x_T)$ at time T .

Now let us compare the welfare maximizing solution to the (non-cooperative) equilibrium effort for the case of the constant hazard rate:

Example 1 (continuing from p. 45). *If we have constant hazard rates, we can directly compare the non-cooperative equilibrium and the socially optimal effort. It turns out that in this case the socially optimal effort is always larger than the equilibrium effort. This can be checked by simply calculating the difference between the welfare-maximizing effort and the equilibrium effort, as stated in Section 3.6.6.*

Given this example, one might suspect that the welfare maximizing solution is to always exert more effort than in the equilibrium. While this can be observed with an increasing or a constant hazard rate, it is not true for decreasing hazard rates, as can be seen in the following example in Figure 3.3 and Figure 3.4 (which uses the incomplete exponential distribution with rate α and failure rate $1 - \beta$). Here we can see that the welfare-maximizing effort starts off being higher but, due to the decreasing belief in the success of the project, decreases much faster than in the equilibrium.

3.4.3 Free-riding, encouragement effect and delay of effort

In this section I analyzing free-riding, the encouragement effect and delay of effort.

Free-riding

Free-riding is usually defined as players exerting less effort because there are others who also exert effort, i.e., a player lets the others do the work. As can be expected there will always be free-riding in this model. One can verify this by looking at the best response function in Section 3.6.2. The instantaneous effort of a player is strictly decreasing in the instantaneous effort of the other players. This is due to the fact that the effort of every player in one period are perfect substitutes. An illustration of this behavior can be found in Figures 3.1 and 3.2

Encouragement effect

This effect is defined as: My actions affect the efficiency of effort for every player in the future and therefore their choice of effort. The direction in which my effort affects the effort of others depends on the hazard rate. For an increasing hazard rate, the effect is called encouragement effect for a good reason:¹⁰ Every effort I spend now increases the efficiency and therefore the effort of everyone in the future. Clearly, this also leads to higher efforts now. However, with a decreasing hazard rate, this effect is a negative encouragement or discouragement effect: If I spend much effort now and we do not succeed, we have a lower belief about the chance of succeeding in the future and therefore we will work less. This leads to less effort, especially in the earlier periods.

Remark (Strategic complements/substitutes and the encouragement effect). *As already mentioned in the introduction, the positive encouragement effect (due to an increasing hazard rate of the breakthrough-effort distribution) is very similar to strategic complementarity: My effort now is a strategic complement to all players' future efforts. In the same sense, the negative encouragement effect is similar to a strategic substitute for future efforts.*

If that is the case, why can't one just say that with increasing hazard rates, current efforts and future efforts are strategic complements, and with decreasing hazard rates they are strategic substitutes? The problem is that current efforts of different players are substitutes (and therefore strategic substitutes), which is the source for free-riding. Hence, we cannot clearly say if, given an increasing hazard rate, efforts are strategic substitutes or complements. However, for decreasing (and constant) hazard rates, we know that the efforts of all players at all times are strategic substitutes.

¹⁰For example by Georgiadis (2014), where past efforts always have a positive effect on the effectiveness of effort, or in terms of this chapter: projects always have a breakthrough effort distribution with an increasing hazard rate.

Delay of effort

It is not surprising that the effort of players increases if they assume an increasing hazard rate. But what does the optimal effort for a decreasing hazard rate look like? From the example depicted in Figure 3.2 we already get a good idea of what the typical optimal effort path might look like:¹¹ At first, we have a decrease in effort. This is due to the encouragement effect: At first the (perceived) probability of success is high, but due to the decreasing hazard rate as more effort is invested, the success rate and therefore the effort level decreases. However, we can see that, after some time, the effort increases again. What might be the reason for this behavior? By investing early we discourage players to invest at every following point in time, so they postpone investment to a later point in time. However, looking at the case of only one player in Figure 3.2 we see the same effect, albeit less pronounced. Therefore, as a single player is neither affected by the encouragement effect nor by free-riding we can clearly identify another effect: delay of effort.

To investigate if this is a general effect let us first define "last-minute rush".

Definition 5 (Last-minute rush). Player i exhibits a last-minute rush if and only if $\exists \delta > 0$ s.t. u_i is increasing on $[T - \delta, T]$.

If we observe a last-minute rush, we know that the delay of effort outweighs every opposing effect (e.g. the effect of strong discounting) near the deadline.

Given Theorem 1, we can show that a last-minute rush can be observed for every variation of the model:

Theorem 3 (Last-minute rush). For every possible breakthrough-effort distribution that fulfills Assumption 1, players delay their effort, as defined in Definition 5, on the symmetric equilibrium path.

For the complete proof I would like to refer to Section 3.6.3. The interpretation, however, is clear: Whatever effects are at work during the first parts of the project, towards the end delay always dominates the other effects.

As this effect is independent of the number of players it can not be explained by free-riding and it can not be a result of discounting, as it does not depend on the discount factor. Furthermore, it is different and independent from the encouragement effect, which can be shown by the following example (and Figure 3.1). Therefore, we know that there is another effect, which shifts effort towards the end.

Example 1 (continuing from p.45). *Let us now have another look at the case of constant hazard rates. This case is special, as current effort is not influenced by past effort at all. Therefore, we have no encouragement effect (only pure free-riding) and*

¹¹In this example the required effort for a breakthrough is distributed according to the incomplete exponential distribution with rate α and a mass point $1 - \beta$ at ∞ .

it is easy to see that, without a deadline (i.e., $T = \infty$), the effort would be constant¹². However, if we introduce a deadline, this changes and we see an increasing effort as depicted in Figure 3.1 and Equation (3.3).

Now we know that we have a last-minute rush in the competitive case, but what about the social optimum? The following proposition shows that we can expect a rational social planner to delay effort.

Theorem 4 (Last-minute rush of the social planner). The welfare-maximizing behavior always leads to a last-minute rush as defined in Definition 5, independent of the breakthrough-effort distribution.

The proof of Theorem 4 is analogous to the proof of Theorem 3 and can be found in Section 3.6.5.

We have seen that even in the welfare maximizing case without memory (i.e. the exponential distribution), more effort is invested close to the deadline. However, it is not present in Bonatti and Hörner (2011), a very similar model with linear instead of quadratic costs.¹³ Therefore, we can safely assume that this effect is only present when we have convex costs, a deadline and when there is uncertainty about the amount of work we have to put into a project to succeed. Section 3.5.3 shows us that a higher discount rate can dampen this effect, but we know from Theorem 3 that it can never be eliminated completely. As this effect is also present in the welfare maximizing case we can conclude: Delaying effort is not only commonly observed in reality, but might also be rational and even part of the welfare-maximizing solution.

3.5 Discussion

In this section, we are going to have a look at a few results and implications of this model. In Section 3.5.1, we discuss the implications of variable deadlines. In Section 3.5.2, we show that delay of effort can not exist without uncertain objectives. Furthermore, we discuss the special case of patient players in Section 3.5.3, the effect

¹²To do so, we can compare the problem at $t_0 = 0$ and any other time t : The only differences between these two problems are the past time t and the effort already exerted $x(t)$. As t in the past has no influence on the best response now, time left is the same and, due to the special properties of the exponential distribution, $x(t)$ has no effect on the beliefs about the success, the problems we are facing at t_0 and t are the same. Therefore, assuming we also have a unique best response, the continuation strategies at t_0 and t are the same for every t , i.e. players exert a constant effort, independently of t .

¹³In their model, the welfare-maximizing effort is as follows: As the chance of success decreases in the invested effort, players invest the maximal amount of effort until the marginal benefits of effort are lower than the marginal costs. After that point is reached, no effort is invested anymore.

of different costs and different qualities of projects (Section 3.5.5) and the effect of changing the number of players (Section 3.5.4).

3.5.1 Deadlines

We have already seen that deadlines induce delay of effort, i.e. an accumulation of effort shortly before the deadline. But is it possible to improve welfare by a deadline? Given that we only consider rational individuals, one would not expect a deadline to be beneficial if the hazard rate of the breakthrough-effort distribution is constant or even increasing. Now, Bonatti and Hörner (2011) have shown that, in their setup (i.e., with certain type of decreasing hazard rates and linear costs), there is always a deadline that improves welfare.

However, with quadratic costs, I was not able to identify any situation in which deadlines improve welfare. Simulations suggest that the welfare maximizing deadline is always the least restrictive (i.e. the deadline which allows the most time to complete the task) one. This is probably due to the smoothing effect of convex costs, which makes it optimal to spread costs over time as evenly as possible. In Figure 3.5 you can




Figure 3.5: Effect of deadlines on efforts

see that the efforts behave as expected: Shorter deadlines lead to overall higher efforts and, given a short enough deadline, we can even prevent the decrease of effort early on. However the effect of a (shorter) deadline on the utility (and therefore the welfare) is,

at least in my simulations, always negative, as shown in the example in Figure 3.6.¹⁴

The question whether deadlines can improve welfare is therefore still open. However, simulations suggest that, (shorter) deadlines are never beneficial.




Figure 3.6: Effect of deadlines on welfare

3.5.2 Full information

One might expect that delaying effort is also optimal in the full information case, i.e. when the players (or the social planner) already know the effort threshold \bar{x} . The following simple example shows that this is not true, if the costs or the threshold are sufficiently low or if the discount rate is sufficiently high:

Example 4. *Assume that the costs are low enough, s.t. (abusing the notation of $c(\cdot)$): $c\left(\frac{\bar{x}}{n}\right) \leq 1 - e^{-rT}$. Then the losses due to delay until the end T are higher than the highest possible costs that can occur in the symmetric equilibrium.*

Let us compare the utility from getting the work done at $t = 0$: $V_0 = 1 - c\left(\frac{\bar{x}}{n}\right)$ and from getting the work done at $t = T$: $V_T = e^{-rT} - c_T$ where c_T are some non-negative costs. Then we know that:

$$V_0 = 1 - c\left(\frac{\bar{x}}{n}\right) > e^{-rT} - c_T = V_T$$

¹⁴The last example uses again a incomplete exponential distribution as in Bonatti and Hörner (2011).

Completing the project at $t = 0$ is strictly better (but not necessarily optimal). Therefore, working until the end can never be optimal.

Solving the problem with full information¹⁵ shows that the optimal solution is to distribute the required effort between 0 and some time $t^* \leq T$, such that the discounted marginal costs are the same at every point in time. This means that, due to discounting, the effort is increasing until t^* and zero afterwards.

3.5.3 Patient players

So far, we have only considered the problem in which players are impatient. For patient players ($r = 0$), the solution from Theorem 1 simplifies to

$$\begin{aligned} \dot{u} &= \frac{2n-1}{2}h(x)u^2 \\ u_T &= \frac{1}{2c}h(x_T) \end{aligned}$$

As we know that $h(x)$ and u are always positive, the following Proposition 1 is an obvious result:

Proposition 1. The equilibrium effort of patient players (i.e. $r = 0$) is increasing everywhere, concave for decreasing hazard rates and convex for increasing hazard rates.

It is not surprising that, without an incentive to work early, we observe even more delay of effort, i.e. effort is shifted to the end. A very nice example for this phenomenon can be seen in Figure 3.7, where equilibrium efforts paths of different discount rates r are shown. Furthermore, we can see that in our model decreasing efforts, which are for example observed in Figure 3.2 and in Bonatti and Hörner (2011) are, only possible if impatience is strong enough to counteract the delay of effort.

Example 1 (continuing from p. 45). *For the exponential distribution with rate parameter λ and $r = 0$, the solution from Theorem 1 reduces to:*

$$x(t) = \frac{2\lambda(\log(2c + (2n-1)T) - \log(2c + (2n-1)(T-t)))}{2n-1}$$

which is clearly increasing in t .

3.5.4 The effect of the number of players

We have assumed a fixed number of players. What happens if the number of players changes? Assuming increasing hazard rates, the effect is pretty clear: we have stronger

¹⁵For example by finding the best path $u^*(t)$ for the problem $\max_{u^*} e^{-rt^*} - \int_0^{t^*} e^{-rt} c(u^*(t)) dt$
s.t. $\int_0^{t^*} c(u^*(t)) dt = \bar{x}$, for every t^* , and then finding the best t^* .




Figure 3.7: Equilibrium effort for different discount rates

free-riding but in turn have more overall effort which, due to the encouragement effect, leads to higher efforts by the players. With a constant hazard rate, we only have free-riding, so every individual will work less but thanks to quadratic costs and the fact that the payout does not depend on the number of players, we can expect higher overall efforts. As the encouragement effect only depends on the cumulative effort, we can also expect the same effect with a decreasing hazard rate, which can be observed in the example of an incomplete exponential distribution in Figure 3.8 and Figure 3.9. Here we can see the lower individual efforts and higher cumulative efforts for higher number of players.

3.5.5 Effect of costs and project quality on different projects

One might expect that costs and the general quality of a project (i.e., the chance of success), might have very simple effects on the efforts of the players: Higher costs should lead to lower efforts, while higher probabilities of success should lead to higher efforts. While this intuition is true for the constant and increasing hazard rates, this effect is not that simple if we have a project that is described by a decreasing hazard rate, for example if there is a chance of failure. In this case, we might observe that due to the negative encouragement effect, the instantaneous effort starts out being higher but declines faster and finally is even lower than in the case with higher costs (Figure 3.10) or a better quality project (Figure 3.11). In the latter case we can even




Figure 3.8: Effort for different numbers of players




Figure 3.9: Cumulative effort for different numbers of players

observe that it is possible that the cumulative effort is lower, due to faster updating of the failure probability.

3.5.6 Concluding remarks

We have analyzed a problem of a team working together on a project where the individual team members are unable to observe each others' efforts and have only a rough idea about the amount of effort that will be needed to complete the project. This leads to free-riding and encouragement, as well as to delay effort. The delay analyzed here is not a result of inefficient behavior but a necessary consequence of the deadline and convex costs, given the information structure. Although different types of projects lead to very different behavior, a last-minute rush always occurs as long as a deadline is present. The encouragement effect on the other hand has very different effects, depending on the type of project. In this model, delay of effort is so prominent that close to the deadline, it is stronger than every other effect.

This chapter opens up a lot of questions for further research, examples being an analysis of the optimal compensation scheme for different projects, or whether choosing deadlines is an efficient tool for a social planner or a principal who is only interested in a breakthrough. Another interesting question is the effect of team size, which is briefly touched upon in Section 3.5.4.




Figure 3.10: Equilibrium effort paths for different costs




Figure 3.11: Equilibrium effort paths for different project qualities

3.6 Appendix

3.6.1 Derivation of the expected utility

We know that the breakthrough effort is drawn from a distribution with the probability distribution function f and the cumulative distribution function F and that $x(t)$ is by Definition 4

$$x(t) = \int_0^t u_i(s) + u_{-i}(s) ds$$

with $u_{-i}(s) = \sum_{\forall j \neq i} u_j(s)$. The breakthrough time \bar{t} is the first time enough effort (i.e., the breakthrough effort) is accumulated:

$$\bar{t} = \inf\{t \geq 0 | x(t) \geq \bar{x}\}.$$

The expected utility is then, given breakthrough time \bar{t}

$$\tilde{V}_i(u_i, \bar{t}) = re^{-r\bar{t}} - r \int_0^{\bar{t}} e^{rt} c(u_i(t)) dt.$$

Therefore we know that the payoff part of the expected utility is equal to the distribution of \bar{t} : $f(\bar{t})$. As we know that the CDF $F(\bar{t})$ of \bar{t} is

$$F_{\bar{t}}(t) = P[t \geq \bar{t}] = P[x(t) \geq \bar{x}] = F(x(t)) \Rightarrow f_{\bar{t}}(t) = f(x(t)) (u_i(t) + u_{-i}(t)),$$

so the expected payoff is

$$\mathbb{E}_{\bar{t}} [re^{-r\bar{t}}] = r \int_0^T e^{-rt} f(x(t)) (u_i(t) + u_{-i}(t)) dt. \quad (3.6)$$

For the expected costs $r \int_0^{\bar{t}} e^{rt} c(u_i(t)) dt$, we have to distinguish between two cases: One in which the project is successful (i.e., $\bar{t} < T$) and we pay until \bar{t} and one in which it is unsuccessful ($\bar{t} \geq T$) and we only pay until T . As we know that $1 - F_{\bar{t}}(\infty) = P[\bar{t} = \infty] = P[\bar{x} > x(\infty)] = 1 - F(x(\infty))$ we get:

$$\begin{aligned}
 & \int_0^\infty \int_0^{\min\{t,T\}} e^{-rs} cu(s) ds dF(x(t)) + (1 - F(x(\infty))) \int_0^T e^{-rs} c(u(s)) ds \\
 &= \int_0^\infty \int_0^\infty \mathbb{1}_{s<t} \mathbb{1}_{s<T} e^{-rs} c(u(s)) ds dF(x(t)) + (1 - F(x(\infty))) \int_0^T e^{-rs} c(u(s)) ds \\
 &= \int_0^\infty \int_0^\infty \mathbb{1}_{s<t} \mathbb{1}_{s<T} e^{-rs} c(u(s)) dF(x(t)) ds + (1 - F(x(\infty))) \int_0^T e^{-rs} c(u(s)) ds \\
 &= \int_0^\infty \mathbb{1}_{s<T} e^{-rs} c(u(s)) \int_0^\infty \mathbb{1}_{s<t} dF(x(t)) ds + (1 - F(x(\infty))) \int_0^T e^{-rs} c(u(s)) ds \\
 &= \int_0^\infty \mathbb{1}_{s<T} e^{-rs} c(u(s)) (1 - F(x(s))) ds \\
 &= \int_0^T e^{-rs} c(u(s)) (1 - F(x(s))) ds. \tag{3.7}
 \end{aligned}$$

If we add the expected payoff (Equation (3.6)) and the expected costs (Equation (3.7)), we get the expected utility, as stated in Equation (3.1):

$$V_i = r \int_0^T e^{-rt} (1 - F(x(t))) \left(\frac{f(x(t)) (\sum_{\forall j} u_j(t))}{1 - F(x(t))} - c(u_i(t)) \right) dt.$$

3.6.2 Theorem 1 (Optimal Effort)

Candidate solution

Finding the best response u_i of some player i to the strategies of the other players u_{-i} in the problem stated in equation (Equation (3.2)) is a discounted optimal control problem of the following form

$$\begin{aligned}
 H(x, u, \lambda, t) &= f(x)(u_i + \sum_{j \neq i} u_j) - cu_i^2(1 - F(x)) + \lambda(u_i + u_{-i}) \\
 \dot{x} &= u_i + \sum_{j \neq i} u_j \\
 x(0) &= 0 \\
 \lambda(T) &= 0 \\
 u_i, x &\in \mathbb{R}_+ \quad \forall i
 \end{aligned}$$

Using the Pontryagin maximum principle (Pontryagin, Boltyanskii, and Gamkrelidze (1962)) in the version of Kamien and Schwartz (2012), we know that the necessary

conditions for a maximum are

$$\frac{\partial H}{\partial u} = 0 \quad (3.8)$$

$$\frac{\partial H}{\partial x} = r\lambda - \dot{\lambda} \quad (3.9)$$

$$\frac{\partial H}{\partial \lambda} = \dot{x} \quad (3.10)$$

$$\lambda(T)x(T) = 0 \quad \Rightarrow \quad \lambda(T) = 0 \quad (3.11)$$

with the Hamiltonian (Equation (3.8)), the equation of motion for the state variable (Equation (3.9)), the equation of motion for the costate variable (Equation (3.10)) and the transversality condition (Equation (3.11)) for $x(T)$ being free. In addition we can see that the optimal control does not depend on the u_j s of the other players but only on the sum $u_{-i} := \sum_{j \neq i} u_j$. Therefore we get

$$\begin{aligned} u_i &= \frac{f(x) + \lambda}{2c(1 - F(x))} \\ \dot{\lambda} &= r\lambda - f'(x)(u_i + u_{-i}) - cu_i^2 f(x) \\ \dot{x} &= u_i + u_{-i} \\ x(0) &= 0, \quad \lambda(T) = 0 \end{aligned}$$

From here on, we only consider symmetric equilibria, therefore we can replace u_{-i} by $(n-1)u_i$. Hence, necessary conditions for a best response are:

$$\begin{aligned} u_i &= \frac{f(x) + \lambda}{2c(1 - F(x))} \\ \dot{\lambda} &= r\lambda - f'(x)nu_i - cu_i^2 f(x) \\ \dot{x} &= nu_i \\ x(0) &= 0, \quad \lambda(T) = 0. \end{aligned}$$

Using u_i we get

$$\begin{aligned} \dot{\lambda} &= r\lambda - f'(x)n \left(\frac{f(x) + \lambda}{2c(1 - F(x))} \right) - c \left(\frac{f(x) + \lambda}{2c(1 - F(x))} \right)^2 f(x) \\ \dot{x} &= n \left(\frac{f(x) + \lambda}{2c(1 - F(x))} \right) \\ x(0) &= 0, \quad \lambda(T) = 0. \end{aligned}$$

So, the equation of motion for the costate and its time derivative are

$$\begin{aligned} \lambda &= \frac{2c}{n} (1 - F(x)) \dot{x} - f(x) \\ \dot{\lambda} &= \frac{2c}{n} (1 - F(x)) \ddot{x} - \frac{2c}{n} f(x) \dot{x}^2 - f'(x) \dot{x} \end{aligned}$$

Using these, we get the boundary value problem:

$$\begin{aligned}(1 - F(x)) \ddot{x} &= -\frac{nr}{2c} f(x) - \frac{1}{2n} f(x) \dot{x}^2 + f(x) \dot{x}^2 + r(1 - F(x)) \dot{x} \\ x(0) &= 0, \\ \lambda(T) &= \frac{2c}{n} (1 - F(x_T)) \dot{x}_T - f(x_T) = 0\end{aligned}$$

Introducing the hazard rate $h(x) := \frac{f(x)}{1-F(x)}$, we have necessary conditions for Equation (3.2)

$$\begin{aligned}\ddot{x} &= -\frac{rn}{2c} h(x) + \frac{2n-1}{2n} h(x) \dot{x}^2 + r \dot{x} \\ x(0) &= 0, \\ \dot{x}_T &= \frac{n}{2c} h(x_T)\end{aligned}\tag{3.12}$$

Or, in terms of the individual effort:

$$\begin{aligned}\dot{u} &= -\frac{r}{2c} h(x) + \frac{2n-1}{2} h(x) u^2 + ru \\ u_T &= \frac{1}{2c} h(x_T),\end{aligned}\tag{3.13}$$

which is a non-linear boundary value problem of the second order.

Existence, uniqueness and sufficiency of the solution

Now that we have necessary conditions for the equilibrium effort, we still have to check for existence of the solution initial value problem and therefore the Nash equilibrium and if the necessary conditions are sufficient.

Checking for existence and uniqueness first gives us the following

Proposition 2 (Existence and Uniqueness). A solution to the initial value problem from Equation (3.12) (and therefore also for Equation (3.13)) exists and is unique.

Proof. As $u_i : [0, T] \rightarrow \mathbb{R}_+$ is continuous and maps from a compact space to a metric space, we know that it is bounded. Therefore, $\frac{r}{2c} h(x) + \frac{2n-1}{2n} h(x) u^2 - ru$ is Lipschitz-continuous in u and t . Thus (by Picard-Lindelöf), we know that a unique solution to the initial value problem for u_i exists.

As $x(t) = \int_0^t nu(s) ds$ and $x_0 = 0$ and u exists and is unique, $x(t)$ also exists uniquely. □

Now we know that the candidate solution from Equation (3.13) exists and is unique, we have to show that it is in fact the maximum:

Proposition 3 (Maximum). The candidate equilibrium strategy u_i as defined in Equation (3.13) is the solution to the maximization problem in Equation (3.2) and therefore maximizes V_i .

Proof. Given Assumption 1 and Proposition 2 we know that u_i exists, is unique and continuous and we know that V_i is continuous in u_i . Furthermore, we know that (abusing the notation of $u_i = c$ as $u_i(t) = c \quad \forall t \in [0, T]$):

$$\begin{aligned} u_i = 0 &\Rightarrow V_i = 0 \\ \exists \varepsilon > 0 : u_i = \varepsilon &\Rightarrow V_i > 0 \\ \lim_{u_i \rightarrow \infty} V_i &= -\infty \end{aligned}$$

Therefore V_i is concave in u_i and, as it is also continuous in u_i , the necessary conditions for a maximum from Equation (3.13) are sufficient. \square

Therefore we have established the necessary and sufficient conditions for an equilibrium and have shown that, given Assumption 1, this equilibrium always exists uniquely, as stated in Theorem 1.

3.6.3 Theorem 3 (Last-minute rush)

To prove Theorem 3, we use continuity of x to show that the negative part of \dot{u} vanishes near the deadline and is therefore strictly positive.

Proof. As $h(x_t)$ is continuous in x , it is also continuous in t . We also know from Theorem 1 that u_t is continuous in t and that it satisfies

$$\dot{u} = \frac{2n-1}{2}h(x)u^2 - r(u - \frac{1}{2c}h(x)), \quad u_T = \frac{1}{2c}h(x_T)$$

As u_t and $h(x_t)$ are continuous in t , we know that:

$$\Rightarrow \lim_{t \rightarrow T} (u - \frac{1}{2c}h(x)) \rightarrow (u_T - \frac{1}{2c}h(x_T)) = 0$$

Now define $\varepsilon(t) = \frac{2n-1}{2r}h(x_t)u_t^2$. Then we know,

$$\begin{aligned} \exists \delta > 0 : \hat{t} := T - \delta &\Rightarrow \left(u_t - \frac{1}{2c}h(x_t) \right) < \varepsilon(\hat{t}) = \frac{2n-1}{2r}h(x_{\hat{t}})u_{\hat{t}}^2 \\ \Rightarrow \dot{u}_{\hat{t}} &= \frac{2n-1}{2}h(x_{\hat{t}})u_{\hat{t}}^2 - r(u_{\hat{t}} - \frac{1}{2c}h(x_{\hat{t}})) > 0. \end{aligned}$$

\square

Therefore, we know that there is always an interval $[\hat{t}, T]$ in which the effort is strictly increasing.

3.6.4 Lemma 1 (Asymmetric equilibria)

Proof. Assume there is an asymmetric equilibrium that is welfare maximizing. Then $\exists i, t, j : u_i(t) > u_j(t)$. However, it would be possible to improve welfare by setting a new $u_i^*(t)$ and $u_j^*(t)$ as follows: $u_i^*(t) = u_j^*(t) = \frac{u_i(t) + u_j(t)}{2}$ as this does not change the overall effort (and therefore the chance of success) but reduces, due to the quadratic costs, the combined expected costs of the project. Therefore, an asymmetric equilibrium can never be welfare maximizing. \square

3.6.5 Theorem 2 (Social planner)

Applying similar methods as in Section 3.6.2, we get the welfare maximizing cumulative effort:

$$\begin{aligned}\ddot{x} &= -\frac{rn^2}{2c}h(x) + \frac{1}{2}h(x)\dot{x}^2 + r\dot{x} \\ x(0) &= 0, \\ \dot{x}_T &= \frac{n^2}{2c}h(x_T)\end{aligned}$$

Or, in terms of instantaneous effort u :

$$\dot{u} = \frac{1}{2}h(x)u^2 + ru - \frac{nr}{2c}h(x)$$

which reaches $u_T = \frac{n}{2c}h(x_T)$ at time T .

Furthermore, the properties derived in Proposition 2 and Proposition 3 for the optimal effort in the game also apply to the solution of the social planner's problem (Theorem 2).

3.6.6 Constant hazard rate

For the exponential distribution with rate parameter λ , the solution from Theorem 1 evolves according to

$$\begin{aligned}\dot{u} &= \frac{2n-1}{2}\lambda u^2 + ru - \frac{r}{2c}\lambda \\ u_T &= \frac{1}{2c}\lambda,\end{aligned}$$

which has the following explicit solution:

$$u(t) = \frac{\lambda \left(e^{\frac{(t-T)\sqrt{cr(cr+\lambda^2(2n-1))}}{c}} + 1 \right) \left(\lambda^2(2n-1) + cr - \sqrt{cr(cr+\lambda^2(2n-1))} \right) + 2\sqrt{cr(cr+\lambda^2(2n-1))}}{2c(cr+\lambda^2(2n-1)) \left(\left(1 - \frac{2cr+\lambda^2(2n-1)}{2\sqrt{cr(cr+\lambda^2(2n-1))}} \right) e^{\frac{(t-T)\sqrt{cr(cr+\lambda^2(2n-1))}}{c}} + \frac{2cr+\lambda^2(2n-1)}{2\sqrt{cr(cr+\lambda^2(2n-1))}} + 1 \right)}.$$

The welfare maximizing effort has the following explicit solution

$$u(t) = \frac{\lambda n e^{-T\sqrt{r\left(\frac{\lambda^2 n}{c} + r\right)}} \left(\left(r - \sqrt{r\left(\frac{\lambda^2 n}{c} + r\right)} \right) e^{t\sqrt{r\left(\frac{\lambda^2 n}{c} + r\right)}} - \left(\sqrt{r\left(\frac{\lambda^2 n}{c} + r\right)} + r \right) e^{T\sqrt{r\left(\frac{\lambda^2 n}{c} + r\right)}} \right)}{\lambda^2 n \left(e^{(t-T)\sqrt{r\left(\frac{\lambda^2 n}{c} + r\right)}} - 1 \right) + 2cr \left(e^{(t-T)\sqrt{r\left(\frac{\lambda^2 n}{c} + r\right)}} - 1 \right) - 2\sqrt{cr(cr + \lambda^2 n)} \left(e^{(t-T)\sqrt{r\left(\frac{\lambda^2 n}{c} + r\right)}} + 1 \right)}.$$

Calculations show that the non-cooperative equilibrium effort is always lower than the welfare maximizing effort.

Chapter 4

Probabilistic Transitivity in Sports¹

4.1 Introduction

In many situations we are confronted with data about a certain set of objects which only include an array of comparisons about two of these objects at a time. Then all too often the task arises to find the "fairest" or "most legitimate" ranking among all of the objects in the considered set reaching from the "best" one to the "worst" one.

The probably most popular application of such paired comparisons is sports. In most sports games two opponents face each other in a duel. The result can be a win for one of the teams or, depending on the sport, also a tie.

An important attribute of a ranking is that it expresses a transitive relation between all of its objects. This means that if object or team A precedes B and B precedes C , it automatically implies that A precedes C . In contrast to this, paired comparison data can include circular relations, which seem to be inconsistent with this property. In a tournament it is possible that A beats B , B beats C , but C beats A . It is easy to imagine that as the number of teams rises, the probability of the occurrence of such inconsistencies rapidly increases. In the literature many suggestions have been made to overcome these inconsistencies and find a ranking with a good fit according to different concepts. A good overview of the classical models for obtaining rankings from data sets gives Brunk (1960). One approach that deserves attention is the one proposed by Slater (1961). Here the observed number of inconsistencies (in the sense mentioned above) is minimized. This nontrivial problem later became known as a particular form of the so called linear ordering problem. For a good survey on the linear ordering problem see for example Charon and Hudry (2010).

¹Parts of this chapter can be found in Tiwisina and Külpmann (2014).

The major issue concerning the mentioned approaches is that despite all of them having some intuitive appeal, they seem to be rather arbitrary in finding the "right" ranking. The difference of our approach is that we assume that there actually *exists a correct ranking*. Of course we cannot directly observe it, but we can try to find the ranking which is most likely identical to it. To be more precise, we first of all make the assumption that the outcome of each match follows a trinomial distribution, with a fixed probability for a loss, a tie, and a win. These unobservable probabilities fulfill a certain form of transitivity. Applying the respective conditions we can then use a likelihood function to gauge the chance of the observed set of results given a particular set of probabilities. Maximizing this likelihood function while fulfilling the transitivity conditions answers the question about the most likely *correct* ranking.

In the literature there can be found plenty of works using the concepts of the so called weak and strong stochastic transitivity. These are definitions, which transfer the very intuitive concept of transitivity to the world of probabilities. Because in our model we consider ties and also home/away asymmetries, we are forced to define our own concept which goes beyond WST and SST.

At this point the optimization problem, which is the main object of the chapter, is completely defined by the set of probabilities for three outcomes for each game, the likelihood function which shall be maximized, and finally the set of constraints imposed by the stochastic transitivity defined above. We are not the first authors trying to find a maximum likelihood ranking while applying probabilistic transitivity conditions. Thompson and Ramage (1964) propose a similar problem of ranking pairwise compared objects. The analysis is extended in Singh and Thompson (1968) by the incorporation of ties. However, Thompson uses only constraints of WST.² This contributes a lot to the simplicity of the problem and enables Decani (1969) to formulate it as a linear program and later propose in Decani (1972) a branch and bound algorithm to solve the problem even more efficiently.

Unfortunately the new set of constraints make things much more complicated. Increasing the number of teams leads to a huge number of constraints. And it is straightforward to see that the space of transitive probability sets of a particular dimension is not convex. So it is not a surprise that state of the art solvers do not succeed in finding the optimal solution to this non-linear, non-convex problem as soon as the number of teams is increased to more than 5 or 6.

This is why we split up the problem in two parts. The first one is to find the probability sets and the likelihood for a fixed ranking and the second one is to find the ranking

²After the incorporation of ties he naturally can't use the WST constraints, but has to alter his concept. However, it still differs substantially from ours which makes a comparison very difficult.

with the greatest likelihood.

When the goal is to find probabilities for a fixed ranking, while still sticking to the transitivity definition, the constraints become much simpler.

The problem we arrive at is now very close to the so called isotonic regression problem in which a set of probabilities needs to be estimated, while one knows their order according to their magnitude (see Barlow and Brunk (1972) or Van Eeden (1996) for an overview). A reference much closer to the subject of this chapter is Brunk (1955). Here the random variables (in our case the match results) are assumed to follow a distribution belonging to an exponential family. The single distribution parameter follows a function depending monotonically on potentially multiple variables. These variables would in this work correspond to the two teams that are playing. The very efficient method developed in this chapter later became known as the pool adjacent-violators algorithm (PAVA). The major difference of Brunk's chapter to our approach is that the trinomial distribution we will be using does not belong to the exponential family he is referring to. It also has not one but two distribution parameters. So we are very unfortunate to not being able to apply the PAVA. To be able to estimate not only ordered binomial but also ordered multinomial distribution parameters Jewell and Kalbfleisch (2004) developed a modification of this algorithm, the so called m-PAVA. This algorithm is technically able to solve our first problem, but turns out to be very inefficient and slow. But there is an alternative. Lim, Wang, and Choi (2009) find that a program of the kind we are facing can be formulated as a geometric program, which then can be transformed into a convex program. By applying state of the art interior point solvers, we are then able to find a solution very efficiently.³

The second part of the problem is more complicated. If we increase the number of teams, the possible number of orderings rises very quickly. For 4 teams there are 24 possibilities, for 5 teams there are 120 and for 18 teams there are more than 6×10^{15} . But even if we're not able to find the optimal ranking, we are still able to compare different rankings created by the application of empirically relevant ranking systems. And this is exactly what we do in the empirical subsection of the chapter. Among the candidates are the classical "three points for a win" and "two points for a win" systems from soccer and also the Elo system applied e.g. in chess.

To be able to make a good judgment about the true quality of the systems when applied to different sports, we develop a statistical test. It assumes the trueness of the null hypothesis stating that one of two ranking systems under consideration is able to find the correct ordering. Then we estimate all the probabilities and simulate a test

³In Lim, Wang, and Choi (2009) investigations geometric programming is more than 150 times faster.

statistic. Combined with the empirically observed likelihoods, we are then ideally able to reject the null hypothesis which lets us state that here the considered system is not able to generate the correct ranking.

The chapter proceeds as follows. In Section 4.2 the formal model is introduced. In Section 4.4 the problem solution for a known ranking is described, before in Section 4.5 we discuss strategies for finding optimal rankings. The next two sections then describe the sports data and provide a thorough empirical analysis. Section 4.9 concludes.

4.2 Setup

Many sports have in common that n teams or individuals are competing in a number of repeated one-on-one games. The results of these games should be aggregated to one final complete ranking. Let p_{ij} be the probability that team i beats team j .

Naturally, we must have $\forall i, j \in \{1, \dots, n\}$

$$\begin{aligned} p_{ij} &\in [0, 1] \\ p_{ij} + p_{ji} &\leq 1 \end{aligned} \tag{4.1}$$

It can be observed that playing at home (meaning in i 's stadium) and playing away makes a difference to the winning probabilities. Therefore we introduce different probabilities for at home and away games: p_{ijh} is the probability that i beats j at home and p_{jia} that team j wins against i at i 's stadium.

Therefore Equation (4.1) changes to

$$p_{ijh} + p_{jia} \leq 1 \quad \forall i, j \in \{1, \dots, n\}$$

Since in many sports there exists the possibility of a draw, there is no strict equality. In fact, the probability of a draw is

$$q_{ijh} = q_{jia} = 1 - p_{ijh} - p_{jia}.$$

In this chapter, we want to make only one assumption concerning a set of those probabilities. This assumption is based on the concept of weak and strong stochastic transitivity, which formalizes the very intuitive thought that if team i is better than team j and j is better than k then i has to be better than k , as well. In a model of symmetric paired comparison without ties this can be translated fairly easily into stochastic terms.

$$p_{ij} \geq 1/2 \quad \wedge \quad p_{jk} \geq 1/2 \implies p_{ik} \geq 1/2 \tag{WST}$$

$$p_{ij} \geq 1/2 \quad \wedge \quad p_{jk} \geq 1/2 \implies p_{ik} \geq \max\{p_{ij}, p_{jk}\} \tag{SST}$$

Where Equation (SST) is equivalent to

$$p_{ij} \geq 1/2 \implies p_{ik} \geq p_{jk}.$$

The concept of stochastic transitivity has been widely used in the literature on paired comparisons, especially in the 60s and 70s (see e.g. Tversky (1969), Chung and Hwang (1978), Morrison (1963) or Davidson and Solomon (1973)).

The introduction of ties and in addition to that the introduction of a home/away asymmetry forbid to use this concept directly. Equation (SST) is best interpreted by saying "if team i is better than team j , it has to have a higher chance of beating any third team k ". But in a world with draws and home advantage we cannot interpret "being better" as $p_{ij} > 1/2$. That's why one has to alter this point. This is done in the following definition.

Definition 6 (Transitivity). A set of probabilities will be called transitive if the following holds for every $i, j, k, l \in \{1, \dots, n\}$, $x, y \in \{a, h\}$ and $\exists i', j', k', l' \in \{1, \dots, n\}$:

$$\begin{aligned} p_{ikx} \geq p_{jkx} &\Leftrightarrow p_{ily} \geq p_{jly} \\ p_{kix} \geq p_{kix} &\Leftrightarrow p_{liy} \geq p_{lly} \\ p_{i'k'x} > p_{j'k'x} &\Rightarrow p_{l'j'x} > p_{l'i'x} \end{aligned} \tag{4.2}$$

The set of transitive probability sets will be called \mathcal{T} .

The first proposition shows that our concept is in fact a generalization of SST.

Proposition 4. Definition 6 is, when assigning 0 to all draw probabilities and ignoring away/home differentiation, equivalent to Equation (SST).

Definition 7 (Transitive Ranking). A ranking will be called transitive if for all i ranked above j the following holds:

$$p_{ikh} \geq p_{jkh}, p_{kih} \leq p_{kjh}, p_{ika} \geq p_{jka}, p_{kia} \leq p_{kja} \quad \forall k \in \{1, \dots, n\} \setminus \{i, j\}$$

The set of probability sets according to this definition will be called \mathcal{T}' .

The fact that a transitive ranking has a set of transitive probabilities and every set of transitive probabilities has a transitive ranking is established in the following Proposition.

Proposition 5. A set of probabilities P is in \mathcal{T} if and only if it is in \mathcal{T}' .

For the proofs of Proposition 4 and Proposition 5 see Section 4.10.1.

The structure of the constraints and hereby the problem we have to solve becomes clearer, if we write down the set of p_{ijx} values in matrix form and add the constraints using one particular ranking.

$$\left(\begin{array}{c} * \leq p_{12h} \leq p_{13h} \leq \dots \leq p_{1nh} \\ \vee | \quad \vee | \quad \vee | \quad \vee | \quad \vee | \\ p_{21h} \leq * \leq p_{23h} \leq \dots \leq p_{2nh} \\ \vee | \quad \vee | \quad \vee | \quad \vee | \quad \vee | \\ \dots \leq \dots \leq \dots \leq \dots \leq \dots \\ \vee | \quad \vee | \quad \vee | \quad \vee | \quad \vee | \\ p_{n1h} \leq p_{n2h} \leq p_{n3h} \leq \dots \leq * \end{array} \right), \left(\begin{array}{c} * \leq p_{12a} \leq p_{13a} \leq \dots \leq p_{1na} \\ \vee | \quad \vee | \quad \vee | \quad \vee | \quad \vee | \\ p_{21a} \leq * \leq p_{23a} \leq \dots \leq p_{2na} \\ \vee | \quad \vee | \quad \vee | \quad \vee | \quad \vee | \\ \dots \leq \dots \leq \dots \leq \dots \leq \dots \\ \vee | \quad \vee | \quad \vee | \quad \vee | \quad \vee | \\ p_{n1a} \leq p_{n2a} \leq p_{n3a} \leq \dots \leq * \end{array} \right)$$

Figure 4.1: Transitivity matrices for home and away probabilities

4.3 The optimization problem

By assumption, each outcome in a set of paired comparisons is trinomially distributed. The probability distribution is

$$Pr\{x_{ij} = w_{ij}\} = p_{ijh}^{w_{ijh}} p_{jia}^{w_{jia}} (1 - p_{ijh} - p_{jia})^{m_{ij} - w_{ijh} - w_{jia}} \quad (4.3)$$

where w_{ij} is the vector consisting of the elements w_{ijh} and w_{jia} . x_{ij} is the analogously defined vector of a realization of the corresponding random variable. Equation (4.3) tells us the probability of a certain outcome of a game between two particular teams in one particular stadium. By taking the exponential of the natural logarithm of the left side, we can write the above equation as

$$\begin{aligned} Pr\{x_{ij} = w_{ij}\} &= \exp(w_{ijh} \ln(p_{ijh}) + w_{jia} \ln(p_{jia}) \\ &\quad + (m_{ij} - w_{ijh} - w_{jia}) \ln(1 - p_{ijh} - p_{jia})) \end{aligned}$$

Let

$$F[x_{ij}, p_{ij}] := w_{ijh} \ln(p_{ijh}) + w_{jia} \ln(p_{jia}) + (m_{ij} - w_{ijh} - w_{jia}) \ln(m_{ij} - p_{ijh} - p_{jia})$$

The likelihood of a set of particular results to occur will be

$$Pr\{(x_{ij}, \dots, x_{i'j'}) = (w_{ij}, \dots, w_{i'j'})\} = \exp(F[w_{ij}, p_{ij}] + \dots + F[w_{i'j'}, p_{i'j'}])$$

Let E be the set of all valid (i, j) combinations $E = \{(i, j) | i, j \in \{1, \dots, n\}, i \neq j\}$. Then Equation (4.2) implies that, in order to maximize the likelihood of a set of outcomes, we have to solve the following maximization problem

$$\max_{p_{ij}} J[p] = \sum_{(i,j) \in E} F[w_{ij}, p_{ij}] \quad s.t. \quad \{p_{ijx} | (i, j) \in E, x \in \{h, a\}\} \in \mathcal{T}$$

This is a rather complicated optimization problem, first because the objective function (the log of the likelihood function) is not linear, and second because we have a huge number of non-linear constraints, which make the space we are dealing with highly convoluted and non-convex. We can achieve convexity by fixing a particular ranking of teams. In this case we face a total number of $2(2(n-2)n + (n-1))$ constraints. Note that a simple transformation of parameters cannot help us making the problem convex. Also it cannot make the problem linear after fixing a ranking. In this highly simplified case, where the untransformed constraints can be expressed in a linear form, a logarithmic transformation would make the objective function linear but take away linearity from the constraints. More details on this will follow in Section 4.4.2.

4.4 Optimization under a known ranking

Note that the probabilities depicted in Figure 4.1 are only the constraints that apply for one ranking. So the optimization problem can be split into first finding the optimal (i.e., likelihood maximizing) probabilities that satisfy the monotonicity constraints from the matrix and second finding the best ranking. It should become clear that if we consider the indices as variables of the functions $p_h(i, j)$ and $p_a(i, j)$, then this function is monotone non-increasing in the first variable and monotone nondecreasing in the second one. In the considered case the two matrices are only insofar dependent on each other as the sum of an element of the upper right half of the first matrix depicted in Figure Figure 4.1 and the corresponding element of the bottom left half of the second matrix has to be less than or equal to unity.

4.4.1 Transitivity without draws

Now, let us again compare the original problem to the one in the much simpler case without ties. Here, the problem of estimating the probabilities is much easier. Given the above assumptions, the number of wins when two teams play each other a particular amount of times follows an elementary binomial distribution. This instant allowed Brunk (1955) to develop an algorithmic approach, building the foundation of what later became known as the Pool Adjacent Violators Algorithm (PAVA). See also Brunk (1960) for an application to paired comparisons. It follows a short description of the estimation procedure.

A lower interval is the set of all points (i, j) for which $i \geq i', j \leq j'$. So it includes a point in one of the above matrices as well as all the points in its south-west quadrant. An upper interval is analogously defined. A lower layer is a union of lower intervals and an upper layer is a union of upper intervals.

The procedure is now to find the largest upper layer within which the average number of wins is maximized. That is, we have to find an upper layer with the property that the number of wins divided by the number of games it comprises is maximal. For each p_{ij} in this layer the maximum likelihood estimate under the monotonicity constraints we defined is this average number of wins. Next step is to repeat the procedure on the remaining set of the matrix of results.

To illustrate the approach, consider the following example of a tournament of 4 teams in which each two teams played each other once. (For simplicity we only consider home games of the row teams, here.)

$$\left(\begin{array}{cc|cc} * & 0 & 1 & 0 \\ 0 & * & 1 & 0 \\ \hline 1 & 1 & * & 1 \\ 1 & 0 & 0 & * \end{array} \right) \quad \left(\begin{array}{cccc} * & \frac{1}{2} & \frac{3}{5} & \frac{3}{5} \\ \frac{1}{2} & * & \frac{3}{5} & \frac{3}{5} \\ \frac{1}{2} & \frac{1}{2} & * & \frac{3}{5} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & * \end{array} \right)$$

Figure 4.2: PAVA example: Result matrix and p-Matrix

On the left there is the matrix of tournament results. The solid line shows the first upper layer with an average number of wins of $3/5$, giving us the p-value listed in the right matrix. The second layer includes all the numbers above and to the right of the dashed line. Here the average value is $1/2$ and so on. Having the p-Matrix at hand, it is straightforward to calculate the maximum likelihood of the tournament to be 0.03888.

Please note that this algorithm, while being very efficient at finding the probabilities for a fixed ranking, does not help finding the optimal permutation of the teams. To find it, one is still forced to apply this algorithm $4! = 24$ times for this example.

Unfortunately including the chance of draws forbids to use this very simple and efficient procedure. In the next subsection we show how to arrive at a solution nonetheless.

4.4.2 Solution process for the case including draws

Again focusing on the part of the problem where the ranking is already fixed, allowing for ties makes the solution procedure much more complicated. Now, the task is not to estimate ordered binomial, but rather ordered trinomial distribution parameters. Jewell and Kalbfleisch (2004) developed an extension of the PAVA algorithm discussed above. The Authors call this algorithm the modified- or m-PAVA algorithm. In the process the problem is iteratively broken down into many one dimensional optimization problems. Since the number of these subproblems grows very quickly with the number of teams and also the number of adjacent violators, the required computational effort also does. This is the main reason for Lim, Wang, and Choi (2009) to reconsider

the problem, finding that it can be formulated as a geometric program. Then it can be transformed into a convex optimization problem, for which one can find a global solution very efficiently with the help of e.g. interior-point algorithms. Lim, Wang, and Choi (2009) compare the computational efficiency of the two approaches and find that geometric programming is much faster than the m-PAVA algorithm. These findings facilitate the choice for us in this chapter.

Let us take a look at it in detail. We define w_{ijh} to be the empirically observed number of times team i beats team j at home and $t_{ijh} = t_{jia}$ as the number of times team i ties team j . Let m_{ij} be the total number of games between i and j at i 's stadium. Consider the optimization problem for a fixed ranking in its raw form.⁴

$$\begin{aligned}
 \min_p \quad & \prod_{(i,j) \in E} p_{ijh}^{-w_{ijh}} p_{jia}^{-w_{jia}} (1 - p_{ijh} - p_{jia})^{-(m_{ij} - w_{ijh} - w_{jia})} \\
 \text{s.t.} \quad & \frac{p_{ijx}}{p_{ikx}} \leq 1 \quad \forall \quad (i,j) \in E, (i,k) \in E, j \succeq k, x \in \{h, a\} \\
 & p_{ijh} + p_{jia} \leq 1 \\
 & p_{ijx} \geq 0
 \end{aligned} \tag{4.4}$$

This is a geometric program. The objective function as well as the left side of the first constraint are monomial and the left side of the second constraint are polynomials. The third constraint reflects the fact that the domain of our objective function is positive, as in all geometric programs. The program can easily be transformed to a convex optimization problem.

$$\begin{aligned}
 \min_p \quad & \sum_{(i,j) \in E} -w_{ijh} \ln(p_{ijh}) - w_{jia} \ln(p_{jia}) - (m_{ij} - w_{ijh} - w_{jia}) \ln(1 - p_{ijh} - p_{jia}) \\
 \text{s.t.} \quad & \ln(p_{ijx}) - \ln(p_{ikx}) \leq 0 \quad \forall \quad (i,j) \in E, (i,k) \in E, j \succeq k, x \in \{h, a\} \\
 & \ln(e^{\ln(p_{ijh})} + e^{\ln(p_{jia})}) \leq 0
 \end{aligned}$$

It is straightforward to show that the logarithm of a posynomial is convex in $\ln(x)$, which proves the fact that this is indeed a convex program. To solve this kind of program we make use of the software package IPOPT (see Wächter and Biegler (2006)). In addition to the program it requires the input of the Jacobian and Hessian matrices of the constraints. It then applies an interior point algorithm and solves our problem very efficiently, given a fixed ranking. This allows us to compare different ranking systems.

⁴The only change made is the conversion to a minimization instead of a maximization problem.

4.5 Ranking methods

4.5.1 The Linear Ordering Problem

At this point, before proceeding with our efforts of finding solutions to the proposed problem, it makes sense to consider a related, but as we will see, clearly different problem. As one of the classical combinatorial optimization problems the linear ordering problem (LOP) attracted many authors resulting in a huge amount of literature on it. See for example Marti and Reinelt (2011) for a good introduction to the problem as well as a review of suitable algorithms. Also feel referred to Charon and Hudry (2010) for a detailed survey.

If one is given a complete directed graph $D_n = (V_n, A_n)$ with arc weights c_{ij} for every ordered pair $(i, j) \in V_n \times V_n$, the linear ordering problem consists of finding an acyclic tournament T (which corresponds to a permutation of the set of objects or teams), which maximizes the sum of the arcs which are in agreement with the direction of the arcs from D_n . So the sum $\sum_{(i,j) \in T} c_{ij}$ has to be maximal. Equivalently one could formulate the problem as minimizing the so called remoteness corresponding to minimizing the arc weights pointing in the opposite direction.

A more illustrative representation of the problem is the maximization of the sum of superdiagonal elements in a matrix by manipulating the row/column ordering. This is the so called Triangulation Problem.

The reader might already be able to grasp a sense of similarity here. To establish a direct connection between the LOP and the problem dealt with in this chapter, consider a situation where we fix the probabilities of wins and losses at homogeneous values below and above the diagonal of the matrix independently of which teams are in question. This means we set $p_{ijh} = \bar{p}_h$ above diagonal and $p_{ijh} = \underline{p}_h$ below it and analogously for the away probabilities. Let us consider the case where $\bar{p}_h > \underline{p}_h$ and $\bar{p}_a > \underline{p}_a$. Remember that the goal is to maximize

$$\begin{aligned} & \sum_{(ij) \in E} w_{ijh} \ln(p_{ijh}) + w_{jia} \ln(p_{jia}) + (1 - w_{ijh} - w_{jia}) \ln(1 - p_{ijh} - p_{jia}) \\ = & \sum_{(ij) \in \bar{E}} w_{ijh} \ln(\bar{p}_h) + w_{jia} \ln(\bar{p}_a) + t_{ijh} \ln(1 - \bar{p}_h - \bar{p}_a) \\ & + \sum_{(ij) \in E} w_{ijh} \ln(\underline{p}_h) + w_{jia} \ln(\underline{p}_a) + t_{ijh} \ln(1 - \underline{p}_h - \bar{p}_a) \end{aligned}$$

where \bar{E} and E represent the sets of elements above and below the diagonals, respectively.

The results of a particular team in his two games against a particular opponent makes a certain contribution to the sum. This contribution might be higher because it is multiplied by higher probabilities if the records are superdiagonal. So we are confronted with a triangulation problem just like the one described above. Many Authors suggest an application of the LOP in sports rankings (see e.g. Marti and Reinelt (2011)). And since it indeed seems well suited for our purposes, we will include it in the analysis.

4.5.2 Branch and Bound Algorithm

Branch and Bound Algorithms are particularly well suited for combinatorial optimization problems. As opposed to the other methods we are proposing, this one leads with certainty to the optimal ranking. For an early survey on Branch and Bound methods feel referred to Lawler and Wood (1966).

The following steps describe the execution of the algorithm:

1. Take the next team from the list of all teams
2. Put it in the list of previously selected teams at each possible position
3. For each position calculate an upper bound \bar{L} above which the likelihood cannot rise going further down the tree (i.e., after all teams were inserted)
4. Leave the team at the position with the highest upper bound
5. If all teams are inserted go to 6., otherwise go to 1.
6. Compare the likelihood to the best one found so far
7. Cut of the tree at all nodes where \bar{L} is below the best likelihood
8. Go to the best of the lowest hanging nodes that could not be deleted and start with 1. from there

Before asking how the upper bound estimate \bar{L} is calculated, lets first focus on the procedure itself. To understand it better, consider a simple example of three teams "a" "b" and "c".

We start by inserting team "a". The upper bound for the log likelihood at this point is still 0, which is indicated in brackets in Figure 4.3. Then team "b" is added at each possible position. We see upper bounds of -1.3 and 0, respectively. So we continue by leaving "b" at the second position and then insert team c at each possible location. Since the example only includes three teams, we can now calculate the value of the real objective function instead of calculating \bar{L} the way it was done previously. The highest value of the objective function is found using the ordering "bac". This value of -1.2 now enables us to cut of all hanging nodes, which have an upper bound below -1.2.

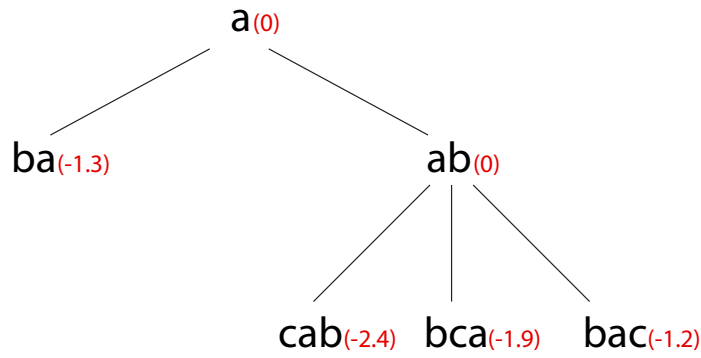


Figure 4.3: Branch and Bound Algorithm: An example

So we cut of the tree at "ba", since there is no way, we could get a better likelihood going down the tree from this node. It is easy to see how the procedure can save computational effort (even in this tiny example) compared to calculating the MLE for all permutations.

The upper bound \bar{L} is calculated as follows. First the optimization problem (for a fixed ranking) is applied to the teams that have been inserted so far.

Lemma 2. Adding an additional team into an existing ranking without changing the relative order of the already existing teams can not increase \bar{L} .

Proof. It is trivial to see that adding a variable (team) to the maximization problem without adding additional constraints (results) does not change the maximum likelihood (i.e., we are multiplying by 1). Now, adding additional constraints without changing the objective function or changing the other constraints can never increase the maximum likelihood and therefore the new \bar{L} has to be less or equal to the \bar{L} with 1 team less. \square

At this stage we could already use this maximum likelihood of the considered subset of teams for \bar{L} . But there is a way to reduce the upper bound even further and thereby make the algorithm a lot more efficient. For each team that is still pending to be inserted we already know a subset of the constraints that will be applied to the corresponding probabilities when going further down the tree, no matter where this particular team will be inserted. Consider a situation where teams $1, \dots, k$ have already been inserted. Now, for each team $l \in \{k + 1, \dots, n\}$ we know that $p_{lix} \leq p_{i'l_x}$ and $p_{li_x} \geq p_{i'l_x}$ for every $i, i' \in \{1, \dots, k\}$ and $x \in \{h, a\}$ such that i is ranked above i' . For $k = 3$ this is depicted in Figure 4.4.

For every team that has not been inserted yet, we know this subset of constraints. So we have another optimization problem for each team. The results of these optimization problems (, having the form of log likelihood values) can be added to the value \bar{L} .

$$\begin{array}{ccc}
 \begin{pmatrix} * & p_{12h} & p_{13h} \\ p_{21h} & * & p_{23h} \\ p_{31h} & p_{32h} & * \end{pmatrix} & \begin{array}{c} p_{11h} \\ \vee \\ p_{21h} \\ \vee \\ p_{31h} \end{array} & \begin{pmatrix} * & p_{12a} & p_{13a} \\ p_{21a} & * & p_{23a} \\ p_{31a} & p_{32a} & * \end{pmatrix} \begin{array}{c} p_{11a} \\ \vee \\ p_{21a} \\ \vee \\ p_{31a} \end{array} \\
 p_{11h} \leq p_{12h} \leq p_{13h} & & p_{11a} \leq p_{12a} \leq p_{13a}
 \end{array}$$

Figure 4.4: Calculation of the upper bound \bar{L}

As mentioned, the algorithm leads for sure to the optimal ordering. The drawback is that despite of the fairly sophisticated upper bound that we are suggesting, it is still not efficient enough to be applied to tournaments with more than 11-12 teams⁵. Nevertheless the branch and bound algorithm deserves to be included in the empirical subsection of this chapter.

4.5.3 Tabu Search

The third ranking algorithm we are suggesting is a heuristic search method. The advantage of tabu search lies in the combination of local search and a diversification mechanism. The local search systematically browses through neighborhood solutions, checking for a possible improvement of the objective value. That the algorithm doesn't get stuck in local optima is assured by a memory structure, avoiding previously visited regions of the solution space, giving a tendency for diversification. A reference with a related application is Laguna, Marti, and Campos (1999).

The algorithm works as follows:

1. Start from a randomly generated order of teams (call it ρ)
2. Calculate the maximum likelihood for the current ranking $L(\rho)$
3. Randomly select a team that is not on the "Tabu List" and remove it from the order
4. Insert the team at position i and calculate difference between the maximum likelihood of the new and the original ranking: $MoveValue = L(\rho') - L(\rho)$
5. Repeat 4. for $1 \leq i \leq n$ except for the original position
6. Insert the team at the position with the highest $MoveValue$
7. Put the team on the "Tabu List" so that it won't be selected for the next "TabuTenure" iterations
8. Go to 2.

⁵Depending on the structure of results in the tournament, as well as the users patience.

Basically what the algorithm does is taking a team from the ranking and trying out every possible position for it, except for the original one. Important is that the best among the new positions is selected even if the "*MoveValue*" is negative. Different convergence criteria are possible for the procedure. Since in our analysis the computational effort in each iteration is fairly large, we use a fixed number of iterations for the algorithm, so that we can best control the amount of time it takes for the algorithm to finish.

4.5.4 Popular ranking methods

Finally, we want to take a more practical approach and compare different ranking systems, which have been used in different fields of sports. We chose the 3 point system (also known as "Three points for a win"), which awards zero points for a loss, one point for a draw and three points for a win. The sum of the points together with the goal difference as a tie breaker then decides upon the ranking. This system has been used in most soccer leagues since it was officially adopted in 1995 by FIFA.⁶ Before the 3 point system was introduced, the analogously structured 2 point system had been widely used in soccer. Here the only difference is that two instead of three points are awarded for a win.

These two systems are fairly easy to apply and (unfortunately) also very similar to each other. As a third candidate for a ranking scheme, we use the Elo rating system. The Elo rating system is a system invented by Elo (1978) originally intended as a rating system for chess. Today it is not only used as for different chess organizations, including the FIDE and the United States Chess Federation, but also the European Go Federation, many different computer games and even the National Collegiate Athletic Association, the organization which is responsible for the organization of many American college sport programs, notable college football and college basketball. The main differences to the three points for a win is that it factors in the strength of the opponent: winning against a strong opponent yields more points than winning against a weak one. This results in the major weakness for our needs: a relatively high number of games is needed to give meaningful results and the order in which the teams play matters a lot.

⁶England introduced the system already in 1981. The first time it was used internationally was in the 1994 World Cup finals.

4.6 Comparing the explanatory power of rankings

To further enhance our comparative analysis of ranking systems, we will apply a statistical hypothesis test. In this test two ranking systems are compared, call them system a and system b . We solve problem Equation (4.4) for both rankings. The p-matrix calculated with the constraints generated by one of the rankings, say system a , will yield a likelihood for the observed season at least as great as the one generated by the other one, say system b .

$$L(\hat{P}_a(w)|w) \geq L(\hat{P}_b(w)|w)$$

where $\hat{P}_a(w)$ and $\hat{P}_b(w)$ are the estimated p-matrices. So we could say, a allows one to calculate a p-matrix with a higher explanatory value, so it must be the better system. But in fact, it might have happened by chance, that this ranking system performed better than the other one. The central question concerns the degree of the odds that a performed better than b by the observed amount. Let us define the likelihood ratio as follows

$$LR_{a,b} = \log(L(\hat{P}_a(w)|w)) - \log(L(\hat{P}_b(w)|w)).$$

We assume a Hypothesis H_0 stating that " b is the correct ranking system". Correct means that it allows us to estimate the right p-values. Using these probabilities for each match, we simulate a complete season and get a new tournament \hat{w} for which we again calculate the likelihoods given $\hat{P}_a(w)$ and $\hat{P}_b(w)$. This way a few thousand seasons are simulated and we receive a distribution over the difference of the log-likelihood. In the ideal case, the probability (suggested by the simulated distribution) of the observed difference between the likelihoods is small enough to be able to reject H_0 with this very test size α .

$$P[LR_{a,b}(\hat{w}) \leq LR_{a,b}(w)] < \alpha$$

So, roughly what we do is assuming that one of the systems is correct, and then we try to reject this hypothesis, by showing that the probability for another system to be as much better as empirically observed is very small.

The weakness of this approach is pretty obvious. We are only able to reject the hypothesis that a particular system is perfectly correct. Even though the data allows us to make a guess about it, the test does not allow us to make a statement about which of the two systems in consideration is actually better. So in fact, both of the systems might be incorrect, but we are only able to reveal the inadequacy of one of them.

4.7 Data

We obtain the data from different sources. For soccer we focus on the German Bundesliga and the British Premier League. For the former we have data from the seasons 1968/69 till 2012/13, for the latter the sample from the seasons 1997/98 till 2012/2013. Additionally, we included the season 2012/2013 from the the Austrian Bundesliga, because of its advantage of having only 10 Teams. The scores for all matches, which are translated to win/draw/loss data, are obtained from the website www.kicker.de. Notable about the soccer data is that each team plays each other team exactly once at home and once away in each season. This introduces a symmetry to the data which, even though it is not necessary, might be considered as desirable and certainly influences the results of our analysis.

Regarding tennis, we face a different situation. Since there is no league of players in which each player faces another one a fixed number of times per season we have to go a different way. We will focus only on the top 10 players according to the official ATP ranking at the end of each year (obtained from <http://www.atpworldtour.com>). Then we collect the data for all the ATP matches played in this season from <http://www.tennis-data.co.uk>. Of course these data sets will be highly asymmetric, because some players play against each other more than once, and some might not face each other at all during a season. Another special fact about the tennis data is that we don't have a real home away situation.⁷ Even more importantly, in tennis there is no possibility of ties. So we face only a binomial distribution for the outcome of each match which considerably facilitates the optimization procedure.

Concerning American football, we will focus exclusively on the NFL. We have data on the scores of every NFL game since 1978 from the website <http://www.repole.com/sun4cast/data.html>. The NFL comprises from 28 in the season 1978 to 32 teams in 2012. This is by far the largest group of teams. Almost naturally it follows that among the samples there is a huge number of teams that don't face each other during a season. Which team is playing which is determined by a complicated system, which shall not be further discussed here. In football draws are possible, but only happen very rarely. Along with the fact that American football enjoys great popularity, this makes NFL data very interesting for our analysis.

⁷Of course some players might feel more at home when a tournament is taking place in their country of origin. But since this is very different to the situation of a team playing in its very own stadium in its city, we will assume that every game takes place on neutral ground.

4.8 Empirical analysis

We now want to apply the presented methods to real data from sports. Countless different types of sports are imaginable and probably the readers preferences for what he would like to see in this section are very heterogeneous. Nevertheless for reasons of space we want to focus on three types, namely soccer, tennis and American football. The main questions we seek to answer are, "Is there a tendency for one of the ranking schemes to be superior to the others according to the criterion we defined?", "If yes, which one is it?", "Does it depend on the type of sport?" and finally "Are we able to improve on the rankings found by the simple ranking methods using one of the algorithms presented in Section 4.5?"

4.8.1 Soccer in Austria: Finding an optimal ranking

With the branch and bound algorithm we find our selves equipped with a very powerful instrument to find the optimal ranking. Unfortunately this algorithm can only be applied to sets of teams that have a limited size. The first object of our investigation shall be the Austrian Bundesliga. Its size of 10 teams enables us to apply the discussed bnb-method. During a season each team plays against each other team four times, two times at home and two times away. This is different from most other soccer leagues, but doesn't increase the computational complexity by much. Here, we consider the season 2012/2013. To draw a first comparison between the performances of the other ranking schemes, Table 4.1 shows the maximum likelihoods that have been calculated.

Method	BnB	2-Point	3-Point	LOP	Elo	Tabu-Search
MLE	-129.844	-131.742	-135.561	-140.024	-131.703	-130.465

Table 4.1: Log likelihood values for the Austrian Bundesliga 12/13

While the ranking corresponding to the solution to the linear ordering problem gives a relatively low likelihood, the two point system as well as the Elo-system seem to explain the results a lot better. Nevertheless, none of the systems generates the optimal ranking found by the branch and bound algorithm. The ranking produced by the Tabu Search gives a higher likelihood than all the systems, but still is not the optimal one.

Figure 4.5 compares the optimal ranking that we found with the actually applied order, namely the 3-Point ranking. One can see that there are indeed some differences. Perhaps most striking is that in this season SV Mattersburg was relegated, while in the optimal ranking Wacker Innsbruck would have been relegated. This team was actually ranked 8th.



Figure 4.5: Rankings resulting from 3-point system and Branch and Bound algorithm

Unfortunately most leagues are larger than the Austrian Bundesliga. The resulting computational effort makes it virtually impossible for us to find optimal rankings, which is why in the next subsection we focus on the other methods and compare the different ranking schemes across panel data from different leagues in different sports.

4.8.2 Ranking systems and maximum likelihood estimates

To give the reader an impression of how a matrix of estimated outcome probabilities for each game looks like after the optimization, Figure 4.6 depicts the probabilities for home game wins for the Bundesliga season 2012/13 estimated using the "three points for a win" system.

	0.18	0.31	0.69	0.8	0.8	0.8	0.8	0.8	0.8	1	1	1	1	1	1	1	1	1	1
	0	0.31	0.69	0.69	0.77	0.77	0.77	0.77	0.77	0.77	0.77	0.77	0.92	0.92	0.92	0.92	0.92	0.92	0.92
	0	0.31	0.46	0.69	0.69	0.69	0.69	0.69	0.71	0.71	0.71	0.71	0.71	0.71	0.92	0.92	0.92	0.92	0.92
	0	0.31	0.31	0.46	0.46	0.46	0.46	0.52	0.71	0.71	0.71	0.71	0.71	0.71	0.92	0.92	0.92	0.92	0.95
	0	0.31	0.31	0.46	0.46	0.46	0.46	0.52	0.7	0.7	0.7	0.7	0.7	0.86	0.92	0.92	0.92	0.92	0.95
	0	0.31	0.31	0.46	0.46	0.46	0.46	0.49	0.49	0.49	0.49	0.49	0.61	0.7	0.85	0.85	0.85	0.85	0.85
	0	0.31	0.31	0.46	0.46	0.46	0.46	0.49	0.49	0.49	0.49	0.57	0.7	0.85	0.85	0.85	0.85	0.85	0.85
	0	0	0.31	0.31	0.31	0.31	0.46	0.48	0.49	0.49	0.49	0.49	0.53	0.58	0.58	0.85	0.85	0.85	0.85
	0	0	0.31	0.31	0.31	0.31	0.46	0.47	0.49	0.49	0.49	0.49	0.49	0.49	0.58	0.85	0.85	0.85	0.85
	0	0	0.31	0.31	0.31	0.31	0.31	0.41	0.41	0.41	0.45	0.49	0.49	0.49	0.49	0.49	0.85	0.85	0.85
	0	0	0.31	0.31	0.31	0.31	0.31	0.41	0.41	0.41	0.41	0.43	0.44	0.44	0.44	0.44	0.44	0.44	0.44
	0	0	0.19	0.31	0.31	0.31	0.31	0.41	0.41	0.41	0.41	0.41	0.42	0.42	0.43	0.44	0.44	0.44	0.44
	0	0	0.19	0.19	0.19	0.19	0.31	0.38	0.41	0.41	0.41	0.41	0.42	0.42	0.42	0.44	0.44	0.44	0.44
	0	0	0.075	0.19	0.19	0.19	0.31	0.36	0.36	0.36	0.36	0.36	0.36	0.42	0.42	0.42	0.42	0.44	0.44
	0	0	0.056	0.19	0.19	0.19	0.19	0.19	0.29	0.29	0.29	0.29	0.36	0.36	0.42	0.42	0.42	0.44	0.44
	0	0	0.037	0.19	0.19	0.19	0.19	0.19	0.29	0.29	0.29	0.29	0.29	0.29	0.29	0.29	0.36	0.44	0.44
	0	0	0	0	0	0.19	0.19	0.19	0.29	0.29	0.29	0.29	0.29	0.29	0.29	0.29	0.29	0.35	0.44
	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0.094	0.17	0.3

Figure 4.6: MLE for p_{ijh} using 3-point system

Generally, a striking feature about the structure of the estimated probability matrices is the occurrence of homogeneous values in certain areas of the matrix, reminding of the layer structure discussed in Section 4.4.1. Remarkable in this particular matrix is the large number of "1"s in the upper right corner and "0"s in the lower left corner. The reader might be tempted to argue that these values are fairly unrealistic, because intuition tells us that even if the strongest team plays the weakest one, in the current case Bayern München against Greuter Fürth, the chance of the former to win against

the latter will be high, but never 100%. The point is that we only hold this intuition, because probably at some point in the past we have seen top teams occasionally loosing against teams that were ranked very low. But since this kind of information is not part of our estimation procedure, it is only natural that estimates look like this.⁸

Next, we want to try to improve this ranking by using one of the algorithms presented in Section 4.5. Unfortunately the sample of 18 teams is too large for an application of the branch and bound algorithm, which would technically allow us to find the optimal ranking. So we use the Tabu search method, which we run for 100 iterations. The resulting ordering as well as the corresponding maximum likelihoods are shown in Figure 4.7.

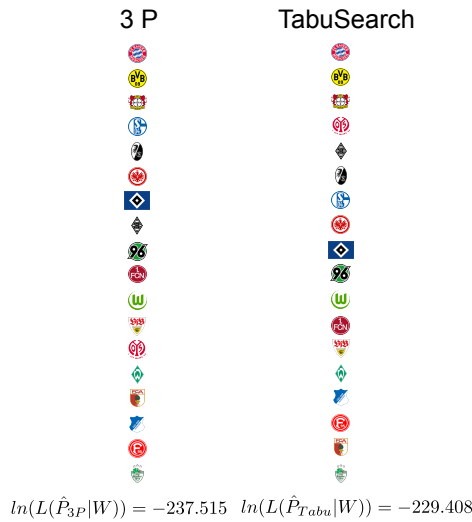



Figure 4.7: Rankings resulting from 3-point system and Tabu Search

The Tabu Search finds a ranking that is partly very different from the one determined using the 3-point system. The biggest difference is the position of "Mainz 05" jumping from the 13th position to the 4th. The cause of this difference can only be that "Mainz 05" has won the matches in this season that were particularly important in the sense of being in accordance with the team having fairly high winning probabilities in general. However, despite of differences in parts, a great similarity between the rankings can be observed. This similarity can be measured using Spearman's rank correlation coefficient defined as $\rho = 1 - \frac{6 \sum (r_i - s_i)^2}{n(n^2 - 1)}$ with r_i being the original (3 point) ranking of team i and s_i the ranking with the highest maximum likelihood as calculated with the Tabu Search algorithm. The correlation between the 3-point ranking and the one found by the Tabu Search is indeed fairly high with a value of about 0.87616. The difference between the

⁸We have to add, that in case the reader has seen Bayern München play in the season 2012/13, he most certainly would agree that estimating some probabilities in the right of the upmost row with a value of 1 most probably only involves a very small error.

maximum likelihood values however, is in fact very large. The probabilities found using the Tabu Search ranking make the observed season 3318 times more likely compared to the probabilities found using the 3 point ranking.


As mentioned above, we have data not only on this one Bundesliga season, but on the ones from the last 50 years.⁹ For every season that we have data on, we calculated the maximum likelihood p-matrices as well as the objective function values using the "2 points for a win", the "3 points for a win", Elo system and the ranking from the solution to the linear ordering problem. Finally, we used the Tabu search method to find out, whether or not one is able to improve on one or all of the ranking schemes. Because from season to season the likelihood values fluctuate heavily, it makes sense to use the likelihood found by one of the systems as a reference value and plot the differences to these values in a diagram. As opposed to just plotting the absolute likelihoods of every system in each year, this technique allows us to better compare the quality of the rankings throughout the panel data. The system of reference will be "2 points for a win".



(a) Bundesliga

Figure 4.8 (a) and (b) reveal that the two and three point systems are in fact very close in the maximum likelihoods they "produce". This is not least because in most cases the rankings determined by the two systems only differ in a few spots. And if the rankings do not differ much, it's only natural that the likelihood values won't be very far apart either. The two point system allows for a calculation of p-matrices that make the observed seasons on average across the Bundesliga samples by about 9.8% more likely than when using the three point system. In the Premier League the three point system has a 5.2% higher explanatory power. The Elo-system also gives us likelihoods

⁹Because in the seasons 1963/64, 1964/65 and 1991/92 the number of teams in the Bundesliga was different from 18, we excluded these seasons from the sample. Sacrificing these three data points for a higher comparability seems reasonable.



(b) Premier League

Figure 4.8: Maximum Likelihoods for Bundesliga and Premier League panel data

in the same range, indicated by the green lines. Actually this is a bit of a surprise, since there were some hopes that the intuitively very reasonable mechanism of getting more points for winning against relatively strong opponents would enable us to explain the observed results better. Still it is not worse than the conventional two and three point systems. But because of its higher complexity we clearly refrain from making a recommendation for using this system. The ranking resulting from solving the linear ordering problem is by far the worst performer in the diagram. One observes it to yield likelihoods that are on average more that 1000 times smaller that the ones from the two point system. So we have to clearly reject the suggestion for a possible application of the LOP in soccer that has been made in the literature.

Another striking feature about the graphs is the position of the likelihood curves corresponding to the tabu search. The heuristic algorithm is able to improve on every single ranking from the sample, except for the Premier League seasons 04/05 and 05/06. On average it helps to explain the results about 457 times better. The graph shows us that even though the simple ranking schemes produce fairly "good" orderings in the sense of a high correlation (as seen above), they are far away from being the most likely correct ones.

Next, Figure 4.9 shows the analogue results for the tennis panel data from the last 14 years. The first thing to note is that the two and three point systems produce the same likelihoods throughout the whole sample, which is why in this graph there is no curve comparing the two, since it would lie on the x-axis. The reason for this is that in tennis we do not have draws, so in both systems the players are only ranked according to their number of victories. In Figure 4.9, in addition to the curves from Figure 4.8, the likelihoods from the official ATP ranking from the end of each year are listed.




Figure 4.9: Maximum Likelihoods for ATP panel data

This ranking is determined by awarding different amounts of points for a stage that is reached in the Grand Slam Tournaments, the ATP World Tour Finals, the Masters 1000, Olympics etc. Of course this method is very sophisticated and includes also the results of the matches of the top 10 players against others that might not be in the top 10. This data is not part of the other systems we are analyzing. According to the criterion of this work, the ATP ranking performs fairly bad in explaining the observed results. Interestingly in this tennis sample, the linear ordering ranking produces fairly high likelihoods, in fact on average higher ones than the n-point and Elo system. Again, in every year the tabu search algorithm is able to improve on all of the discussed rankings.

Finally, Figure 4.10 illustrates the results from the same calculations as above, now for American football results from the National Football League in the US. There is little difference between 2- and 3 point systems, because draws are very unlikely to occur. However, the 3 point system is almost at every point at least as good as the 2 point system. The LOP and Elo systems operate in the same range of likelihoods as well. With NFL data, applying the tabu search is more effortful and thus takes more time for the same number of iterations, because of the higher number of teams. However, again, the tabu search improves upon all the rankings in the sample.

In general, the difference in the relative likelihoods when applying the Elo/LOP system and the n-point systems between soccer on the one hand and tennis and American football on the other could be due to the heterogeneity in the number of games played between the teams in tennis and football as opposed to the symmetric situation in soccer. Certainly a system like "two points for a win" doesn't seem to be particularly well suited in a situation where teams play different amounts of matches. And as




Figure 4.10: Maximum Likelihoods for NFL panel data

explained further above, here it could be justified to give 1 and -1 points instead of 0 and 2 for a win and a loss, respectively. However, implementing this changes not much and even reduces the average likelihood a bit. Another explanation could be the sport itself. It might be due to the result generating probabilities themselves, that for one sport different ranking schemes are better suited than others. Indeed, it is easy to show that in the space of transitive probability matrices, there are areas where each of the considered systems is most likely to generate a ranking closest to the real one. This is an interesting direction for further theoretical research.

4.8.3 Hypothesis testing

Now we are going further in the analysis of ranking systems than just observing which ordering scheme is able to generate a higher maximum likelihood value. We will consider two examples, which will help deepen the understanding of the problem, but will also clearly highlight the limitations of this hypothesis testing approach, as described in Section 4.6.

Consider the Bundesliga season 2011/2012. Looking at Figure 4.8 reveals that for this data set the 3 point system performed better than the 2 point system. The difference between the two maximum likelihood logs is 0.564. But the central question is "did this MLE difference appear because the underlying unobservable probabilities make the 3 point system more appropriate than the 2 point system in this season or could it in fact be the other way around with the observation just happening by chance?".

To answer this question, assume the correctness of the Hypothesis H_0 : "The 2 point system puts the teams in the correct order". We will test H_0 against the alternative

Hypothesis H_1 : "The 3 point system puts the teams in the correct order". Now, for the two systems the probability matrices $\hat{P}_{2p}(w)$ and $\hat{P}_{3p}(w)$ are estimated. Using $\hat{P}_{2p}(w)$ 5000 seasons are simulated. Then $L(\hat{P}_{2p}(w)|\hat{w})$ and $L(\hat{P}_{3p}(w)|\hat{w})$ are calculated for each of the seasons. Their respective frequency distribution is depicted in Figure 4.11 (a) and (b). The distribution of their difference, which corresponds to the ratio of the likelihoods without logs is plotted in Figure 4.11 (c).

Looking closely at the first two diagrams reveals that the distribution of $L(\hat{P}_{3p}(w)|\hat{w})$ is shifted a little bit to the left relative to the one of $L(\hat{P}_{2p}(w)|\hat{w})$. This is intuitively correct because it is only natural that the probability matrix that generated the seasons of the sample gives the higher likelihood values than the matrix $\hat{P}_{3p}(w)$, which has nothing to do with the season simulation. Now to find out the confidence level with which we would be able to reject H_0 one has to compare the observed likelihood ratio to the likelihood ratio distribution in Figure 4.11 (c). This procedure shows us that assuming the correctness of H_0 , the probability of the likelihood ratio being ≤ 0.564 is only 11%. So we are able to reject the hypothesis that the 2 point system gives the correct ranking with test size $\alpha = 0.11$, meaning that the probability of not making an error of the first kind is 0.89. One has to be careful not to misinterpret this result. It means that we are able to reject the hypothesis that the 2 point system gives the correct ranking. However, this does by no means imply that the 3 point system gives the correct ranking.


Now let us conduct a second hypothesis test, this time using tennis data. A good experiment would be to test for the correctness of the LOP system against the 3 point/2 point system in the year 2012. In this year the LOP produced a considerably higher likelihood than the 2 point system (see Figure 4.9), so we would like to know if this was just a random result or if we can actually conclude that the underlying probabilities favor the LOP scheme in the sense of telling us the truth about the ordering of tennis players. The hypothesis are:

1. H_0 : "The 2 point system puts the teams in the correct order"
2. H_1 : "The solution to the LOP puts the teams in the correct order"


Assuming the correctness of H_0 , we again estimate the probability matrices and then simulate 5000 seasons. Hereby we always assume that the m_{ij} values stay constant, i.e., the amount of times players meet is the same in every simulation. We proceed as above by calculating the test statistic for the likelihood ratio and then comparing it to the empirically observed one. We have:

$$LR_{2p,LOP} = \log(L(\hat{P}_{2p}(w)|w)) - \log(L(\hat{P}_{LOP}(w)|w)) = -6.9103$$


The simulated test statistic tells us that in case H_0 is correct, the probability of an



(a) $L(\hat{P}_{3n}(w)|\hat{w})$



(b) $L(\hat{P}_{3p}(w)|\hat{w})$



(c) $LR_{2p,3p}(\hat{w})$

Figure 4.11: Simulated test statistic for Bundesliga hypothesis test

occurrence of such a small likelihood ratio is only 0.02%. It follows that we can reject H_0 with test size $\alpha = 0.02$ (i.e., a confidence level of 99.98%).

This method is, as already mentioned in Section 4.6, only useful to test if a ranking scheme gives us the correct ranking. If one is very confident that a ranking system gives a very good approximation to the optimal ranking, one could test the results of this ranking system against the optimal ranking (either approximated by the Tabu Search Algorithm or calculated by the Branch and Bound Algorithm). If this test fails to reject that the ranking system is the correct ranking system, one could be very certain, that there is no (significantly) better ranking system. However, in the data we have analyzed so far there has been no candidate for this good ranking system and every test performed like this would lead to rejection of the hypothesis.

4.9 Conclusion

We constructed a statistical model describing the outcomes of sports matches. The model assumes a transitive relationship between the relative strengths of the teams. The resulting constraints turn out to be very restrictive, which is illustrated by the rapidly shrinking size of the parameter space shown in appendix B. The incorporation of ties as well as home/away asymmetries makes our model much more complicated than the related isotonic regression problem. The discussed branch and bound algorithm is capable of solving the problem for up to 12 teams. For larger data sets, a tabu search heuristic has been proposed. The empirical subsection of the chapter first illustrates the structure of an optimized probability matrix with an example. We have shown that in the example the maximum likelihood produced by the tabu search is more than 3000 times higher than the one resulting from an application of the 3-point system. But this does not mean that the two rankings are strongly uncorrelated as seen from the high value of Spearman's rank correlation coefficient. Panel data has been used to compare different ranking systems in three types of sports. In soccer, data from German Bundesliga and English Premier League have shown that the 2- and 3-point systems are very close to each other in the maximum likelihoods they produce, which is not a surprise when considering their structural similarity. Hopes were higher for the performance of the Elo system, because as opposed to the traditional point systems it considers the opponents strength. However, on average the generated MLEs were in the same range as the ones from the n-point systems. This result also applies for ATP tennis and NFL American football data. So the additional degree of complexity seems to be enough of a justification for not giving a recommendation towards an introduction of the Elo system. A difference worth mentioning is that the ranking,

which results from the LOP performs fairly well in tennis and American football, but worse than everything else in soccer. We show that almost in every sample across all considered types of sports we are able to improve on the rankings produced by the considered systems by using tabu search. This illustrates that there might be a system that is much better at finding the most likely correct ranking, possibly without the inclusion of a great complexity. As a final remark, we want to mention that the framework presented in this chapter has its natural limitations and leaves out many important aspects that should be considered when choosing or designing a ranking scheme. Things like opponents incentives during a match and the resulting effects on the observers level of thrill or the occurrence of winning decision as late as possible during a season could be interesting points for further research.

4.10 Appendix

4.10.1 Proofs

Proof of Proposition 4: Ignoring the away/home differentiation, we can write p_{ikx} as p_{ik} . With 0 probabilities of draws, Equation (4.1) is now

$$p_{ik} = 1 - p_{ki}$$

and therefore Equation (4.2) is then equivalent to

$$\begin{aligned} p_{ik} \geq p_{il} &\Leftrightarrow p_{jk} \geq p_{jl} \\ p_{i'k'} > p_{j'k'} &\Rightarrow p_{i'j'} > p_{i'j'} \end{aligned} \quad (4.5)$$

Now we have to show that $(SST) \Rightarrow (4.5)$ and $(4.5) \Rightarrow (SST)$.

$(SST) \Rightarrow (4.5)$:

We are dividing this case into two cases: For $p_{ik} \geq \frac{1}{2} \geq p_{jk}$ we can see:

$$p_{ij} \geq p_{kj} = 1 - p_{jk} \geq \frac{1}{2} \xrightarrow{SST} p_{ix} \geq p_{jx} \quad \forall x$$

For every other case we can assume wlog that $p_{ik} \geq p_{jk} \geq \frac{1}{2}$

$$\begin{aligned} p_{ij} &\geq \frac{1}{2} \xrightarrow{SST} p_{ix} \geq p_{jx} \quad \forall x \\ p_{ij} < \frac{1}{2} &\Rightarrow p_{ji} > \frac{1}{2} \xrightarrow{SST} p_{jk} > p_{ik} \end{aligned}$$

Which is a contradiction to the assumption, therefore $p_{ij} \geq \frac{1}{2}$.

$(4.5) \Rightarrow (SST)$:

$$\begin{aligned} p_{jk} > p_{ik} &\xrightarrow{(5)} p_{li} > p_{lj} \quad \forall l \\ \Rightarrow p_{ii} > p_{ij} &\xrightarrow{p_{ii}=1/2} p_{ij} < \frac{1}{2} \end{aligned}$$

□

Proof of proposition Proposition 5: Define a ranking from best to worst $\rho(i) : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ such that $p_{ikx} \geq p_{jkx} \Rightarrow \rho(i) < \rho(j)$ and $p_{kix} \leq p_{kix} \Rightarrow \rho(i) < \rho(j)$.

$$\begin{aligned} p_{ikx} \geq p_{jkx} &\Leftrightarrow \rho(i) < \rho(j) \Leftrightarrow p_{ily} \geq p_{jly} \quad \forall i, j, k, l, x, y \\ p_{kix} \geq p_{kix} &\Leftrightarrow \rho(j) < \rho(i) \Leftrightarrow p_{ily} \geq p_{jly} \quad \forall i, j, k, l, x, y \\ p_{i'k'x} > p_{j'k'x} &\Leftrightarrow \rho(j) < \rho(i) \text{ and } \rho(i) \not\leq \rho(j) \exists i', j', k', x \end{aligned}$$

□

4.10.2 Parameter space

In this subsection we explore the effect of transitivity conditions on the parameter space of winning probabilities to illustrate the limitations enforced by it. To do that we compare the size of the parameter space with transitivity to the space of unrestricted winning probabilities \bar{S}_n , e.g. every p_{ij}, p_{ji} fulfilling $p_{ij} + p_{ji} = 1$.

The space of parameters including the transitivity conditions is a subset of this set \bar{S}_n . $S_n(R)$ is hereby defined as the size of this space relative to \bar{S}_n only considering the restrictions for $p_{ij} \in R$. The unrestricted parameter space is in this simple case: $\bar{S}_n = [0, 1]^{\frac{n(n-1)}{2}}$ which can be easily seen by the fact that every p_{ji} is completely determined by p_{ij} . The restricted space for n players and the transitivity conditions for every $(i, j) \in K_n$ with $K_n = \{(i, j) | i, j \in \{1, 2, \dots, n\}, i < j\}$ is therefore

$$S_n(K_n) = \int_{b_{i+1,j}}^{b_{i,j+1}} S_n(K_n \setminus \{(i, j)\}) dp_{ij}$$

with

$$S_n((i_0, j_0)) = \int_{b_{i_0+1,j_0}}^{b_{i_0,j_0+1}} dp_{i_0j_0}$$

and

$$b_{i,j} := \begin{cases} p_{ij}, & \text{for } (i, j) \in K_n \\ 0.5, & \text{for } i = j \\ 0, & \text{else} \end{cases}$$

As this fairly complicated recursive integral may be hard to interpret, Table 4.2 gives the values for the relative size of the transitive parameter space for up to five teams. It can be seen that the size rapidly shrinks and it is not hard to imagine that for a league comprising e.g. 18 teams the conditions are in this sense very strict.

n	2	3	4	5	6	7
Relative size	1	$\frac{1}{4}$	$\frac{1}{120}$	$\frac{1}{40320}$	$\frac{1}{203212800}$	$\frac{1}{19313344512000}$
Approximation	1	0.25	8.3×10^{-3}	2.5×10^{-5}	4.9×10^{-9}	5.2×10^{-14}

Table 4.2: Relative size of the transitive parameter space

4.10.3 Code

The first code listing shows the problem definition of an optimization with fixed team ordering, so that the ipopt framework will understand it.

```

1 #include "nfl_nlp.hpp"
2
3 #include <cassert>
4 #include <iostream>
5 #include <math.h>
6
7 static int t=50;
8
9 static int w[3][50][50];
10 static double p[2][50][50];
11
12
13 using namespace Ipopt;
14
15 // constructor
16 nfl_NLP::nfl_NLP(int myw[][50][50], double* myp[][50][50], double*& ←
    zielwert, int myt)
17 {
18     zielwert = &zw;
19     t=myt;
20
21     for (int h=0; h<3; h++) {
22         for (int k=0; k<t; k++) {
23             for (int l=0; l<t; l++) {
24                 w[ht][k][l]=myw[ht][k][l];
25                 myp[ht][k][l]=&p[ht][k][l];
26             }
27         }
28     }
29
30 }
31
32 //destructor
33 nfl_NLP::~~nfl_NLP()
34 {}
35
36 // returns the size of the problem
37 bool nfl_NLP::get_nlp_info(Index& n, Index& m, Index& nnz_jac_g,
38                             Index& nnz_h_lag, IndexStyleEnum& ←
39                             index_style)

```

```

40 // The problem described in nfl_NLP.hpp has 4 variables, x[0] ←
    // through x[3]
41 n = 2*pow(t,2);
42
43 // one equality constraint and one inequality constraint
44 m = pow(t,2) + 4*t*(t-1);
45
46 // in this example the jacobian is dense and contains 8 nonzeros
47 nnz_jac_g = 2*m;
48
49 // the hessian is also dense and has 16 total nonzeros, but we
50 // only need the lower left corner (since it is symmetric)
51 nnz_h_lag = 2*n-4*t;
52
53 // use the C style indexing (0-based)
54 index_style = TNLP::C_STYLE;
55
56 return true;
57 }
58
59 // returns the variable bounds
60 bool nfl_NLP::get_bounds_info(Index n, Number* x_l, Number* x_u,
61                               Index m, Number* g_l, Number* g_u)
62 {
63 // here, the n and m we gave IPOPT in get_nlp_info are passed back ←
    // to us.
64 // If desired, we could assert to make sure they are what we think ←
    // they are.
65
66 // the variables have lower bounds of 0
67 for (Index i=0; i<2*t*t; i++) {
68     x_l[i] = 0.0;
69 }
70
71 // the variables have upper bounds of 1
72 for (Index i=0; i<2*t*t; i++) {
73     x_u[i] = 1.0;
74 }
75
76
77 Index i = 0;
78 for (Index k=0; k<t; k++) {
79     for (Index l=0; l<t; l++) {
80         g_l[i] = -2e19;
81         g_u[i] = 1.0;
82         i++;

```



```

83     }
84 }
85 for (Index h=0; h<2; h++) {
86     for (Index k=0; k<t; k++) {
87         for (Index l=0; l<t; l++) {
88             if (l<t-1) {
89                 g_l[i] = -2e19;
90                 g_u[i] = 0.0;
91                 i++;
92             }
93             if (k<t-1) {
94                 g_l[i] = -2e19;
95                 g_u[i] = 0.0;
96                 i++;
97             }
98         }
99     }
100 }
101
102 return true;
103 }
104
105 // returns the initial point for the problem
106 bool nfl_NLP::get_starting_point(Index n, bool init_x, Number* x,
107                                 bool init_z, Number* z_L, Number* ←
108                                 z_U,
109                                 Index m, bool init_lambda,
110                                 Number* lambda)
111 {
112     assert(init_x == true);
113     assert(init_z == false);
114     assert(init_lambda == false);
115
116     // initialize to the given starting point
117     for (Index i=0; i<2*t*t; i++) {
118         x[i] = 0.4;
119     }
120
121     return true;
122 }
123
124 // returns the value of the objective function
125 bool nfl_NLP::eval_f(Index n, const Number* x, bool new_x, Number& ←
126     obj_value)
127 {
128     assert(n == 2*t*t);

```

```

127
128   obj_value=0;
129   Index i=0;
130   for (Index k=0; k<t; k++) {
131       for (Index l=0; l<t; l++) {
132           if (k!=l) {
133               obj_value += (-w[0][k][l]*log(x[t*k+l]+0.000001) - w[1][l][k]*log(x[t*t+t*l+k]+0.000001) - (w[2][k][l]-w[0][k][l]-w[1][l][k])*log(1-x[t*k+l]-x[t*t+t*l+k]+0.000001));
134
135               i++;
136           }
137       }
138   }
139
140   return true;
141 }
142
143 // return the gradient of the objective function grad_{x} f(x)
144 bool nfl_NLP::eval_grad_f(Index n, const Number* x, bool new_x, Number*& grad_f)
145 {
146     Index i=0;
147     for (Index h=0; h<2; h++) {
148         for (Index k=0; k<t; k++) {
149             for (Index l=0; l<t; l++) {
150                 if (k!=l) {
151                     if (h==0)
152                         grad_f[i] = -w[ht][k][l]/(x[t*t*h+t*k+l]+0.000001) + (w[2][k][l]-w[ht][k][l]-w[1-h][l][k])/(1-x[t*t*h+t*k+l]-x[t*t*(1-h)+t*l+k]+0.000001);
153                     else
154                         grad_f[i] = -w[ht][k][l]/(x[t*t*h+t*k+l]+0.000001) + (w[2][l][k]-w[ht][k][l]-w[1-h][l][k])/(1-x[t*t*h+t*k+l]-x[t*t*(1-h)+t*l+k]+0.000001);
155                 }
156                 else {
157                     grad_f[i] = 0;
158                 }
159                 i++;
160             }
161         }
162     }

```

```

163
164     return true;
165 }
166
167 // return the value of the constraints: g(x)
168 bool nfl_NLP::eval_g(Index n, const Number* x, bool new_x, Index m, ←
    Number* g)
169 {
170     Index i = 0;
171     for (Index k=0; k<t; k++) {
172         for (Index l=0; l<t; l++) {
173             g[i] = x[t*k+1]+x[t*t+t*l+k];
174             i++;
175         }
176     }
177     for (Index h=0; h<2; h++) {
178         for (Index k=0; k<t; k++) {
179             for (Index l=0; l<t; l++) {
180                 if (l<t-1) {
181                     g[i] = x[t*t*h+t*k+1]-x[t*t*h+t*k+1+1];
182                     i++;
183                 }
184                 if (k<t-1) {
185                     g[i] = x[t*t*h+t*(k+1)+1]-x[t*t*h+t*k+1];
186                     i++;
187                 }
188             }
189         }
190     }
191     return true;
192 }
193
194 // return the structure or values of the jacobian
195 bool nfl_NLP::eval_jac_g(Index n, const Number* x, bool new_x,
196     Index m, Index nele_jac, Index* iRow, Index ←
        *jCol,
197     Number* values)
198 {
199     if (values == NULL) {
200         // return the structure of the jacobian
201
202         // this particular jacobian is dense
203
204
205         Index z=0;
206         Index r=0;

```

```

207     for (Index k=0; k<t; k++) {
208         for (Index l=0; l<t; l++) {
209             iRow[z] = r;
210             jCol[z] = t*k+l;
211             z++;
212             iRow[z] = r;
213             jCol[z] = t*t+t*l+k;
214             z++;
215
216             r++;
217         }
218     }
219     for (Index h=0; h<2; h++) {
220         for (Index k=0; k<t; k++) {
221             for (Index l=0; l<t; l++) {
222                 if (l<t-1) {
223                     iRow[z] = r;
224                     jCol[z] = t*t*h+t*k+l;
225                     z++;
226                     iRow[z] = r;
227                     jCol[z] = t*t*h+t*k+l+1;
228                     z++;
229                     r++;
230                 }
231                 if (k<t-1) {
232                     iRow[z] = r;
233                     jCol[z] = t*t*h+t*k+l;
234                     z++;
235                     iRow[z] = r;
236                     jCol[z] = t*t*h+t*(k+1)+l;
237                     z++;
238                     r++;
239                 }
240             }
241         }
242     }
243     assert(z==nele_jac);
244 }
245 else {
246     Index z=0;
247     Index r=0;
248     for (Index k=0; k<t; k++) {
249         for (Index l=0; l<t; l++) {
250             values[z] = 1;
251             z++;
252             values[z] = 1;

```

```

253         z++;
254
255         r++;
256     }
257 }
258 for (Index h=0; h<2; h++) {
259     for (Index k=0; k<t; k++) {
260         for (Index l=0; l<t; l++) {
261             if (l<t-1) {
262                 values[z] = 1;
263                 z++;
264                 values[z] = -1;
265                 z++;
266                 r++;
267             }
268             if (k<t-1) {
269                 values[z] = -1;
270                 z++;
271                 values[z] = 1;
272                 z++;
273                 r++;
274             }
275         }
276     }
277 }
278 assert(z==nele_jac);
279 }
280
281 return true;
282 }
283
284 //return the structure or values of the hessian
285 bool nfl_NLP::eval_h(Index n, const Number* x, bool new_x,
286                     Number obj_factor, Index m, const Number* lambda←
287                     ,
288                     bool new_lambda, Index nele_hess, Index* iRow,
289                     Index* jCol, Number* values)
290 {
291     if (values == NULL) {
292         // return the structure. This is a symmetric matrix, so we fill←
293         // the lower left
294         // triangle only.
295         Index i=0;
296         for (Index h=0; h<2; h++) {
297             for (Index k=0; k<t; k++) {
298                 for (Index l=0; l<t; l++) {

```

```

297         if (k!=1) {
298             iRow[i] = t*t*h+t*k+1;
299             jCol[i] = t*t*h+t*k+1;
300             i++;
301
302             iRow[i] = t*t*h+t*k+1;
303             jCol[i] = t*t*(1-h)+t*1+k;
304             i++;
305         }
306     }
307 }
308 }
309 }
310 else {
311     // return the values. This is a symmetric matrix, fill the ←
312     // lower left
313     // triangle only
314
315     Index i=0;
316     for (Index h=0; h<2; h++) {
317         for (Index k=0; k<t; k++) {
318             for (Index l=0; l<t; l++) {
319                 if (k!=1) {
320                     if (h==0){
321                         values[i] = obj_factor *( w[ht][k][l]*pow(x←
322                             [t*t*h+t*k+1]+0.000001, -2) + (w[2][k][l]←
323                             [ht][k][l]-w[1-h][l][k]) *pow((1-x[t*t*←
324                             h+t*k+1]-x[t*t*(1-h)+t*1+k]+0.000001)←
325                             ,-2));
326                         i++;
327                         //std::cout << "values[" << i << "]" = " << ←
328                         values[i] << std::endl;
329
330                         values[i] = obj_factor *0.5*( (w[2][k][l]-w←
331                             [ht][k][l]-w[1-h][l][k]) *pow((1-x[t*t*←
332                             h+t*k+1]-x[t*t*(1-h)+t*1+k]+0.000001), -2))←
333                             ;
334                     }
335                     else{
336                         values[i] = obj_factor *( w[ht][k][l]*pow(x←
337                             [t*t*h+t*k+1]+0.000001, -2) + (w[2][l][k]←
338                             [ht][k][l]-w[1-h][l][k]) *pow((1-x[t*t*←
339                             h+t*k+1]-x[t*t*(1-h)+t*1+k]+0.000001)←
340                             ,-2));
341                         i++;
342                         //std::cout << "values[" << i << "]" = " << ←
343                         values[i] << std::endl;

```

```

329         values[i] = obj_factor * 0.5 * ( (w[2][1][k] - w[1][1][k] - w[1-h][1][k]) * pow((1 - x[t*t*h+t*k+1] - x[t*t*(1-h)+t*1+k] + 0.000001), -2))
330         ;
331     }
332     i++;
333 }
334 }
335 }
336 }
337
338     return true;
339 }
340
341 void nfl_NLP::finalize_solution(SolverReturn status,
342                                Index n, const Number* x, const Number* z_L, const Number* z_U,
343                                Index m, const Number* g, const Number* lambda,
344                                Number obj_value,
345                                const IpoptData* ip_data,
346                                IpoptCalculatedQuantities* ip_cq)
347 {
348     // here is where we store the solution to variables
349     // so we could use the solution.
350
351     for (Index h=0; h<2; h++) {
352         for (Index k=0; k<t; k++) {
353             for (Index l=0; l<t; l++) {
354                 //std::cout << x[t*t*h+t*k+1] << " ";
355                 p[ht][k][l] = x[t*t*h+t*k+1];
356             }
357             //std::cout << " " << std::endl;
358         }
359         //std::cout << " " << std::endl;
360     }
361
362     zw = obj_value;
363 }

```

The second code snippet shows how the program reads 5 NFL seasons, puts them in different orders, and then optimizes the probabilities and prints them.

```

1 int main(int argv, char* argc[])

```

```

2 {
3
4 // Create a new instance of the nlp
5 SmartPtr<TNLP> mynlp;
6 SmartPtr<TNLP> mynlp2;
7
8 // Create a new instance of IpoptApplication
9 SmartPtr<IpoptApplication> app = IpoptApplicationFactory();
10
11 // Change some options
12 app->Options()->SetIntegerValue("print_level", 0);
13 app->Options()->SetNumericValue("tol", 1e-4);
14 app->Options()->SetStringValue("mu_strategy", "adaptive");
15 app->Options()->SetStringValue("output_file", "ipopt.out");
16
17 // Intialize the IpoptApplication and process the options
18 ApplicationReturnStatus status;
19 status = app->Initialize();
20 if (status != Solve_Succeeded) {
21     std::cout << std::endl << std::endl << "*** Error during ↵
22         initialization!" << std::endl;
23     return (int) status;
24 }
25
26 srand(time(NULL));
27
28 double* z; //Variable for the likelihood
29 int tempW50[3][50][50]; //temporary result matrices
30
31 for (Index h=0; h<3; h++)
32     for (Index k=0; k<50; k++)
33         for (Index l=0; l<50; l++)
34             tempW50[ht][k][l]=w050[ht][k][l];
35
36 double l[14][6];
37
38 //for the years 2000 till 2005, the nfl data is read from the files
39 for (int jahr=100; jahr<105; jahr++) {
40     int t;
41     if (jahr<95)
42         t=28;
43     else if (jahr<99)
44         t=30;
45     else if (jahr<102)
46         t=31;

```



```

47     else
48         t=32;
49
50     t=10;
51
52
53     einlesenNfl(jahr,t);
54
55     //order according to the 2 point system and run the ↔
56         optimization
57     ordneNachPunkteSystem50(2,w50,t);
58     mynlp2 = new nfl_NLP(w50,p50,z,t);
59     app->OptimizeTNLP(mynlp2);
60     l[jahr][1]=-*z;
61
62     //order according to the 3 point system and run the ↔
63         optimization
64     ordneNachPunkteSystem50(3,w50,t);
65     mynlp2 = new nfl_NLP(w50,p50,z,t);
66     app->OptimizeTNLP(mynlp2);
67     l[jahr][2]=-*z;
68
69     //order according to the LOP system and run the optimization
70     ordneNachLOP50(w50,t);
71     mynlp2 = new nfl_NLP(w50,p50,z,t);
72     app->OptimizeTNLP(mynlp2);
73     l[jahr][3]=-*z;
74
75     //order according to the ELO system and run the optimization
76     ordneNachSchach50(w50,t);
77     mynlp2 = new nfl_NLP(w50,p50,z,t);
78     app->OptimizeTNLP(mynlp2);
79     l[jahr][4]=-*z;
80
81     //run the tabu search for 100 iterations and then run the ↔
82         optimization
83     l[jahr][5]=tabuSearch50(100,t);
84
85     //print out the results
86     cout<<l[jahr][0]<<" "<<l[jahr][1]<<" "<<l[jahr][2]<<" "<<l[jahr]↔
87         |[3]<<" "<<l[jahr][4]<<" "<<l[jahr][5]<<endl;
88     cout<<endl;
89 }

```

Bibliography

- Anat R Admati and Motty Perry. Joint projects without commitment. *The Review of Economic Studies*, 58(2):259–276, 1991.
- Marina Agranov, Elizabeth Potamites, Andrew Schotter, and Chloe Tergiman. Beliefs and endogenous cognitive levels: An experimental study. *Games and Economic Behavior*, 75(2):449–463, 2012.
- George A Akerlof. Procrastination and obedience. *The American Economic Review*, pages 1–19, 1991.
- Larbi Alaoui and Antonio Penta. Endogenous depth of reasoning. *The Review of Economic Studies*, page rdv052, 2015.
- Carlos Alós-Ferrer and Christoph Kuzmics. Hidden symmetries and focal points. *Journal of Economic Theory*, 148(1):226–258, 2013.
- Ayala Arad and Ariel Rubinstein. The 11–20 money request game: a level-k reasoning study. *The American Economic Review*, 102(7):3561–3573, 2012.
- Michael Bacharach. *Variable Universe Games*. MIT Press, 1993.
- Sandeep Baliga and Tomas Sjöström. Arms races and negotiations. *The Review of Economic Studies*, 71(2):351–369, 2004.
- RE Barlow and HD Brunk. The isotonic regression problem and its dual. *Journal of the American Statistical Association*, 67(337):140–147, 1972.
- Dirk Bergemann and Ulrich Hege. The financing of innovation: learning and stopping. *RAND Journal of Economics*, pages 719–752, 2005.
- V Bhaskar. Egalitarianism and efficiency in repeated symmetric games. *Games and Economic Behavior*, 32(2):247–262, 2000.
- Cristina Bicchieri. *The grammar of society: The nature and dynamics of social norms*. Cambridge University Press, 2005.

- Andreas Blume. Coordination and learning with a partial language. *Journal of Economic Theory*, 95(1):1–36, 2000.
- Andreas Blume and Uri Gneezy. An experimental investigation of optimal learning in coordination games. *Journal of Economic Theory*, 90(1):161–172, 2000.
- Andreas Blume and Uri Gneezy. Cognitive forward induction and coordination without common knowledge: An experimental study. *Games and Economic Behavior*, 68(2):488–511, 2010.
- Patrick Bolton and Christopher Harris. Strategic experimentation. *Econometrica*, 67(2):349–374, 1999.
- Alessandro Bonatti and Johannes Hörner. Collaborating. *American Economic Review*, 101(2):632–63, 2011.
- Allison A Brennan and James T Enns. When two heads are better than one: Interactive versus independent benefits of collaborative cognition. *Psychonomic bulletin & review*, 22(4):1076–1082, 2015.
- Hugh D Brunk. Maximum likelihood estimates of monotone parameters. *The Annals of Mathematical Statistics*, 26(4):607–616, 1955.
- Hugh D Brunk. Mathematical models for ranking from paired comparisons. *Journal of the American Statistical Association*, 55(291):503–520, 1960.
- Antonio Cabrales, Rosemarie Nagel, and Roc Armenter. Equilibrium selection through incomplete information in coordination games: an experimental study. *Experimental Economics*, 10(3):221–234, 2007.
- Colin F Camerer, Teck-Hua Ho, and Juin-Kuan Chong. A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, pages 861–898, 2004.
- Hans Carlsson and Eric Van Damme. Global games and equilibrium selection. *Econometrica: Journal of the Econometric Society*, pages 989–1018, 1993.
- André Casajus. Focal points in framed strategic forms. *Games and Economic Behavior*, 32(2):263–291, 2000.
- André Casajus. Static focal points in framed strategic games. In *Focal Points in Framed Games*, pages 9–54. Springer, 2001.
- Irène Charon and Olivier Hudry. An updated survey on the linear ordering problem for weighted or unweighted tournaments. *Annals of Operations Research*, 175(1):107–158, 2010.

- FRK Chung and FK Hwang. Do stronger players win more knockout tournaments? *Journal of the American Statistical Association*, 73(363):593–596, 1978.
- Michael Suk-Young Chwe. *Rational ritual: Culture, coordination, and common knowledge*. Princeton University Press, 2013.
- Russell W Cooper, Douglas V DeJong, Robert Forsythe, and Thomas W Ross. Selection criteria in coordination games: Some experimental results. *The American Economic Review*, 80(1):218–233, 1990.
- Camille Cornand. Speculative attacks and informational structure: an experimental study. *Review of International Economics*, 14(5):797–817, 2006.
- Giancarlo Corsetti, Amil Dasgupta, Stephen Morris, and Hyun Song Shin. Does one Soros make a difference? a theory of currency crises with large and small traders. *Review of economic Studies*, pages 87–113, 2004.
- Miguel Costa-Gomes, Vincent P Crawford, and Bruno Broseta. Cognition and behavior in normal-form games: An experimental study. *Econometrica*, pages 1193–1235, 2001.
- Vincent P Crawford and Hans Haller. Learning how to cooperate: Optimal play in repeated coordination games. *Econometrica: Journal of the Econometric Society*, pages 571–595, 1990.
- Vincent P Crawford, Uri Gneezy, and Yuval Rottenstreich. The power of focal points is limited: Even minute payoff asymmetry may yield large coordination failures. *The American Economic Review*, pages 1443–1458, 2008.
- Vincent P Crawford, Miguel A Costa-Gomes, and Nagore Iriberrí. Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications. *Journal of Economic Literature*, 51(1):5–62, 2013.
- Roger R Davidson and Daniel L Solomon. A bayesian approach to paired comparison experimentation. *Biometrika*, 60(3):477–487, 1973.
- John S Decani. Maximum likelihood paired comparison ranking by linear programming. *Biometrika*, 56(3):537–545, 1969.
- John S Decani. A branch and bound algorithm for maximum likelihood paired comparison ranking. *Biometrika*, 59(1):131–135, 1972.
- John Duffy and Jack Ochs. Equilibrium selection in static and dynamic entry games. *Games and Economic Behavior*, 76(1):97–116, 2012.

- Arpad E Elo. *The rating of chessplayers, past and present*, volume 3. Batsford London, 1978.
- Urs Fischbacher. z-tree: Zurich toolbox for ready-made economic experiments. *Experimental economics*, 10(2):171–178, 2007.
- Sidney D Gamble. The disappearance of foot-binding in tingsien. *American Journal of Sociology*, pages 181–183, 1943.
- George Georgiadis. Projects and team dynamics. *The Review of Economic Studies*, page rdu031, 2014.
- David Gill and Victoria L Prowse. Cognitive ability, character skills, and learning to play equilibrium: A level-k analysis. *Journal of Political Economy*, 2016.
- Itay Goldstein and Ady Pauzner. Demand–deposit contracts and the probability of bank runs. *the Journal of Finance*, 60(3):1293–1327, 2005.
- John C Harsanyi and Reinhard Selten. A general theory of equilibrium selection in games. *MIT Press Books*, 1, 1988.
- Harvard Business School Press. *Managing Teams: Forming a Team that Makes a Difference*. 2004.
- Frank Heinemann, Rosemarie Nagel, and Peter Ockenfels. The theory of global games on test: experimental analysis of coordination games with public and private information. *Econometrica*, 72(5):1583–1599, 2004.
- Christian Hellwig. Imperfect common knowledge of preferences in global coordination games. *UCLA Economics Online Papers*, 211, 2002.
- Teck-Hua Ho, Colin Camerer, and Keith Weigelt. Iterated dominance and iterated best response in experimental "p-beauty contests". *American Economic Review*, pages 947–969, 1998.
- Bengt Holmstrom. Moral hazard in teams. *The Bell Journal of Economics*, pages 324–340, 1982.
- Tom Hutchinson. Introduction to project work, 2001. URL http://www.oupskyline.com/downloads/project_work.pdf. http://www.oupskyline.com/downloads/project_work.pdf.
- Andrea Isoni, Anders Poulsen, Robert Sugden, and Kei Tsutsui. Focal points in tacit bargaining problems: Experimental evidence. *European Economic Review*, 59:167–188, 2013.

- Nicholas P Jewell and John D Kalbfleisch. Maximum likelihood estimation of ordered multinomial parameters. *Biostatistics*, 5(2):291–306, 2004.
- Morton I Kamien and Nancy L Schwartz. *Dynamic optimization: the calculus of variations and optimal control in economics and management*. Courier Corporation, 2012.
- Margaret E Keck and Kathryn Sikkink. *Activists beyond borders: Advocacy networks in international politics*, volume 6. Cambridge Univ Press, 1998.
- Godfrey Keller, Sven Rady, and Martin Cripps. Strategic experimentation with exponential bandits. *Econometrica*, 73(1):39–68, 2005.
- Urme Khan and Maxwell B Stinchcombe. The virtues of hesitation: Optimal timing in a non-stationary world. *The American Economic Review*, 105(3):1147–1176, 2015.
- Nicolas Klein. Strategic learning in teams. *Games and Economic Behavior*, 82:636–657, 2013.
- Nicolas Klein and Sven Rady. Negatively correlated bandits. *The Review of Economic Studies*, 2011.
- Katrin B Klingsieck. Procrastination. *European Psychologist*, 2015.
- Terri Kneeland. Identifying higher-order rationality. *Econometrica*, 83(5):2065–2079, 2015.
- Terri Kneeland. Coordination under limited depth of reasoning. *Games and Economic Behavior*, 2016.
- Philipp Külpmann. Procrastination and projects. *Center for Mathematical Economics Working Paper*, (544), 2015.
- Christoph Kuzmics, Thomas Palfrey, and Brian W Rogers. Symmetric play in repeated allocation games. *Journal of Economic Theory*, 154:25–67, 2014.
- Manuel Laguna, Rafael Marti, and Vicente Campos. Intensification and diversification with elite tabu search solutions for the linear ordering problem. *Computers & Operations Research*, 26(12):1217–1230, 1999.
- David Laibson. Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, pages 443–477, 1997.
- Eugene L Lawler and David E Wood. Branch-and-bound methods: A survey. *Operations research*, 14(4):699–719, 1966.

- Patrick Legros and Steven A Matthews. Efficient and nearly-efficient partnerships. *The Review of Economic Studies*, 60(3):599–611, 1993.
- Johan Lim, Xinlei Wang, and Wanseok Choi. Maximum likelihood estimation of ordered multinomial probabilities by geometric programming. *Computational Statistics & Data Analysis*, 53(4):889–893, 2009.
- Ben Lockwood and Jonathan P Thomas. Gradualism and irreversibility. *The Review of Economic Studies*, 69(2):339–356, 2002.
- Robert Duncan Luce and Howard Raiffa. *Games and decisions: Introduction and critical surveys*. Wiley New York, NY., 1957.
- Ching-To Ma, John Moore, and Stephen Turnbull. Stopping agents from “cheating”. *Journal of Economic Theory*, 46(2):355–372, 1988.
- Gerry Mackie. Ending footbinding and infibulation: A convention account. *American sociological review*, pages 999–1017, 1996.
- Gerry Mackie and John LeJeune. Social dynamics of abandonment of harmful practices: a new look at the theory. *Special Series on Social Norms and Harmful Practices, Innocenti Working Paper*, 6, 2009.
- Rafael Marti and Gerhard Reinelt. The linear ordering polytope. In *The Linear Ordering Problem*, pages 117–143. Springer Berlin Heidelberg, 2011.
- Judith Mehta, Chris Starmer, and Robert Sugden. The nature of salience: An experimental investigation of pure coordination games. *The American Economic Review*, 84(3):658–673, 1994.
- Dov Monderer and Dov Samet. Approximating common knowledge with common beliefs. *Games and Economic Behavior*, 1(2):170–190, 1989.
- Stephen Morris and Hyun Song Shin. Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review*, pages 587–597, 1998.
- Stephen Morris and Hyun Song Shin. Coordination risk and the price of debt. *European Economic Review*, 48(1):133–153, 2004.
- H William Morrison. Testable conditions for triads of paired comparison choices. *Psychometrika*, 28(4):369–390, 1963.
- Rosemarie Nagel. Unraveling in guessing games: An experimental study. *The American Economic Review*, pages 1313–1326, 1995.

- John Nash. Non-cooperative games. *Annals of mathematics*, pages 286–295, 1951.
- Ted O’Donoghue and Matthew Rabin. Choice and procrastination. *Quarterly Journal of Economics*, pages 121–160, 2001.
- Martin T Orne. On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications. *American psychologist*, 17(11):776, 1962.
- LS Pontryagin, VG Boltyanskii, and RV Gamkrelidze. The mathematical theory of optimal processes. 1962.
- Eugenio Proto, Aldo Rustichini, and Andis Sofianos. Higher intelligence groups have higher cooperation rates in the repeated prisoner’s dilemma. 2014.
- Ariel Rubinstein. The electronic mail game: Strategic behavior under "almost common knowledge". *The American Economic Review*, pages 385–391, 1989.
- Thomas C Schelling. The strategy of conflict. *Cambridge, Mass*, 1960.
- Jagbir Singh and WA Thompson. A treatment of ties in paired comparisons. *The Annals of Mathematical Statistics*, 39(6):2002–2015, 1968.
- Patrick Slater. Inconsistencies in a schedule of paired comparisons. *Biometrika*, 48 (3/4):303–312, 1961.
- Dale O Stahl and Paul W Wilson. On players’ models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 10(1):218–254, 1995.
- Tomasz Strzalecki. Depth of reasoning and higher order beliefs. *Journal of Economic Behavior & Organization*, 108:108–122, 2014.
- Robert Sugden. A theory of focal points. *The Economic Journal*, pages 533–550, 1995.
- Andrew Taylor. It projects: sink or swim. *The computer bulletin*, 42(1):24–26, 2000.
- WA Thompson and Russell Remage. Rankings from paired comparisons. *The Annals of Mathematical Statistics*, 35(2):739–747, 1964.
- Johannes Tiwisina and Philipp K ulpmann. Probabilistic transitivity in sports. *Center for Mathematical Economics Working Paper*, (520), 2014.
- Tostan. Female genital cutting. <http://www.tostan.org/female-genital-cutting>, 2016. Accessed: 2016-05-21.
- Amos Tversky. Intransitivity of preferences. *Psychological review*, 76(1):31, 1969.

- UNICEF. Female genital mutilation/cutting: a global concern. http://www.unicef.org/media/files/FGMC_2016_brochure_final_UNICEF_SPREAD.pdf, 2016.
- C Van Eeden. Estimation in restricted parameter spaces-some history and some recent developments. *CWI Quarterly*, 9(1):69–76, 1996.
- John B Van Huyck, Raymond C Battalio, and Richard O Beil. Tacit coordination games, strategic uncertainty, and coordination failure. *The American Economic Review*, pages 234–248, 1990.
- Andreas Wächter and Lorenz T Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical programming*, 106(1):25–57, 2006.
- Jonathan Weinstein and Muhamet Yildiz. Impact of higher-order uncertainty. *Games and Economic Behavior*, 60(1):200–212, 2007a.
- Jonathan Weinstein and Muhamet Yildiz. A structure theorem for rationalizability with application to robust predictions of refinements. *Econometrica*, 75(2):365–400, 2007b.
- Eyal Winter. Incentives and discrimination. *American Economic Review*, pages 764–773, 2004.
- Christopher A Wolters. Understanding procrastination from a self-regulated learning perspective. *Journal of Educational Psychology*, 95(1):179, 2003.
- World Bank Group. *World Development Report 2012: Gender Equality and Development*. World Bank Publications, 2012.
- Stefan Wuchty, Benjamin F Jones, and Brian Uzzi. The increasing dominance of teams in production of knowledge. *Science*, 316(5827):1036–1039, 2007.