Strategic Thinking: Experimental Investigation and Economic Theory

Stefan Peter Penczynski

Dissertation submitted for the degree of

Doctor of Philosophy in Economics at

The London School of Economics and Political Science

May 2009

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others. The first chapter draws on work that was carried out jointly with equal share by Konrad B. Burchardi and me.

I warrant that this authorization does not, to the best of my belief, infringe the rights of any third party.

Stefan Peter Penczynski

Abstract

Strategic interaction has traditionally been modelled in economics with game theoretic equilibrium models. In these models, strategies constitute best responses to beliefs that are consistent with other players' strategies. While this consistency is realistic in settings familiar to the players, it is less appropriate in situations that are encountered for the first time. This shortcoming has led to the conception of models of bounded rationality, in particular the level-k model of levels of reasoning.

While experimental studies usually employ only action data to test the level-k model, in this thesis, a team setup with electronic communication between participants allows for a qualitatively richer insight in actual reasoning processes.

Two different games are played to investigate different notions of strategic thinking. The first study uses a dominance-solvable 'beauty contest' game in which 6-8 teams compete for a prize. This game lends itself naturally to the observation of levels of reasoning. In addition, the communication allows to analyse the anchoring level-0 belief and the population belief of individual players.

The second study uses a zero-sum 'hide and seek' game that two teams play against each other. Both the influence of non-neutral framing on the level-0 belief and the task-dependence of the level of reasoning can be brought to light in this study.

The third and final chapter considers an application of the equilibrium concept in the theory of implicit incentives, a situation of complex strategic interaction. The method and results of the study are viewed against the background of the limitations of equilibrium models to reflect a situation of inherent one-shot nature.

Für meine Eltern, Hiltrud und Manfred Penczynski

Acknowledgements

First of all, I am greatly indebted to my advisors, Gilat Levy and Georg Weizsäcker, for the support and guidance that they provided me with.

I like to thank my coauthor, Konrad B. Burchardi, for his great support and many intense discussions.

Many thanks are due to Madhav S. Aney, Eric Eyster, Leonardo Felli, Nagore Iriberri, Elena Manzoni, Guy Mayraz, Rosemarie Nagel and Michele Piccione for plenty of helpful comments and discussions.

I also thank all my friends and fellow researchers at the LSE, and STICERD/EOPP in particular, for discussions and comments on my work, and for creating such a stimulating research environment.

This work benefitted a lot from my stay at Universitat Pompeu Fabra, Barcelona, Spain, between April and August 2007. The exchange was organised within the European Doctoral Program in Quantitative Economics.

Financial support from the LSE, the Michio Morishima fund (STICERD), the German Association of Experimental Economists (GfeW) and the German Academic Exchange Service (DAAD) is gratefully acknowledged.

I owe a lot to my friends for their support over the past few years. I am also grateful for my family's support and understanding. I have not been able to spend as much time with them as I would have liked.

Contents

List of Figures

List of Tables

Preface

Many strategic decisions are taken by economic agents after only a short history of similar situations or without such history at all. In trying to understand and predict behavior in these situations, equilibrium concepts used in economics often turn out to be of limited use. In equilibrium models the players find their strategy by optimally responding to their belief about the strategies of the other players – and these beliefs are thought to be consistent with the other players' strategies. This is well suited to characterise behaviour in settings familiar to the player. However, in strategic situations with a short or no history of play there is little reason to think that the players' beliefs and the opponents' strategies ought to be consistent.

Alternative models of bounded rationality have been proposed for economic applications that are closer to one-shot games than repeated games. The model of levels of reasoning features prominently in this area, because it precisely relaxes the assumption of mutually consistent beliefs, but retains the idea of strategic behaviour with optimal responses. Applications include the 'beauty contest' game (Nagel, 1995), normal-form games (Stahl and Wilson, 1995; Costa-Gomes, Crawford and Broseta, 2001), two-person zero-sum games with non-neutral framing (Crawford and Iriberri, 2007a), common-value auctions (Crawford and Iriberri, 2007b), and coordination games (Crawford, Gneezy and Rottenstreich, 2008). Related structures have been introduced in the area of social learning (Eyster and Rabin, 2008).

The main idea of the level-k model is to impose a certain structure on the players' beliefs about the other players' strategies, and to allow for heterogeneity among the players in this respect: Some players, so-called 'level-0 reasoners', do not best respond to any belief. This is modeled as playing randomly according to some distribution ('level-0 distribution') over the action space. Other players, 'level-1 reasoners', recognise this behaviour, form the belief that all other players are level-0-reasoners and best-respond to this belief. A third type of players, 'level-2 reasoners', believe to be playing against both level-0 and level-1 reasoners, form a belief about their proportions in the population ('population belief') and best respond to this belief. Similarly, 'level-3 reasoners' believe to be playing against level-0, level-1 and level-2 reasoners and so on. Every agent believes that others are on a strictly lower level. Hence, the population belief for level-0 and level-1 is trivial. For players of level-2 and higher, different forms of population beliefs can be conceived. A degenerate population belief, as for example assumed in the model by Nagel (1995), is one where all other players are believed to be exactly one level of reasoning below oneself.

Various versions of the level-k model have been introduced, with differing assumptions about the population beliefs, the precision of best responses, the sophistication of players and the presence of types with dominance reasoning. Generally, they were shown to explain reasonably well the action data from experimental one-shot games. However, this explanatory power of the level-k model is not surprising: Once the consistency of beliefs assumption is relaxed, the model features more degrees of freedom. The action is thought to be the outcome of a more or less involved reasoning process, which depends on the player's own level of reasoning, the belief about how far other players reason and the belief about the level-0 play. In fact, the model is able to explain *any* experimental action data if one just chooses appropriately the level-0 action distribution, the level-0 belief and/or level-k population distribution. The flip-side of this are two major short-comings: Firstly, any test of the model which is based only on action data will have little or no power. Secondly, and most importantly, unless all the model's parameters are pinned down, it has little predictive value – which is what one is ultimately interested in.

For a variety of applications, different level-0 beliefs have been proposed in the literature to incorporate specific factors of the situation under consideration. The initial models of Nagel (1995) and Stahl and Wilson (1995) did not impose particular distributions but assumed a uniform random level-0 belief. Bacharach and Stahl (2000) propose a level-0 belief that captures the salience of options which is not reflected in the payoffs, but only in the labelling. Crawford (2003) considers a certain attraction of level-0 players to truthful statements in strategic communication. So far, a clear definition as to what kind of considerations can enter the level-0 has not yet developed. To which extent the proposed level-0 beliefs reflect the naturally occurring level-0 beliefs is the subject of the first two chapters of this thesis.

Since data about choices made is on its own not sufficient to estimate the parameters of a general level-k model, further efforts have been made to test the theory's assumptions by considering information search patterns and relating them to types of players (Costa-Gomes et al., 2001; Costa-Gomes and Crawford, 2006) or by eliciting beliefs (e. g. Costa-Gomes and Weizsäcker, 2008).

The first chapter, which draws on joint work with Konrad B. Burchardi, proposes an experimental design that allows to closely observe the reasoning process in situations of one-shot games by making use of intra-team communication. In particular, the games are played in teams of two players with one simultaneously exchanged message and a suggested decision. After reading the team partner's message both players make their final decision independently, one of which is then chosen randomly to be the team's action. Importantly, this design provides strong incentives for the players to communicate their reasoning fully because the message is the only opportunity to convince the partner of one's reasoning. Furthermore, due to the simultaneity in the exchange of the message, any influence that the team partners might have on each others' reasoning can only come into effect once the own message is sent. This set-up is hence suitable to get a written account of the individual's reasoning.

The first chapter presents the results from the 'beauty contest' game. This game is particularly suited to investigate the shape of the level distribution, the nature of the level-0 belief and the prevailing population belief. The second chapter presents results from the 'hide and seek' game. This game allows to investigate a non-random level-0 belief as assumed to result from the non-neutral framing of the options. Furthermore, the population is divided into hiders and seekers according to the players' tasks. It can be analysed to which extent the task of the player influences the level of reasoning and the level-0 belief.

The third chapter considers a theoretical model of career concerns. Career concerns have been studied by economists since Holmström (1999, first published 1982) formalised the idea of Fama (1980) according to which concerns about reputation in the labour market suffice to mitigate moral hazard problems. Holmström shows that costly effort is exerted because it increases the agent's expectation of future wages by increasing the market's belief of the worker's productivity.

In this chapter I explore implications of heterogeneity among the agents in addition to varying talent. The agents differ in the magnitude of career concerns, which leads them to exert different levels of effort in a given situation. If principals are uncertain about this aspect of the agent's type, I find that depending on the circumstances, this additional uncertainty might decrease total effort exertion over time.

1 Estimating the level-k model: the 'beauty contest' game[1]

1.1 Introduction

The contribution of this chapter is twofold: (a) it presents an experimental design which allows to obtain an incentivised written account of individual reasoning in one-shot games and (b) it applies this design to the 'beauty contest' game in order to estimate the structural parameters of the most prominent non-equilibrium model of individual reasoning in one-shot situations, the level-k model.

We will show that the written accounts are rich in the sense that they convey information about the parameters of the level-k model which were previously hard to bring to light. In particular, we apply our design to the classic 'beauty contest' game. In this game all players choose an integer between 0 and 100. The winner is the one whose choice is closest to p times the average of all numbers, $p < 1$. The unique Nash equilibrium is 0, however, in experiments most choices are between 20 and 50. It is the contribution of Nagel (1995) to point out how this can be explained by a level-k model of reasoning.

We classify the transcribed communication from the first period of the game along the lines of a general level-k model that nests previous contributions. Following this procedure, the vast majority of players was classified to follow a level-k type of reasoning. The level of reasoning was unambiguously determined for 50 of the 84 subjects in our experiment. Of those subjects 17 were identified as level-0, 26 as level-1, 6 as level-2, and 1 as level-3 reasoner. Importantly this shows that the level-0 play is not only a thought experiment of higher level players but does actually exist. Further 20 subjects were identified as level-k reasoners, but only an interval of likely levels could be classified. For 36 subjects we could identify level-0 beliefs from the communication. For almost all players of at least level-2 the message allows to identify their population belief, which is degenerate on level $k-1$ in 75% of the identified cases. Only two players identified the Nash equilibrium and only two players were classified to apply iterated deletion of dominated strategies.

We then use the information on the individuals' levels to estimate econometrically the structural parameters of the level-k model. In particular we show how this information can be used as constraints in a maximum-likelihood estimation of a generalised version of the level-k model. For reasons outlined in section 1.6, a maximum likelihood estimation (MLE) of the level-k model is not instructive under the standard assumptions in the literature. The model is generalised in the sense that rather than assuming that level-0 beliefs take the same value across players and that the level-0

[1]This chapter draws on work that was carried out jointly with equal share by Konrad B. Burchardi and me.

play is uniformly distributed, we model both level-0 belief means and level-0 play as draws from appropriate distributions which are estimated jointly with the level-k distribution. In this way we get estimates for all relevant parameters of the model. The estimated level-0 beliefs and actions are consistent with the information obtained in the communication classification. We use Monte Carlo studies to show that the estimator used is unbiased.

The estimates from this exercise shed light on which versions of the level-k model best reflect the individuals' reasoning process. The observed levels of reasoning are surprisingly low. Our maximum likelihood estimate for the average level of reasoning is only 0.82. The distribution of the actions of level-0 players is estimated to be centered around 60 and does not appear to be uniform. The belief of higher-level players about the mean of the actions of level-0 players is centered at 53 and not consistent with the actions. These estimates discriminate between prevailing versions of the level-k model, shed doubts on assumptions made by all versions of the model and lead us to different conclusions than we would have drawn had we applied estimation methods used elsewhere.

The remainder of the chapter is organised as follows: Section 1.2 introduces prominent versions of the level-k model. Section 1.3 relates our research methodology to other approaches in the literature. The description of the experimental design and the classification procedure is detailed in section 1.4. The main results are presented subsequently in section 1.5. The MLE estimation is presented in section 1.6 and followed by the conclusion.

1.2 The Level-k Model and its Variants

The seminal paper by Nagel (1995) introduced the level-k model as outlined in the preface in order to explain the choice distribution in 'beauty contest' games. Subsequently, numerous studies have analysed experiments involving variants of the beauty contest game.[2] Object of the variations are monetary incentives, equilibrium actions, subject populations, information environments and many more.

A recent variant of the level-k model is the model of cognitive hierarchy (CH) introduced by Camerer, Ho and Chong (2004). This model links various aspects of the level-k concept using only one parameter. The model by Stahl and Wilson (1995) considers probabilistic best responses in normal-form games, relaxing the full optimisation assumption. The setup by Costa-Gomes et al. (2001) and Costa-Gomes and Crawford (2006) allows for and can identify dominance reasoners as well as sophisticated types, who know of the equilibrium but play a non-equilibrium best response.

Subsequently the shared characteristics of those non-equilibrium models of strategic reasoning in the tradition of Nagel (1995) will be laid out and formalised. We will point out where the most prominent versions differ from each other.

[2] For an overview see the survey by Nagel (1999) and the paper by Camerer, Ho and Chong (2004).

Let there be players $i = 1, ..., N$ and let their action be denoted as $a_i \in A$. In our case, the action space A consists of all integers in the interval $[0, 100]$. Let the level of reasoning be $k = 0, 1, 2, 3, ...$ Let the belief of a level-k player with $k \geq 1$ about the proportion of level-i players in the population be denoted as $b_k(i)$, $i \leq k-1$.

Level-0 players are defined as not best-responding to any belief. So while they might or might not have some belief about what others do, they do not take a belief into account when deciding on their strategy. Rather they are assumed to play non-strategically which is modeled as random draws from a distribution $r^0(a)$ defined over A.

Players of higher levels $k \geq 1$ form beliefs about the mean of the level-0 distribution $\tilde{a}$. The distribution of level-0 mean beliefs is $\tilde{r}(\tilde{a})$. Given his level-0 mean belief a level k, $k \geq 1$, player will figure out how lower level players play and given his belief on their distribution, $b_k(i)$, he will find his best response. Note, that to be able to calculate the optimal strategy of some lower level player, a player with $k \geq 3$ needs to form a belief about the beliefs of player i, $2 \leq i \leq k-1$. In particular, a level-3 player needs to know what a level-2 player believes the relative proportions of level-0 and level-1 players are. These higher-order beliefs are in all versions of the level-k model assumed to be consistent with the beliefs of a player of level i. Similarly a level-k player, $k \geq 2$, needs to form a belief over the lower players' beliefs of the mean of the level-0 distribution. This is generally assumed to be the same as a level-k player's own level-0 belief.

Given $b_k(\cdot)$ and $\tilde{r}(\tilde{a})$, the players are assumed to 'best respond', meaning that their strategy is found as s_k

$$s_k(b_k(\cdot), \tilde{r}) = \underset{a \in A}{\operatorname{argmax}}\ u(a; b_k(\cdot), \tilde{r}). \tag{1.1}$$

Lastly, the true distribution of level-k types is denoted as $l(k)$.

All models assume the level-0 distribution $r^0(a)$ to be a uniform distribution over the action space. This assumption is often justified by Bernoulli's principle of insufficient reason. As well they share the assumption of rational expectations for the level-0 mean beliefs, implying a belief $\tilde{r}(\tilde{a})$ degenerate at 50.

Under an assumption on the anchoring level-0 distribution, the level-k distribution and a specification for $b_k(\cdot)$ such a model is able to make a probabilistic prediction about the frequencies of strategies.[3]

1.2.1 The Model of Nagel (1995)

The original model by Nagel (1995) assumes the players' beliefs for $k > 0$ to be

$$b_k(i) = \begin{cases} 1 & \text{if } i = k-1 \\ 0 & \text{otherwise.} \end{cases}$$

[3]The prediction is probabilistic since the level-0 play is probabilistic.

Given the common assumptions of a uniform $r^0(a)$ and rational expectations with respect to the level-0 play, a level-1 type would hence believe that everybody else is playing 50 on average. If $p = \frac{2}{3}$, his optimal strategy given this belief is to play 33. A level-2 type would, given the above $b_k(i)$ specification, believe that everybody else is level-1, hence playing 22. A level-3 type again believes that everybody else is a level-2 type and so on. Together with any assumption on the distribution of types $l(k)$ this gives a probabilistic prediction on the observed actions in a p-'beauty contest' and other games.

1.2.2 The Model of Stahl and Wilson (1995)

Stahl and Wilson (1995) consider a version where the best response is calculated with error. The level-0 distribution, $r^0(a)$, is assumed to be uniform across the action space again, which they think of as a fully imprecise best-response. Levels 1 and 2 best respond with error, playing actions with a higher expected payoff with a higher probability, such that the optimisation principle behind equation 1.1 is relaxed. Higher levels than $k = 2$ are not considered. For the population distribution belief of level 2 players, they allow for any combination of level-0 and level-1 players, not putting any restriction as opposed to all other models considered.

Apart from level reasoners, Stahl and Wilson (1995) consider some other reasoning types: 'Naïve Nash' types that play the Nash equilibrium; 'Worldly' types that best-respond to a population of level-k and 'naïve Nash'-types; and 'rational expectation' types that best-respond to a correct belief over the frequencies of all other types.[4]

1.2.3 The Model of Camerer, Ho and Chong (2004)

The cognitive hierarchy model of Camerer et al. (2004) differs from the previous models in its assumptions about the belief formation. Rather than assuming that players believe all others to reason exactly one level lower than themselves, it is assumed that players believe the others to be a mixture of lower-level types. Further they assume that the players' beliefs reflect the true relative frequencies of lower level types and that the true distribution of types $l(k)$ follows a Poisson distribution with parameter τ. This elegantly links the population distribution, the population beliefs and higher-order population beliefs[5] with only one parameter. Formally, for all $k > 0$,

$$b_k(i) = \frac{l(i;\tau)}{\sum_{m=0}^{k-1} l(m;\tau)},$$

where

$$l(k;\tau) = \frac{\tau^k e^{-\tau}}{k!}. \tag{1.2}$$

[4]Note that all of those types, except the 'rational expectations' type, can be thought of as being nested in the level-k framework.

[5]These are modeled to be consistent with the actual population beliefs of lower-level types.

A level-1 type still believes everybody else to be level-0 reasoners and plays 33. However, a level-2 type now believes both level-0 and level-1 reasoners to exist, their exact relation depending on the true population distribution, i. e. on the parameter τ. Therefore, a level-2 type chooses an action higher than 22.

1.2.4 The Model of Costa-Gomes, Crawford and Broseta (2001)

Costa-Gomes et al. (2001) model the level-k types in the same fashion as Nagel (1995).

An innovation of Costa-Gomes et al. (2001) and Costa-Gomes and Crawford (2006) is to allow for a type that reasons along the lines of dominance. The original concept of dominance is not leading towards a point prediction of play. They propose a closure of the dominance reasoning by assuming a best response to a uniform distribution of actions that are remaining after a number of rounds of deletion of dominated strategies. In their setup, they can empirically identify this form of dominance reasoning.[6] Further they consider 'equilibrium' types that play the Nash equilibrium strategy, similar to Stahl and Wilson's 'naïve Nash'-types, and they call 'sophisticated' types those players that best respond to the actual distribution of others' responses, similar to Stahl and Wilson's 'worldly' and 'rational expectation' types.

The general model outlined at the beginning of this section nests all these possible specifications of the level-k model. Note that 'equilibrium' and 'sophisticated' types are as well nested in the level-k models as individuals with a high level of reasoning and appropriate population beliefs. The only exception are 'dominance' reasoners, which were identified separately in the classification procedure. Otherwise we classified the messages along the lines of this general model (see section 1.4.3 on the classification procedure) and section 1.6 will present the equivalent econometric model. Estimating this allows us to discriminate between the different sets of assumptions.

1.3 Literature: Estimating and Testing the Level-k Model

Most of the early studies which 'test' the level-k model make use of choice data and conclude that the model does well in explaining the observations. The major problem in the falsification of the theory is the fact that the level-k model offers many potential explanations for observed actions. In fact, if you just make assumptions on the – so far – unobserved parameters of interest, like the level-0 distribution, the level-k distribution and the system of population beliefs, the model is able to explain any data. Hence, when using choice data only, the single way to estimate one of the model's parameters is by making assumptions on the remaining subset of parameters.

In the early literature, qualitative insights about the underlying reasoning were obtained by asking subjects after the experiment as to how they made their choice in

[6]Whether this particular kind of closure models behaviour correctly is an empirical question. The two 'dominance' types in our data do not best respond to the actions remaining after one round of deletion, they rather play the midpoint of the remaining action interval.

the game. A descriptive analysis of optionally given comments received in the context of laboratory, classroom and field experiments is presented by Bosch-Domènech, Montalvo, Nagel and Satorra (2002). Apart from the absent incentives to state one's thoughts, a drawback of this kind of data is the temporal gap between the actual reasoning process and the report of the reasoning. Between the reasoning and the statement, influences coming from the game's outcome, forgetfulness, further thoughts that occurred later or even direct communication might make the comments unreliable representatives of the true reasoning. Our design is immune to such critique since it allows to get verbal accounts of reasoning processes which are written during the actual thinking period and minimises the potentially biasing influences on the way from thinking to verbalising. Indeed, the standard theme in protocol analysis, the psychological research on methods of eliciting verbal accounts from participants, is that "[t]he closest connection between thinking and verbal reports should be found when participants were instructed to focus on the task while verbalising their ongoing thoughts" (Ericsson, 2002, p. 983).

Progress on this initial approach was made by analysing data on information search patterns as in Costa-Gomes et al. (2001), Costa-Gomes and Crawford (2006) and Costa-Gomes and Crawford (2009). The two recent papers use decisions made in a 16-period two person 'beauty contest' game without feedback and match the players' 16 choices with a typical fingerprint of, say, a level-1 player. In addition, they link search patterns to types. This interesting design allows them to obtain individual level estimates, however it does not allow to investigate non-uniform level-0 beliefs. An advantage of our setting is that there is no need for game repetitions, under which individual reasoning might evolve even without feedback. Furthermore, we are able to identify level-0 players, which is not possible in the multi-period 'fingerprint' analysis, where only non-random patterns with an underlying mechanism can be classified.

Attempts to obtain data other than mere choices are also prominent in the learning literature, where interview questions and attentional data are used to delimit different learning behaviour (Salmon, 2001, 2004). However, interview questions are not incentivised and attentional data can usually only be used in repeated games, in order to allow for players' information search patterns to stabilise.

A further approach to identify the underlying parameters of the model is presented in the econometric study by Bosch-Domènech, Montalvo, Nagel and Satorra (2004). Here, the action data of many different 'beauty contest' games is used for a finite mixture analysis, a collection of superimposed bounded normal distributions to best fit the action data. Some of the estimated normal distributions can be associated with underlying levels of reasoning. While the procedure benefits from the high number of observations, it is limited in its interpretation of the underlying reasoning.[7] In addi-

[7]For example, a very flat distribution centered around 36 is said to be part of the level-0 belief. Our results would suggest that this is not the case. Both because in our study the level-0 seems to be rather above 50 and the optimisation of level-1 players is subject to mistakes, which is the more likely effect caught in this distribution.

tion, the data is aggregated among others from non-laboratory experimental settings which, due to the uncontrolled informational environment, makes it difficult to infer general patterns in individual reasoning.

At the heart of our approach lies the disclosure of reasoning processes thanks to the incentives to reveal the own thoughts in the team communication. The experimental literature has used team setups on various occasions to get better insights into the reasoning process of subjects and to investigate the performance of teams as opposed to individuals. In that respect, our experimental design is related to the innovative study by Cooper and Kagel (2005) in which team players communicate via an instant messenger, allowing the experimenters to observe the speed of learning to play strategically. Video analyses are increasingly used to judge decision making in the Prisoner's dilemma, the power-to-take game, the ultimatum game, etc. (Hennig-Schmidt and Geng, 2005; Geng and Hennig-Schmidt, 2007; Brosig, 2002). Kocher and Sutter (2005) play a repeated 'beauty contest' with teams of 3 players, showing that teams are better in playing strategically over time, though not initially. However, their experiment was not designed to obtain data which can be used to test the parametric assumptions of the level-k model. In general, the main advantage of our design is the possibility to infer about individual's reasoning by having only one, simultaneously exchanged message.

1.4 Observing Individual Reasoning: The Experiment

In this section we describe in detail our experimental design, lay out the experimental procedures and present our subject pool. In order to process the information contained in the written accounts of the individual reasoning in a structured way we classified the messages along the lines of a general level-k model. We explain the details of this procedure in the last part of this section.

1.4.1 Experimental Design

The main difference to the standard 'beauty contest' treatments in the literature was the team play and the communication structure between the team players. Participants were randomly assigned in teams of 2 players. The two members of a team were connected through the chat module of the experimental software.[8] The team setup was essential to the investigation of the individuals' reasoning thanks to the following communication structure reflected in figure 1.1 (steps in brackets): After the explanation of the rules of the game and the indication of being either a hider- or seeker-team, each team member could state a so-called 'suggested decision' and justify this in a message (I). Neither the size of the message nor the time to write it were limited. Once both team members finished entering their message, the suggested decisions and the messages were simultaneously exchanged (II). Since it was only possible to send this

[8]The experiment was programmed and conducted with the software z-Tree (Fischbacher, 2007).

one message, the entire communication was undertaken without any influence from the team partner, therefore reflecting individual reasoning.

In a next step (III), both team members would individually state their 'final decision', knowing both the suggested decision and the message of the team partner. It was known to them that one of the two final decisions would be chosen randomly by the computer to count as the 'team's action' (IV). Facing the 50% chance of having the team partner determine one's action in the game gave the message the importance of being the only way to influence the partner's decision.

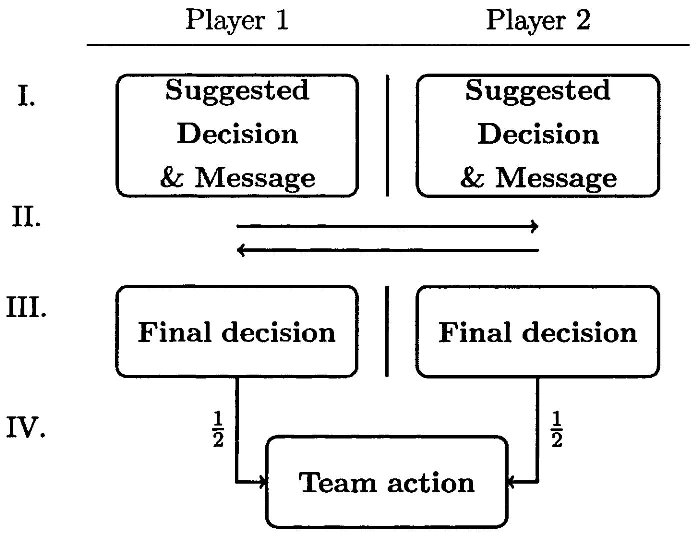

Figure 1.1: Structure of the team communication.

This design ensures that there is a strong incentive to write down the full reasoning underlying the suggested decision in a clear and convincing way. Since this study considers exclusively individual reasoning and no team interaction, the analysis will make use of the suggested decision and the message from step I only.

1.4.2 Experimental Setup

We conducted a 'beauty contest' experiment which was run in 6 sessions in the Experimental Economics Laboratory of the Department of Economics in Royal Holloway (University of London). The 84 participants were mainly undergraduate students in Royal Holloway and all participants were recruited by the host institution. We had invited 18 students per session. Since some students did not show up we allowed the highest possible even number of students into the laboratory on a first-come-first-serve basis. This resulted in 2 sessions with 12, 14 and 16 participants each.

Out of the 84 students 15 were studying Economics, 13 of them being in their first year of studies, one in the second year and one being in the third year. 16 of the 84 students had received some form of training in game theory, but only 5 had been confronted with the beauty contest game before the experiment. The majority of students had participated in an economic experiment before. Table 1.1 summarises some of the participants' background characteristics.

Table 1.1: Players' Backgrounds.

	Economics	*Other*	*NA*	
Field of Study	15	67	2	
	Bachelors	*Masters*	*Ph.D.*	*Other*
Degree	59	13	3	9
	1st	*2nd*	*3rd*	*4th or more*
Year in Degree	54	14	12	4
	Female	*Male*		
Gender	36	48		
	Mean	*St. Dev.*		
Age	21.48	4.38		
	Yes	*No*		
Game Theory Training	16	68		
	Yes	*No*		
Experimental Experience	65	19		
	Yes	*No*		
Beauty Contest Experience	5	79		

Before the start of the 'beauty contest' game the participants were made familiar with the structure of the experiment and the messaging system in two practice rounds. These used the exact same software as above, but asked the teams to find the answer to two unrelated questions. Since it was important to avoid any pre-treatment sensitisation to strategic considerations, we asked them to provide the year of two historic events. The questions in the test round were chosen relatively difficult to stimulate the use of the messaging system.[9]

We played the standard 'beauty contest' game where the winning team was the one closest to $p = \frac{2}{3}$ of the average and integer numbers between 0 and 100 were allowed. We repeated this game 3 times, since in this study we are interested in individual reasoning, we are exclusively concerned with the first round suggested decision and the accompanying first message, i. e. with the actions that took place before any interaction of the team players.[10]

The participants of the experiment were paid a show-up fee of £5 and the winning team won a prize of £20 (£10 per team player).

1.4.3 Classification of communication transcripts

As outlined before, the motivation of our experimental design is to learn from the players' messages about their individual reasoning. For this we had the messages

[9]The questions were: "Which is the year of birth of Nicolaus Copernikus, the first astronomer to formulate a scientific heliocentric cosmology?" (1473) and "In which year died Wolfgang Amadeus Mozart, the famous composer of the Classical Era?" (1791).

[10]The analysis of the learning process is subject of a further project.

classified by two research assistants (RA).

We are interested in 5 pieces of information: (i) How many steps of reasoning does the player apply? (ii) Which level-0 belief does he state, if any? (iii) Which population belief does he state, if any? Further, in order to investigate the prevalence of reasoning patterns different from level-k reasoning we attempt to distinguish level, dominance, equilibrium and sophisticated players, by asking: (iv) Did the player recognize 0 as the unique equilibrium? (v) Did the player apply an iterated elimination of dominated strategies? If the equilibrium is recognised, the player will be classified as 'equilibrium' type if she played 0 and as 'sophisticated' type if she played a best response to a more realistic population belief.

We believe that questions (ii) and (iii) can be, if those beliefs are stated, unambiguously answered. Likewise questions (iv) and (v) can be unambiguously answered with a yes or no. However, we are aware that in some cases it might not be identifiable from the communication how many steps of reasoning were applied *exactly*.[11] Therefore, we ask the classifiers to only indicate a lower and upper bound on the steps of reasoning, if the message lends itself to such inference. These were defined as the lowest and highest level of reasoning which could possibly be interpreted into the messages. The details of the classification instructions can be viewed in appendix A.5 where they are reproduced.

When designing the classification procedure we intended to avoid two potential concerns: Firstly, the classifiers might try to extract more information than the messages actually contain. Our guiding principle in the classification procedure was therefore to take a cautious approach and instruct the classifiers not to note any information if it was not obviously contained in the message. And secondly, we were concerned that in the case of an ambiguous statement relatively low 'suggested decisions' might lead the classifiers to indicate a higher lower-bound on the level of reasoning than was clearly exhibited. We therefore did not reveal the choice data to the classifiers when asking for the lower-bound.[12] Likewise we did not reveal the choice data when asking to indicate the level-0 and population-belief, hence ensuring that these would only be indicated when stated in the message.[13] In contrast, when indicating the upper-bound, knowing about low choices should, if anything, lead the classifiers to indicate a higher upper-bound. Similarly, we revealed the choice data when asking the classifiers to indicate whether the message exhibited any elimination of dominated strategies or equilibrium reasoning, which should, if anything, bias the evidence against the

[11] Think for example of the imaginary statement: "I presume everybody else will play 33, so let us play 22." This clearly exhibits one step of reasoning. But it seems possible, too, that the player skipped the first step of his reasoning when writing down his argument.

[12] Other studies applying a classification do apply a similar procedure in order to avoid any unconscious alignment of the classification with the choice data that might result from implicit assumptions (Rydval, Ortmann and Ostatnicky, 2007, for example).

[13] In some messages the level-0 belief was given as an interval. We asked the classifiers to indicate the mid-point in this case.

prevalence of level-k reasoning.[14]

The classification was undertaken by two Ph.D. students in the Department of Economics at LSE. The instructions were written by the two authors, of which one had taken a look at the communication transcripts beforehand. The instructions were self-contained and were not complemented by verbal comments. Remaining questions were answered via an e-mail list that included all four persons involved and which can be obtained from the authors. As detailed earlier, the classification was split in two parts. The documents for the second part were only given out once the first part was completed. After considering each part individually, the two research assistants met to reconcile their judgements and come up with a joint classification, if possible. We had specified rules how to make cautious use of their individual classifications in case of disagreement.

The two RA's did not have major discrepancies in their judgements and the data we use is exclusively agreed upon by both. This hints towards a robust classification which can be replicated easily. Examples of the classification will be given in section 1.5.

1.5 Results

Before presenting the results obtained from the classification of the players' communication we will first describe their choices in the game played.

1.5.1 Action data

Figure 1.2 shows histograms of the suggested decisions by session from the first period of the beauty contest game. Table 1.2 gives aggregate summary statistics for all 6 session. Figures 1.3 and 1.4 show histograms of the suggested decision and the final decisions aggregated over all sessions.

Table 1.2: Summary statistics.

	Mean	Std. Dev.	Median	Min.	Max.	N
All Sessions						
Suggested Decision	43.93	21.14	40	0	100	84
Final Decision	39.73	18.75	35	0	100	84
Team Action	40.02	18.98	35.5	16	100	84

The data on 'suggested decisions' is comparable to data from the first periods of other experiments with individual participants. However, the data on 'final decisions' is not, since the participants have, at the time of taking this decision, already received a message from their team partner. For the subsequent analysis and classification of the individual reasoning we will use exclusively the 'suggested decision'. The data on

[14] For this part of the classification we provided the classifiers with the message and choice data from the first period. Further we provided all data from later periods, which – according to the classifiers – was not consulted for any participant.

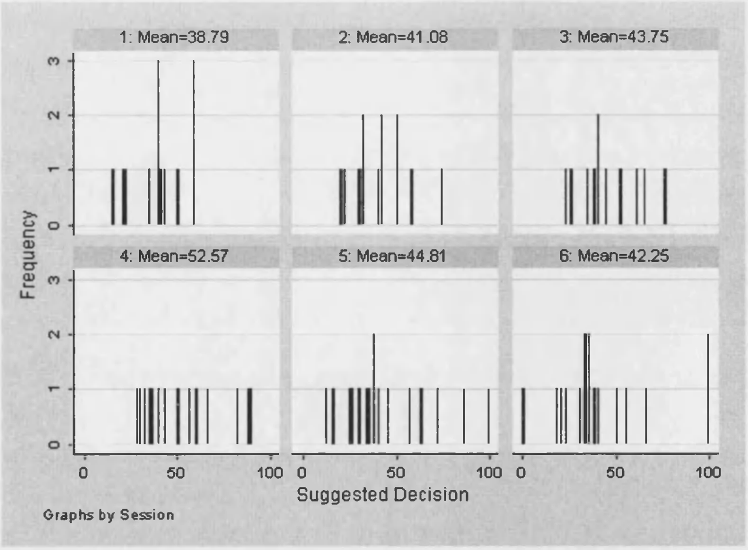

Figure 1.2: Suggested decisions by session.

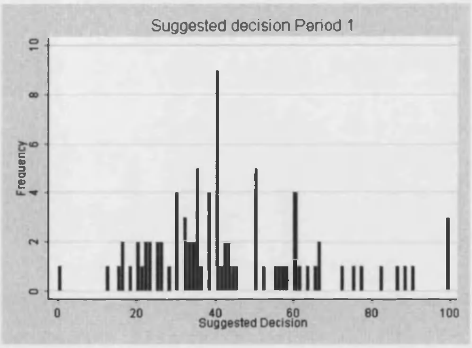

Figure 1.3: Suggested decisions.

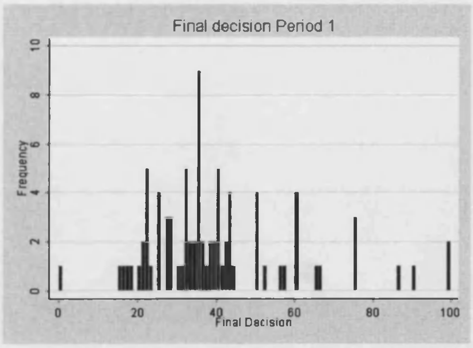

Figure 1.4: Final decisions.

the 'suggested' and 'final decisions' looks similar to data generated in other comparable experiments in having similar means and a high fraction of choices between 22 and 50.[15]

Several other aspects of our choice data are in interesting ways similar to the results of other experiments. The 'final decisions' of the participants are on average lower.[16] This supports the intuition that exchanging messages in a group-decision making process might increase the level of reasoning and hence, in the beauty contest game, decrease the chosen numbers on average. Moreover the variance of decisions across participants decreases after the exchange of the message and outliers become fewer, again supporting the idea that level-0 players, who potentially chose the outliers, become fewer. This effect of group decision making adds to the results of Kocher and Sutter (2005), who find an average of 34.99 when playing with 17 individuals but an

[15]It might be puzzling to see the spike at 40, but since the design allowed to look behind the scenes of reasoning, we can show that some level-0 players simply chose 40 and some level-1 players had a level-0 belief mean of 60. This shows how our design relaxes the dependence on the action data, allowing an analysis of the structural characteristics of reasoning.

[16]The exception to this rule is Treatment 2 where we find a slightly higher average.

average of only 30.78 when playing with 17 groups. However, they allowed for free communication while we restricted the communication to one message only and their groups were composed of 3 individuals each. Both could explain the slightly stronger effect of group play in their study.

Furthermore the literature suggests that for treatments with higher numbers of players the average tends to decrease. Nagel (1995) and Kocher and Sutter (2005) played games with 14-16 and 17 participants, and found an average first round play of 36.6 and 34.99, respectively. Indeed this is below the average of 43.93 observed in our data with groups sizes of 12-16.

We discussed possible influences the design with explanatory messages might have on the reasoning at an earlier point. A possible objection to our procedure is that having to write down their reasoning might increase the participants' level of reasoning, e. g. because the participants would examine the task at hand more thoroughly in order to be able to state sensible arguments in the communication. The fact that our data exhibits if anything higher action choices comforts us in this respect.

Note that the 'team's action' is a random draw of the 2 'final decisions' in each group and hence, not surprisingly, their relation to the 'final decision' shows no systematic pattern.

1.5.2 Communication information

Let us now present the additional data we were able to obtain through the experimental procedure outlined earlier. The following representative examples of communication give an idea of how the classification of the bounds works:

> *"it's quite random, so just guess"*
> (Suggested decision: 58; lower and upper bound: 0)
>
> *"Well. i suppose mopst of the field will choose big numbers so that average will go up to 60-70. 2/3 of this number would be in range of 40-45"*
> (Suggested decision: 42; lower and upper bound: 1; level-0 mean belief: 65)
>
> *"lets say everyone thinks the average will be 50 based on random probability and 2/3rds of that will be about 34 hence they will mostly be choosing around 34 only making the average 34 and its 2/3rd to be around 21"*
> (Suggested Decision: 21; lower and upper bound: 2; level-0 mean belief: 50; degenerate population belief)

Level of Reasoning Table 1.3 presents the lower and upper bounds on the individuals' levels of reasoning. For 70 subjects both a lower and an upper bound was indicated. Eight subjects do not appear in the table due to a non-classified upper bound. Another 6 participants did not make any statement. We classified the bounds on the level of reasoning both when an individual applied a (possibly iterated) elimination of dominated strategies, which leaves him with a set of non-eliminated actions, and a level-k type of reasoning, which derives from a level-0 belief and gives a unique action prediction. However, only 2 participants were found to apply an elimination of dominated strategies.

For 50 of the 84 participants the lower and the upper bounds coincide ($\approx 60\%$) and we could hence fully determine their level of reasoning. As can be seen on the diagonal of the table, these are 17 level-0, 26 level-1, 6 level-2 and 1 level-3 players. For further 20 players we can restrict their level of reasoning to be one of two possibilities. Only for one subject we have an interval between 0 and 2. None of those subjects for who both a lower and upper bound is indicated was classified as potentially reasoning higher than level 3. Note that those players who identified the equilibrium were not given any upper-bound. The bounds on the level of reasoning will be used as constraints in the maximum likelihood estimation of section 1.6. The data by subject is presented in tables 1.8-1.10 in appendix A.3 on pages 39-41.

These results allow us to draw two important conclusions on debates in the literature: Firstly, the big majority of subjects did follow a level-k type of reasoning and only very few follow a deletion of dominated strategies type of reasoning. And secondly, level-0 players are not only a thought experiment of higher-level players but do actually exist.

Table 1.3: Classification results.

		Level Upper Bounds				
		0	1	2	3	Total
	0	17	11	1	0	29
Level Lower Bounds	1		26	3	0	29
	2			6	5	11
	3				1	1
	Total	17	37	10	6	70

Level-0 Belief The results of the classification of the belief about the mean of the level-0 distribution are depicted in the histogram in figure 1.5. For 36 subjects the mean of the level-0 belief could be identified. While the large majority of 21 participants mentioned a belief of exactly 50 or an interval centered around 50, 13 others had a higher level-0 belief, between 55 and 100, in their mind. Two players indicated a lower belief of 40 and 45, respectively. Two players expected other players to play 'something above 50' and were not classified.

Level-0 Actions Figure 1.6 shows the suggested decisions of players who were classified to be level-0 reasoners. In order to check whether the beliefs are consistent

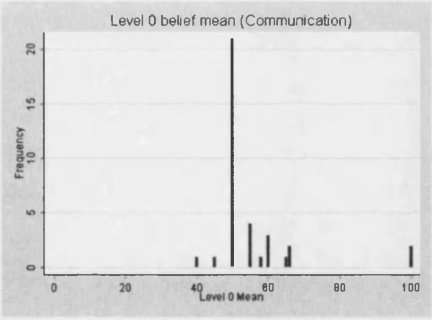

Figure 1.5: Level-0 beliefs from the communication.

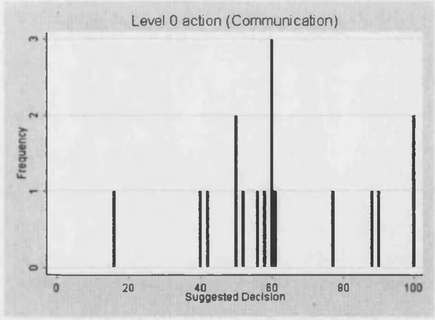

Figure 1.6: Suggested decisions of players classified to be level-0.

with these level-0 actions, table 1.4 shows the summary statistics of the suggested decisions of participants that were classified as level-0 and of the level-0 beliefs. It can be seen that both of them are substantially above 50. When looking at the standard deviation and the extrema, the distribution of level-0 players' suggested decisions is less concentrated, as can be seen from the histogram in figure 1.6. We will come back to this interesting feature in the estimation section.

Table 1.4: Level-0 action and level-0 beliefs.

	Mean	Std. Dev.	Median	Min.	Max.	N
Level-0 suggested decision	62.35	22.39	60	16	100	17
Level-0 belief	55.26	12.33	50	40	100	36

The above shows the suggested decisions of players classified to be level-0 and the level-0 beliefs of higher level players.

The information about the individuals' level-0 mean beliefs will allow us to analyse whether the best responses are best thought of as probabilistic as proposed by Stahl and Wilson (1995). Appendix A.2 will show that there are frequent deviations between the suggested decision and the optimal best response given the belief and the level estimate.

Population Belief The final relevant parameter to characterise the players' level reasoning is the population belief. Among the 12 players with a lower bound of at least 2, 9 were having a degenerate population belief, expecting everybody else to be of one level below. 3 participants were classified as having a non-degenerate population belief. It is unclear whether these 3 are modeled appropriately by the non-degenerate population belief as postulated by Camerer et al. (2004).[17] In any case, all three

[17] The messages sent were as follows:

> *"I reckon the average will be players saying 66/100 approx due to it being two thirds. I believe we should say two thirds of that number which would range between 40 and 50 The lower range will probably be a wiser move as a few others will estimate lower. Personally i would be happy with anything between 38-42."*

> *"provided that a number of people will go for answer of 2/3s of a 100 (66) then we should go for an answer that is 2/3's of that (44), however, i suggest we scale the number down*

statements were given an upper bound of 3. In particular none of them was clearly classified as level-2, the lowest level of reasoning with higher-order beliefs. So in our arguably small sample the rather complex non-degenerate higher-order population-beliefs correlate positively with higher levels of reasoning.

Equilibrium and Dominance Reasoning Lastly, the analysis of the equilibrium and dominance indicators shows that only two players identified the Nash equilibrium, one of them being an 'equilibrium'-type that suggests playing 0, another one being 'sophisticated' in the sense that she imitated a level-2 player and suggested 20. Two other players were classified to apply iterated deletion of dominant strategies. Notably, the one round of deletion of dominated strategies was not followed by a best response. In that sense the assumption of Costa-Gomes et al. (2001) that the deletion of dominated strategies is followed by a best response to the remaining actions is not confirmed in our data.

In conclusion, we believe that the descriptive information about the reasoning process underlying the players' choices is in itself informative. Despite a careful classification our experimental approach allows us to find tight bounds on the individual levels of reasoning with the majority of players not making more than one step of reasoning. The players' belief about the mean of the level-0 players' action is 50 or slightly higher, empirically supporting the assumptions made in previous contributions. The population belief of the classified players above level-1 is mostly identified to be of the degenerate type. It seems unnecessary to consider more complex types of population beliefs in the level of reasoning model, which ultimately enables us to have a workable model with less parameters.

1.6 Maximum Likelihood Estimation

In the previous section we have presented the results of the classification procedure. We saw how we can learn about the level-k parameters of individuals. Now we use the obtained intervals on the levels of reasoning to estimate a structural model for the entire population. We assume a degenerate population belief as suggested by the communication, but the classified level-0 beliefs do not enter the estimation.[18] We present a generalised model which nests the previously presented models and estimate its parameters with a maximum likelihood estimation.

slightly if other teams do the same as ourselves, therefore I suggest 38. I am open to your theory"

"i guess if most people assume the teams average will be 50 then they will expect (2/3) to be aprox 33 and thiss will be their guess which would mean our guess should be 22 however if other people assume the same then (2/3) average will be more like 7 want to go with somewhere in the middle...# 12?".

[18]They enter the analysis on the precision of best responses in appendix A.2.

1.6.1 An Estimable Model

When trying to evaluate the level-k model empirically and testing hypotheses about its parameters, standard maximum likelihood methods cannot be reliably applied. The problem is that for any distribution $l(k)$ and $r^0(a)$ most versions of the model make stark predictions on the action of any player with $k > 0$. Any action just slightly off this predicted action will be identified as random level-0 play. Small perturbations of the action data might hence lead to substantially different results when performing the maximum likelihood estimation.[19]

To understand the root of this problem, note first that only the mean of the players' level-0 belief is relevant for their decision. Let this be denoted as $\tilde{a}$. Now, underlying the stark predictions is the fact that all players have the same belief about $r^0(a)$ and hence respond to the exact same mean. Therefore, all players of the same level of reasoning will choose the same 'optimal' response. Since the level-k model usually retains the optimisation principle, room for error is neither in the decision-making nor in the belief formation.

We estimate a modified model along the lines of the generalised model presented in section 1.2. In particular we assume that the players' belief about the mean of the level-0 play comes itself from a distribution $\tilde{r}(\tilde{a})$. This might – as assumed in the versions of the level-k presented earlier – or might not be collapsed at the mean of $r^0(a)$. It means that we allow the actions of, say, two level-1 players to differ if their belief about the level-0 mean is different.[20] This avoids the stark predictions described earlier and allows to apply a maximum likelihood estimation. Furthermore estimating both the level-0 belief distribution and the level-0 action distribution jointly with the level-k distribution allows us to check whether the level-0 actions and beliefs are consistent. We will specify different functional forms for $r^0(a)$ and $\tilde{r}(\tilde{a})$, each of which will depend on two parameters α, β and $\tilde{\alpha}$, $\tilde{\beta}$, respectively.

Let us lay out how we use this generalisation in order to estimate the level-k model with a maximum likelihood estimation. We observe a sequence of N actions $\{a_i\}_{i=1}^N$ and use MLE to get the estimates $\hat{\tau}$ for the population distribution $l(k,\tau)$, $\hat{\alpha}, \hat{\beta}$ for the level-0 distribution $r(a;\alpha,\beta)$ and $\hat{\tilde{\alpha}}$, $\hat{\tilde{\beta}}$ for the distribution of beliefs about the mean of the level-0 distribution $\tilde{r}(\tilde{a};\tilde{\alpha},\tilde{\beta})$. Assuming independent actions, the likelihood function of the sequence $\{a_i\}_{i=1}^N$ is

$$L(\{a_i\}_{i=1}^N) = \prod_{i=1}^{N} \mathcal{F}(a_i;\tau,\alpha,\beta,\tilde{\alpha},\tilde{\beta})$$

[19]In the discussion of the variation in the suggested decisions in appendix A.2, it can be seen in figure 1.10 that only 4 of 31 players played a 'precise' best response. Most players would be identified as level-0 in the mentioned estimation, leading to an upward biased estimate of the proportion of non-strategic players.

[20]This indirectly softens the optimisation principle in the spirit of Stahl and Wilson (1995). Our model allows for heterogeneous beliefs, but not for imprecise best responses. The latter are, however, picked up in the form of a higher level-0 mean belief variance.

The log-likelihood function is then

$$\log L(\tau, \alpha, \beta, \tilde{\alpha}, \tilde{\beta}; \{a_i\}_{i=1}^{N}) = \sum_{i=1}^{N} \log \mathcal{F}(\tau, \alpha, \beta, \tilde{\alpha}, \tilde{\beta}; a_i)$$

Note how the function $\mathcal{F}$ takes into account the structure of the level-k model: Generally an action a_i can be justified by $k = 0$ and many combinations of k and $\tilde{a}$ with $k > 0$. An action a_i of for example 40 could come from a level-0 player. Then the possible likelihood contribution is

$$l(0; \tau) \cdot r^0(40; \alpha, \beta).$$

But as well it can be justified by, e. g. $k = 1$ and $\tilde{a} = 60$. The likelihood contribution of this would be

$$l(1; \tau) \cdot \tilde{r}(60; \tilde{\alpha}, \tilde{\beta}).$$

Among the likelihood contributions of all possible levels k, the pair (k, a^0) or $(k, \tilde{a})$ with the highest likelihood contribution will be selected in order to proceed in the calculation of the log-likelihood function.[21] Formally, this can be reflected as

$$\mathcal{F}(\tau, \alpha, \beta; a_i) = \max\left(l(0; \tau) \cdot r^0(a^0; \alpha, \beta), \max_{\{(k,\tilde{a})|k>0\}} l(k; \tau) \cdot \tilde{r}(h(k, a_i); \tilde{\alpha}, \tilde{\beta}) \right).$$

H$\mathcal{F}(\tau, \alpha, \beta; a_i) = \max\left(l(0; \tau) \cdot r^0(a^0; \alpha, \beta), \max_{\{(k,\tilde{a})|k>0\}} l(k; \tau) \cdot \tilde{r}(h(k, a_i); \tilde{\alpha}, \tilde{\beta}) \right)$irst, given each possible tuple of parameters, we find for each individual the combination of $(k, \tilde{a})$ which gives the highest likelihood contribution and calculate the 'maximum' log-likelihood given the tuple of parameters. In a second step we then find the maximum likelihood estimate for the structural parameter of the model. When we subsequently talk about the 'level identification of individuals from the estimation' we will refer to the values of k which were obtained in the first step of this procedure for the maximum-likelihood estimates.

While this is the way the estimation is constructed, the first step of the maximisation can be controversial because the second step is then taken as if the subject's level was known with certainty, which it is not. It might be better justified to maintain a probabilistic view on the level throughout the estimation. A mixture model with the following log-likelihood function would reflect that a subject is of level-k with a certain probability.

[21]Since we are interested in estimating a general model of reasoning in one-shot games, we included all players in this estimation, even if they were identified as dominance or equilibrium reasoners. We have checked that our later analysis is robust to the exclusion of equilibrium players.

$$\log L(\tau, \alpha, \beta, \tilde{\alpha}, \tilde{\beta}; \{a_i\}_{i=1}^N) = \sum_{i=1}^{N} \log \left(\sum_k l(k, \tau) \cdot r^k(h(k, a_i), \alpha, \beta, \tilde{\alpha}, \tilde{\beta}) \right).$$

It is assumed that this alternative approach will not lead to major differences in the estimates to be presented at a later point.

For both approaches, note that the population belief pins down the mean of the level-0 belief $h(k, a_i)$ as a function of the level k and the played action a_i. These beliefs will generally differ for degenerate or non-degenerate population belief structures.[22] Note as well that the number of possible levels might be limited by the fact that for higher levels, the corresponding mean of the level-0 belief will eventually be higher than 100.[23] Furthermore, the information from the communication data will be imposed as constraint on the possible levels. Any levels that can be excluded thanks to the communication analysis will be associated with a zero likelihood.

We mostly assume the population distribution to be Poisson as in equation 1.2, but we also show general results without a distribution assumption. For the distributions $\tilde{r}(\tilde{a})$ and $r^0(a)$ we use a bounded normal distribution with mean α and standard deviation β, respectively. We also ran estimations with a beta distribution with parameters α and β which nests the uniform distribution as a special case. There are no notable differences between the estimation outcomes.

1.6.2 MLE with communication information

When using the lower and upper bounds on the levels of reasoning as constraints in the estimation the likelihood function is well-behaved and exhibits an interior global maximum as can be seen in figure 1.7.[24] Furthermore, Monte Carlo studies show that this estimator is indeed unbiased for all five parameters of interest. This is not the case for estimators which use only the lower bound or no bound at all, which is precisely why MLE could not be applied in other studies.

The estimates obtained when applying this procedure are presented in table 1.5.[25] They allow us to scrutinise some of the assumptions made in the literature and hence test the different versions of the level-k model presented earlier.

Level-0 Actions and Beliefs Consider first our estimates on the anchoring element of the level-k model: the distribution of non-strategic ('level-0') play. In the previous literature this has been assumed to be uniformly distributed on the action space. The belief of other players about the mean of the non-strategic play was then assumed

[22]If we decided to have heterogeneous types of population beliefs, this would be reflected by $h_i(k, a_i$.

[23]Whether this can happen depends on the type of population belief.

[24]Note that the points in the graph represent the likelihood for different values of the level-0 mean and the population τ after maximising over the parameters which are not represented in the graph.

[25]The parameters are the mean and standard deviation of the bounded normal distribution.

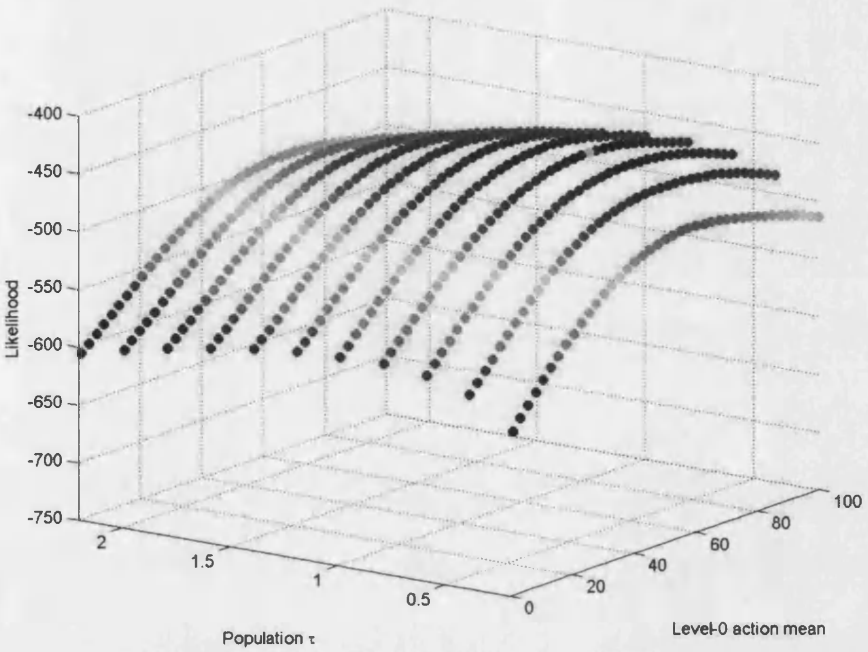

Figure 1.7: Log-likelihood contours under lower and upper bounds from communication.

to be 50. In line with the latter assumption, our estimate of the mean of the level-0 belief of those players who are estimated to have a positive level of reasoning is indeed estimated close to 50. Figure 1.8 shows the pdf of the estimated level-0 belief distribution.

However, this belief is not consistent with the actual mean of the non-strategic play, which is estimated to be substantially higher at 60.14. When simulating the distribution of our estimator under the null hypothesis of a true level-0 action mean of 50, the 97.5th percentile is 58.07, below our estimate. We can therefore reject the hypothesis that the true level-0 action mean was 50 at the 5% level. See appendix A.1 and Study 2 in table 1.7 for details. Furthermore, the level-0 action distribution does not resemble a uniform distribution. We also generate data according to a uniform distribution and use an ML-estimator based on the beta distribution. The 97.5th percentile estimate for the distribution's α-parameter is 1.98, substantially below the estimate of 2.48 for our data. Figure 1.9 shows the resulting shape of the distribution with a mean of the bounded normal of 60.14 and mode of 68.32. Note that both figures 1.8 and 1.9 are consistent with the according data from the communication which was presented in figures 1.5 and 1.6. This suggests that obtaining bounds on the level of reasoning is sufficient to back out the level-0-actions and -beliefs econometrically.

It has been raised that the mean of the level-0 action distribution is estimated to be higher than 50 because the estimation by design attributes actions of levels higher

Table 1.5: Maximum-likelihood estimates.

	Global max
Level-0 action mean	60.14
Level-0 action s.d.	18.85
Level-0 belief mean	53.16
Level-0 belief s.d.	15.04
$\hat{\tau}$	0.82

The estimation was constrained by lower- and upper-bounds, using separate level-0 action- and belief-distributions.

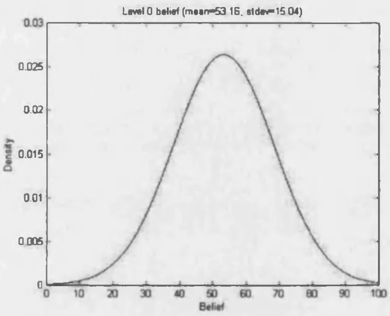

Figure 1.8: Estimated level-0 belief distribution.

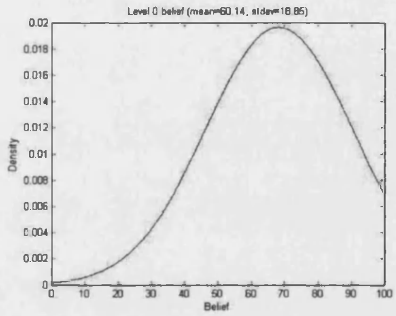

Figure 1.9: Estimated level-0 action distribution.

than 66 to level-0. We believe that this is not of concern for us. Firstly, the levels are often observed and indeed many players have been identified as level-0 while playing a number lower than 66. Secondly, in Monte Carlo studies our estimator is unbiased for the means of the level-0 action and level-0 belief distribution.

The estimates of the level-0 beliefs and actions raise interesting questions. The first one is why the actions are drawn to high numbers. This might suggest a surprising attraction to the upper bound 100, which some not only have as a level-0 belief, but also suggest as action. Furthermore, the number $\frac{2}{3}$ in the rules of the game might make numbers close to 66 somehow salient. But it should be noted that messages of individuals who were classified as level-0 players and play numbers around 66 show no sign of any strategic consideration like dominance reasoning. A second interesting issue concerns the classified level-0 beliefs depicted earlier in figure 1.5. When different from 50 these are almost always higher than 50 – to some extent consistent with the actual level-0 play. However, although both estimates are above 50, the level-0 beliefs do not seem to be consistent with the actions of level-0 players. Thirdly, the distribution of beliefs *about the mean* of the non-strategic play is more concentrated than the non-strategic play itself.[26] Overall, the results suggest that it is problematic to impose consistency between actions and beliefs, as it is to assume a level-0 play centered around 50.

[26]This might hint to the belief formation being adequately modeled as making several draws from the underlying level-0 action distribution and taking the mean over these imaginary draws.

Table 1.6: Our and Other Results.

	Our Data					CGC Data	
	Poisson		*General*		*CHC*		
	(1)	(1b)	(2)	(2b)	(3)	(4)	(4b)
$\hat{\tau}$	0.82		–		0.43		
$\hat{\mu}_0$	0.44^a		0.36		0.65^a		
$\hat{\mu}_1$	0.36^a		0.47		0.28^a		
$\hat{\mu}_2$	0.14^a		0.15		0.06^a		
$\hat{\mu}_3$	0.04^a		0.02		0.01^a		
L0	30	(35.7%)	29	(34.5%)		–	
L1	37	(44.0%)	38	(45.2%)		27	(30.7%)
L2	12	(14.3%)	12	(14.3%)		17	(19.3%)
L3	1	(1.2%)	1	(1.2%)		1	(1.1%)
D1	2	(2.4%)	2	(2.4%)		1	(1.1%)
S2	1	(1.2%)	1	(1.2%)		1^b	(1.1%)
E	1	(1.2%)	1	(1.2%)		11	(12.5%)
n.c.	–		–			30	(34.1%)
N	84		84		84	88	

[a] Values implied by the Poisson distribution with the $\hat{\tau}$ parameter.
[b] 'Sophisticated' without level indication.

Level-k Distribution Finally, let us turn to the second crucial parameter of the level-k model, the average level of reasoning τ. We estimate this to be 0.82. Compared to other studies with a similar student population, this is a low value (See Camerer et al., 2004). Table 1.6 presents our results and those of Costa-Gomes and Crawford (2006). Column (1) presents the types estimated in the constrained maximum likelihood estimation with an underlying Poisson distribution. The first row gives the τ estimate obtained, the next 4 rows give the frequencies of level-types implied by the τ estimate. The bottom of the table presents the frequencies of individual level-estimates. Combined with the classification information on dominance reasoning and equilibrium identification these give the type estimates presented. Column (1b) gives their relative frequency and shows that more than 80% are classified to be either level-0 or level-1 players. When comparing the Poisson frequencies in (1) with the type frequencies in (1b) is can be seen that the shape of the Poisson distribution fits only roughly the estimated type distribution. We therefore performed an MLE estimation without the Poisson assumption, estimating the frequencies of level-0 to level-3 types with 3 parameters. The results are presented in column (2). Column (3) shows the results we would have obtained, had we applied the method of moments approach used by Camerer et al. (2004) to our average 'suggested decision' of 43.93. This method would have led us to think that many individuals were level-0 reasoners, while we actually identify them to have made 1 or 2 steps of reasoning.

Comparing our results to the type distribution of Costa-Gomes and Crawford (2006) in the two last columns (CGC, 88 subjects in total), it is striking that the frequencies of L1, L2, L3, D1 and Sophisticated are very similar.[27] But in contrast to their results we identify a substantial fraction of the population to be playing non-strategically.

1.7 Conclusion

This chapter has presented data and estimation results from a beauty-contest game with team-communication.

Applying the experimental design to the 'beauty-contest' game allows us to find estimates of parameters of a prominent model of boundedly rational play: the level-k model. These estimates cannot be obtained by only looking at choice data. Generally, we can view the communication transcripts as clear evidence for level reasoning by the large majority of subjects. Furthermore, level-0 reasoners do actually exist and are not only a thought experiment of higher level players. The implied level-k distribution is contrasted to the ones found in recent studies. Further, we confirm a level-0 belief distribution centered around or slightly above 50. The actions of level-0 players are, however, still higher than those, implying an inconsistency between beliefs and actions. We do find that higher level players mostly believe other players to be one level below them, an assumption formulated in Nagel (1995). This sheds doubts on the assumption of consistency of the population beliefs with the true populations distribution as is underlying the model of cognitive hierarchy proposed by Camerer et al. (2004).

[27] We have fewer people playing equilibrium, but the equilibrium play in the standard 'beauty contest' might be a bit more subtle since it is not reached after a finite number of rounds of iterated deletion of dominated strategies.

A Appendix

A.1 Monte Carlo studies

We use Monte Carlo studies in order to analyse the properties of our maximum likelihood estimators in section 1.6. By design, the Monte Carlo results are specific to the parameters for which the data is generated. To claim unbiasedness of our estimators, we investigate this property of our estimator for a set of parameter values in the vicinity of the estimate values.

We generate data with 84 datapoints reflecting the total number of datapoints in our study. In a first step, we assume that the true data generating process for the level-0 mean is reflected correctly by random draws from a bounded normal distribution.[28] Similarly, the level of reasoning is assumed to be correctly represented by random draws from the Poisson distribution.

In addition we need to calibrate the data to the output of the classification of the communication. The lower and upper bounds on the levels are generated by assuming uniformly distributed errors in the observation. Unless otherwise stated, they are specified such that the probability of being one level away from the truth is 0.28 for each lower and upper bound. In our sample, 72% of the complete level classifications are precise up to one remaining level value. Using a probability of imprecision of 0.28 for both lower and upper bound gives a higher overall probability that the level classification is only precise up to two level values.[29] Furthermore, there are subjects whose communication does not lend itself to any classification. We reflect this by not giving a value for the bounds in 17% of the observations, matching our 14 out of 84 non-classifications.

In the following, we show results indicating that our estimators are unbiased for a sample size of 84. One estimate of particular interest is the level-0 action mean. We show results for three different values of this mean in the true data generating process. We add another set of results that assumes a more successful classification procedure in order to analyse the estimates' sensitivity in this respect (Study 2).

It can be seen that the estimates are generally very close to the true underlying parameter values, indicating that potential biases are of negligible magnitude. We are therefore confident that the results presented in section 1.6 are informative about the parameters of the level-k model in our sample.

[28] We also run studies with an underlying beta distribution that nests the uniform distribution. The estimation fares slightly worse but delivers qualitatively very similar results.

[29] In our sample we have one classification with three possible level values, which corresponds to 1.2%. In our calibration such imprecision arises with probability 7.9%.

Table 1.7: Monte Carlo results for different level-0 action means.

		Study 1		Study 2	
	Data generation	Mean	St. dev.	Mean	St. dev.
Level average	0.80	0.69	0.16	0.76	0.11
Level-0 action mean	50.00	47.75	6.20	49.82	4.21
Level-0 action st. dev.	19.61	18.35	2.47	17.89	2.21
Level average	0.80	0.75	0.13	0.79	0.07
Level-0 action mean	56.94	57.39	5.22	57.21	2.68
Level-0 action st. dev.	19.13	18.20	2.51	18.58	1.74
Level average	0.80	0.76	0.10	0.79	0.12
Level-0 action mean	61.61	61.61	5.72	61.42	4.35
Level-0 action st. dev.	18.14	17.66	2.92	17.48	2.61

A.2 Variability in the best responses

Our data allows us to look at the discrepancy between the observed suggested decision on the one hand and the decision that would be prescribed by the stated level-0 belief and the estimated level on the other hand. By comparing these two values we can analyse to which extent the assumption of probabilistic best responses as in Stahl and Wilson (1995) is appropriate. For this purpose we use the ML-estimates of the levels from section 1.6 and the stated level-0 beliefs from the communication to calculate the '*precise*' best-response.

We restrict ourselves to the 31 observations for which we have the stated level-0 belief and which are higher than level-0. Figure 1.10 depicts the absolute deviation of the suggested decision from this theoretically expected decision. It shows a roughly symmetric distribution centered at 0. The median deviation is 0, implying a balance between deviations to both sides. The mean deviation of 1.55 is slightly higher than 0 and influenced by some outliers on the right hand side of the distribution. The standard error is 6.75, which indicates that there is a substantial fraction of players playing an action distant from a 'precise' best response. It indicates that the assumption of a probabilistic best response as in Stahl and Wilson (1995) probably reflects actual decision making in a better way than the optimisation principle.[30]

At the same time, it might be a stretch to take the stated number for the level-0 belief too literal. Rather than picking up noise in the best response, we might be picking up temporal variations in the beliefs.

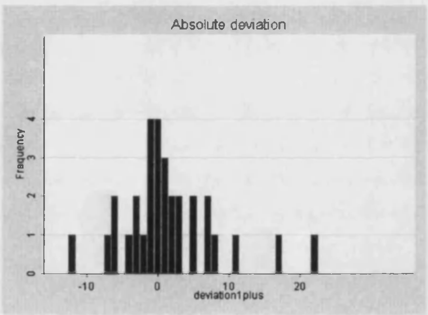

Figure 1.10: Deviation of suggested decision from 'precise' best response.

[30]This result can be seen as an additional piece of evidence in the literature on imprecision in choices and valuations, extended to best responses to beliefs. See Butler and Loomes (2007) and the literature cited therein.

A.3 Classification and estimation by subject

Table 1.8 to 1.10 on the next pages show the estimation and classification by subject.

Table 1.8: Classification and estimation by subject (Sessions 1-2).

Session	Subject	Suggested decision	Lower bound	Upper bound	Estimates	L0 mean	Degen. pop. belief	Equilibrium Id	Dominance	Type
1	1	35	1	1	1	50		0	0	L1
1	2	43	1	1	1	60		0	0	L1
1	3	22	2	2	2	50	1	0	0	L2
1	4	60	0	0	0	55		0	0	L0
1	5	60	0	0	0			0	0	L0
1	6	40	1	1	1	60		0	0	L1
1	7	15	1	2	2			0	0	L2
1	8	50	1	1	1	50		0	0	L1
1	9	16	0	0	0			0	0	L0
1	10	60	0	1	0	60		0	0	L0
1	11	40	0	0	0			0	0	L0
1	12	41	0	1	1			0	0	L1
1	13	21	2	2	2	50	1	0	0	L2
1	14	40			1					L1
2	1	32	2	2	2	55	1	0	0	L2
2	2	40			1					L1
2	3	42	0	0	0			0	0	L0
2	4	58	0	0	0			0	0	L0
2	5	42	1	1	1	65		0	0	L1
2	6	50	0	0	0			0	0	L0
2	7	20	1		2			1	0	S2
2	8	75	0	1	0	50		0	0	L0
2	9	30			1					L1
2	10	32	1	1	1	50		0	0	L1
2	11	50	0	0	0			0	0	L0
2	12	22	1	1	1			0	0	L1

Table 1.9: Classification and estimation by subject (Sessions 3-4).

Session	Subject	Suggested decision	Lower bound	Upper bound	Estimates	L0 mean	Degen. pop. belief	Equilibrium Id	Dominance	Type
3	1	26	1	1	1	50		0	0	L1
3	2	40	1	1	1			0	1	D1
3	3	34	1	1	1	50		0	0	L1
3	4	38	1	2	1	50		0	0	L1
3	5	61	0	0	0			0	0	L0
3	6	44	1	1	1	66		0	0	L1
3	7	65	0	1	0	50		0	0	L0
3	8	52	0	0	0			0	0	L0
3	9	77	0	0	0			0	0	L0
3	10	23	2	2	2	50	1	0	0	L2
3	11	40			1					L1
3	12	25	2	2	2	50	1	0	0	L2
4	1	56	0	0	0			0	0	L0
4	2	28	1	1	1			0	0	L1
4	3	30			1					L1
4	4	60	0	0	0			0	0	L0
4	5	43	0	1	1			0	0	L1
4	6	90	0	0	0			0	0	L0
4	7	66	0	1	0			0	0	L0
4	8	32	1	1	1	50		0	0	L1
4	9	35	1	1	1	50		0	0	L1
4	10	82	0	1	0			0	0	L0
4	11	36	0	1	1			0	0	L1
4	12	88	0	0	0			0	0	L0
4	13	40	0		1			0	0	L1
4	14	50	0	1	0	50		0	0	L0

Table 1.10: Classification and estimation by subject (Sessions 5-6).

Session	Subject	Suggested decision	Lower bound	Upper bound	Estimates	L0 mean	Degen. pop. belief	Equilibrium Id	Dominance	Type
5	1	57	0		0			0	0	L0
5	2	38	0		1			0	0	L1
5	3	25	1	1	1	55		0	0	L1
5	4	30	0		1			0	0	L1
5	5	45	1	1	1	57.5		0	0	L1
5	6	72	0		0			0	0	L0
5	7	63			0					L0
5	8	16	2	3	2		1	0	0	L2
5	9	34	1	1	1	50		0	0	L1
5	10	35	1	2	1			0	1	D1
5	11	38	2	3	2	100	0	0	0	L2
5	12	86	0	1	0			0	0	L0
5	13	12	3	3	3	50	0	0	0	L3
5	14	40	1	1	1	50		0	0	L1
5	15	26	1	1	1			0	0	L1
5	16	100	0		0			0	0	L0
6	1	35	1	1	1	55		0	0	L1
6	2	100	0	0	0			0	0	L0
6	3	40	2	3	2	66	1	0	0	L2
6	4	0	1					1	0	E
6	5	33	1	1	1			0	0	L1
6	6	38	2	3	2	100	0	0	0	L2
6	7	33	1	1	1	50		0	0	L1
6	8	35	1	1	1	45		0	0	L1
6	9	18	2	3	2		1	0	0	L2
6	10	20	2	2	2		1	0	0	L2
6	11	23	1	1	1	40		0	0	L1
6	12	66	0	1	0			0	0	L0
6	13	100	0	0	0			0	0	L0
6	14	30	1	1	1	50		0	0	L1
6	15	55	1	1	1	50		0	0	L1
6	16	50	0	2	0			0	0	L0

A.4 Experiment instructions

Welcome to the experiment!

Introduction

You are about to participate in an experiment in team decision making. The experiment is funded by the Michio Morishima fund, the London School of Economics and the German Society of Experimental Economic Research. Please follow the instructions carefully.

In addition to the participation fee of £5, you may earn a considerable additional amount of money. Your decisions and the decisions of the other participants determine the additional amount. You will be instructed in detail how your earnings depend on your and the others' decisions. All that you earn is yours to keep, and will be paid to you in private, in cash, after today's session.

It is important to us that you remain silent and do not look at other people's screens. If you have any questions or need assistance of any kind, please raise your hand, and an experimenter will come to you. If you talk, exclaim out loud, etc., you will be asked to leave. Thank you.

Since this is a team experiment, you will at various times be matched randomly with another participant in this room, to form a team that plays as one entity. Your team's earnings will be shared equally between you and your team partner.

The experiment consists of two parts (**Part I** and **Part II**) which are independent of each other and feature different tasks. Part I consists of three rounds and Part II consists of four rounds. However, the way you interact as a team to take decisions will be the same throughout the two parts.

Now, let us explain how your ***Team's Action*** is determined. In fact, both your team partner and you will enter a ***Final Decision*** individually and the computer will choose randomly which one of your two final decisions counts as your team's action. The probability that your team partner's final decision is chosen is equal to the probability that your final decision will be chosen (i. e. your chances are 50:50). However, you have the possibility to influence your partner's final decision in the following way: Before you enter your final decision, you can propose to your partner a ***Suggested Decision*** and send him one and only one text ***Message***. *Note that this message is your only chance to convince your partner of the reasoning behind your suggested decision. Therefore, use the message to explain your suggested decision to your team partner.* After you finish entering your suggested decision and your message, these will be shown to your team partner. She/he will then make her/his final decision. Similarly, you will receive your partner's suggested decision and message. You will then make your final decision. As outlined above, once you both enter your final decision, the computer chooses randomly one of your final decisions as your team's action.

If you have any questions at this point, please raise your hand. In order for you to get familiar with the messaging system, you will now try it out in a **Test Period**. Please turn the page for further instructions.

Test period

A participant in this room is now randomly chosen to be your team partner. The **Test Period** has two rounds, with one question to answer in each round. Since this is only a test, your earnings will not depend on any decision taken now. In both test rounds you will need to answer a question about the year of an historic event. The team that is closest to the correct year wins.

As described, you will be able to send one ***Suggested Decision*** with your proposed year and an explaining ***Message***. After having read your partner's suggested decision and message, you will enter your ***Final Decision***. As described earlier, either your or your partner's final decision will be chosen randomly to be your ***Team's Action***.

The messenger allows ***Messages*** of any size. However, you have to enter the message line by line since the input space is only one line. Within this line you can delete by using the usual "Backspace" button of your keyboard. By pressing "Enter" on the keyboard, you add the written sentence to the message. Please note that only added sentences will be sent and seen by your partner. *The words in the blue input line will **not** be sent.* You can always delete previously added sentences by clicking the "Clear Input" button. The number of lines you send is not limited. You can therefore send messages of any length. You finally send the message to your partner by clicking the "Send Message" button.

When you are ready, please click the "Ready" button to start the **Test Period**.

Start Part I

You are about to start Part I of the experiment. You are now randomly matched with a new team partner. For each of the next three rounds you will be matched with a new team partner, i. e. in each of the following rounds you will play with a different person.

In each round, once all teams' actions have been taken, the computer will let you know which of your final decisions has been chosen randomly as your team's action. It will also let you know all other teams' decisions, whether your team won the round and your personal earnings. In each round, the winning team earns £20 (£10 per team player).

Then the next round of the game follows. It will feature an identical task, but you will be matched with a new team partner.

Your task is the following:

Your team and all other teams will take their ***Team's Action*** by choosing a number between 0 and 100. 0 and 100 are also possible. Only whole numbers will be accepted. From all teams' actions, the computer will calculate the ***Average***. Two thirds ($\frac{2}{3}$) of this average will be your ***target*** number. The winning team will be the one that is ***closest to two thirds of the average***. If two or more teams are equally close, the prize will be randomly given to one of these teams.

As described earlier, you will send your team partner a ***Suggested Decision*** and a ***Message***. Remember to explain in the message your reasoning behind your suggested decision. *(And note again that the words in the blue input line will **not** be sent. Press "Enter" to add them to the message.)* After this information is exchanged, both of you enter your ***Final Decision***, from which the computer randomly chooses the ***Team's Action***.

When you click the "Ready" button, you will start the first round of **Part I** of the experiment.

Start Part II

You are about to start Part II of the experiment. You are now randomly matched with a new partner. For each of the next four rounds you will be matched with a new team partner, i. e. in each of the following rounds you will play with a different person.

In this part of the experiment your team will play against only one other team. In each of the four rounds you play against a different team. From the four rounds, one round is chosen randomly and will be considered for determining the payoff. If your team wins this selected round, your team will earn £10 (£5 per team player). Please note that you will be informed of your opponent's team action of the chosen round at the end of Part II. There will be *no feedback* after the individual rounds.

Your task is the following:

In the beginning, the computer will tell you whether your role throughout Part II is "***Hider***" or "***Seeker***".

If you are ***Hider***, your task is to hide an object behind one of four items. In rounds 1 and 2, the object is a ***Treasure***. In rounds 3 and 4, the object is a ***Mine***. The hider team wins the round if the treasure was not found by the seeker or if the mine was found by the seeker. The seeker does not observe where you hide the object. The seeker will look behind one item in each round, not more and not less.

If you are ***Seeker***, your task is to find the treasure in rounds 1 and 2 and to avoid the mine in rounds 3 and 4. The seeker team wins the round if it chooses the particular item behind which the treasure ***was*** hidden or if it chooses an item behind which the mine ***was not*** hidden.

Just like in Part I, you can send a ***Suggested Decision*** and an explaining ***Message*** to your team partner. *(And note again that the words in the blue input line will* ***not*** *be sent. Press "Enter" to add them to the message.)* From your two ***Final Decisions*** the computer again chooses the ***Team's Action***.

When you click the "Ready" button, you will start the first round of **Part II** of the experiment.

A.5 Classification instructions

It is worthwhile to mention that both RAs had seen the level-k model in a course on Behavioural Game Theory at LSE.

Classification Part 1

In the following we will describe the classification process for the analysis of our experiment. We, Konrad and Stefan, assume that you are familiar with the level-k model as it has been introduced by Nagel (1995) or represented by Camerer et al. (2004). However, in order to clarify potential questions of terminology, the appendix A.5 reproduces the main features of the model in the terminology used in this document.

The classification proceeds in two steps, Part 1 and Part 2. You are now provided by us with the transcripts for Part 1. The transcripts differ in the amount of information about the decisions taken. Only in Part 2 will you see the choices of the players that were made.

After your individual classification of each part, you will meet with your co-classifier to reconcile your classification. In this process, try to agree on common classifications if possible and note them in the third sheet. If an agreement is not possible and you keep your initial individual classification, simply note nothing in the third sheet. After you finished this process for Part 1, you will hand in the three sheets and we will provide you with the material for Part 2. If you have questions about the procedure at any point, simply write an email to us and we will clarify any point in a mail to both of you.

For Part 1, follow the instructions of this booklet now. Read them entirely to get an overview and then start the classification. Please read the messages of each player in Period 1 and note for each player the minimum level of reasoning, the level-0 belief mean and the population distribution. Below you find detailed instructions for classifying each player. Please limit yourself to making inferences only from what can clearly be derived from the message stated, i. e. do not try to think about what the player *might have thought.*

IMPORTANT: When you think that the information does not clearly lend itself to any inference, simply do not note any classification. Consequently, do not note anything if no statement has been made! It seems that this time the statements are less clear than in the first round. Please note only those classifications for which you are certain. Make use of the comments space if you are not certain but still want to indicate a feature of the reasoning. Similarly, please comment if the statement exhibits some argument that does not fit the level-k model as we present it here.

Levels: Minimum lower bounds

For the minimum lower bound on the level of reasoning, you should ask yourself: "What is the minimum level of reasoning that this statement clearly exhibits?" Once noted, you should be able to say to yourself: "It seems impossible that the players' level of reasoning is below this number!"

Here we ask you to be very cautious with the classification, not giving away high levels easily. ***Please only write down the highest lower-bound for which you are absolutely certain!*** In part 2 you will be asked to classify the maximum lower bounds of the level of reasoning. This will be the time to be generous with the interpretation of the statements.

Level 0 The player does not exhibit any strategic reasoning whatsoever. Different versions of this might be randomly chosen numbers, misunderstanding of the game structure or giving other non-strategic 'reasons' for picking a number, e. g. taste. Important is that no best-responding to the others' play occurs. There could be considerations of what others might play, but without best responding to it. Examples[31]: "Let's use 50. This is the average between 0 and 100." "It's random, so let's guess something." "My favourite number is 74."

Level 1 This player best responds to something (calculates 'two thirds'). However, he does not realise that others will be strategic as well. Example: "They will all go for a number of about 50-55. So we should do something like 35."

Level 2 This player not only best responds (calculating 'two thirds'), but also realises that other players best respond as well. At level 2 the question about the extent of strategic reasoning of other players can come up. In the theory, this is reflected as a population belief on levels 0 and 1. Example: "Thinking that others play 60, everybody will play 40. So, we should be more clever and play two thirds of 40." "Some will play just play 90, while others will think and play 60 in response. We should therefore play somewhere between 60 and 40."

Level 3 This player realises that others could be level 2 and reacts by best responding to this as well. Put differently, he realises that others realise that others best respond as well. As for all players above level 1, the extent of strategic reasoning by others is important for level 3 reasoners as well. In addition, they have to ask themselves how the level-2 players think about the distribution of level 1 and level 0 players.

Level 4, 5, ... The process continues to higher levels. More levels of best responses and higher orders of beliefs become relevant.

[31]All examples have been made up for illustrative purposes.

Level 0 belief mean

If the comment hints toward a value of the mean of the level-0 distribution, then indicate this value as level-0 mean. Remember, the level-0 mean is the starting point of the reasoning. Players of positive level start best responding on this number. Please note a number as level-0 mean only when this number is not logically derived through level reasoning or the like. If an interval is indicated, please note the average of the lower and upper bound. For example, 'I think the others play around 50-60.' can be noted as a mean of 55. If only a qualitative statement is made about the level-0 mean, try to quantify it if possible. Otherwise, please write a short comment that indicates what is written down. Similarly, if a distribution is specified, please comment precisely on the relevant passage.

The literature usually assumes a mean of 50. Be reminded that this is only a common assumption which should not influence your considerations at this point.

Population belief

The population belief distribution $g_k(h)$ of a player of level k gives the fraction in the population he expects to be of level h. By definition, for level 0 players, the population belief is irrelevant. Level 1 players are defined as believing that all others are level 0. Hence, differences in the population belief distribution can only show up for players who are level 2 or higher. Therefore, we do not expect any statement from you for reasoners below level 2. But even if the level is 2 or higher, be reminded that at points where you think the information does not lend itself to any inference, simply do not note any classification.

We want to distinguish two sorts of population beliefs, distinguished by the degenerateness of the population distribution.

Degenerate Under a degenerate population belief, a player believes that all other players reason exactly one level below themselves. A level 5 player believes everybody else to be level 4. Higher order beliefs are also degenerate. So a level 3 player would think that all others (who are believed to be level 2 players) will believe that all others are level 1 and so on. Example: "Thinking that others play 60, everybody will play 40. So, we should be more clever and play two thirds of 40."[32]

Non-degenerate A non-degenerate population belief gives non-zero probability to more than one lower level. In such a case, a level 2 player believes that both level 0 and level 1 subjects are in the game with positive probability. Higher order beliefs could be either degenerate or non-degenerate. An example would be: "Some will play just play 90, while others will think and play 60 in response. We should therefore play somewhere between 60 and 40." (In the unlikely case,

[32]Nagel (1995) proposed such a population belief.

that the population belief is of the 'degenerate' type but some higher order beliefs are not, please make a note.)[33]

Model and terminology

The level-k model of bounded rationality assumes that players only think through a certain number (k) of best responses. The model has four main ingredients:

Population distribution This distribution on $\mathbb{N}_0$ reflects the proportion of types with a certain level k.

Level-0 distribution By definition, a level-0 player does not best respond. Hence, his actions are random to the game and distributed over the action space, which in our case is $\mathcal{A} = \{\{0\}, \{1\}, \{2\}, \dots, \{99\}, \{100\}\}$.

Level-0 belief In the model, players with $k > 0$ best respond to what they believe the level-0 players play. Their level-0 belief might not be consistent with the level-0 distribution. For best responding, all that matters of the level-0 belief is the mean, which lies in $[0, 100]$. It is frequently assumed that the level-0 distribution and the level-0 belief are consistent, but for the classification this is irrelevant.

Population belief Players do not expect other players to be of the same or a higher level of reasoning. For a level-k player, the population belief is therefore defined on the set of levels strictly below k. It follows that level-0 players have no defined belief, level-1 players have a trivial belief with full probability mass on $\{0\}$, level-2 players have a well defined belief on $\{\{0\}, \{1\}\}$. From level 3 higher order beliefs are relevant as level-3 players have to form a belief about level-2's beliefs.

Classification Part 2

The classification proceeds in two steps, Part 1 and Part 2. You are now provided by us with the transcripts for Part 2. You can now see the choices of the players that were made.

Please consider the information on each player in Period 1 and note for each player the maximum lower bound of level of reasoning and whether the equilibrium has been identified and whether dominance reasoning has been applied in the excel sheet provided. Below you find detailed instructions for classifying each player. Please limit yourself to making inferences only from what can clearly be derived from the message and the action data, i. e. do not try to think about what the player might have thought.

[33] Camerer et al. (2004) proposed a truncated Poisson distribution as population belief distribution.

Be reminded that when you think that the information does not clearly lend itself to any inference, simply do not note any classification. Consequently, do not note anything if no statement has been made! It seems that this time the statements are less clear than in the first round. Please note only those classifications for which you are certain. Make use of the comments space if you are not certain but still want to indicate a feature of the reasoning. Similarly, please comment if the statement exhibits some argument that does not fit the level-k model as we present it here.

Levels: Upper bounds

The upper bounds should give the maximum level of reasoning that could be interpreted into the statement. Therefore, you should ask yourself: "What is the highest level of reasoning that can be underlying this statement?" Once noted, you should be able to say: "Although maybe not clearly communicated, this statement could be an expression of this level. If the player reasoned higher than this number, this was not expressed in the statement!"

Please refer to the level characterisations in Part 1 of the instructions.

Type: Equilibrium identification

With this dummy, you indicate whether the player realised that the unique equilibrium is 0. For this he has to mention the equilibrium action 0. It is not enough to describe a process of downward convergence. The equilibrium might be mentioned anywhere in the statement, so it is irrelevant whether he stops reasoning when he found the equilibrium strategies or whether finds further arguments not to play 0. People will not necessarily use the word 'equilibrium', but they might describe that 'theoretically everybody should play 0' or that 'the process will be going down to 0'.

Set the dummy to 0 if the equilibrium was not identified and to 1 if it was.

Type: Dominance reasoning

With this dummy, you indicate whether the reasoning applied the concept of dominance for the explanation. This is defined as involving iterative deletion of dominated strategies and randomly playing one of the remaining actions or best responding to a distribution over the partner's remaining action space. People will not necessarily use the word 'dominance', but they might describe that 'playing above 66 makes no sense'. Note that 'everybody plays on average 50 so I should not play higher than 34' is not a dominance reasoning, because it rules out strategies based on a distribution on the full action space. Dominance reasoning rules out first and plays a best response then.

Set the dummy to 1 if dominance was used in the argument and to 0 if it was not.

2 Estimating the level-k model: the 'hide and seek' game

2.1 Introduction

The level-k model of reasoning has so far proven to be the most successful model in explaining experimental data on strategic interaction. In particular, it usually provides a framework that reflects the nature of individuals' strategic thinking in one-shot games better than the Nash equilibrium concept. Its structure consists of a level-0 player that is assumed to play randomly, truthfully, or according to salience depending on the situation at hand. Level-1 players play one best response to what they believe a level-0 player will do. Higher level reasoners play iterated best responses to expected lower level play.

This structure is very versatile in explaining data, but the flexibility is also a major disadvantage for a scientific model that wants to be used for predictions. The most fluid element in the model is the level-0 belief, which is also the most important ingredient of the model. If it is desirable to avoid the "circular concept" of equilibrium (Selten, 1998, p. 421), where the starting point of reasoning is given by the correct belief about the action of the opponent, another belief structure has to be defined. I think that it is necessary to let the model assumptions be guided by the actual beliefs exhibited by individuals.

In order to investigate this crucial starting point empirically, this chapter applies the novel experimental design which was introduced in the first chapter of this thesis. It obtains insights in the level reasoning by making use of the transcripts from intra-team communication. In particular, teams of 2 players are playing together as one entity. Each team member can initially send a so-called 'suggested decision' and a justifying 'message' to her team partner in order to explain her arguments. She has an incentive to do so because after the *simultaneous* exchange of their messages, both players give – again individually – a 'final decision', of which one is chosen randomly by the computer to be the 'team's action' in the game.

The reasoning that is revealed in the communication transcripts is new evidence for level-k reasoning in the 'hide and seek' game, a zero-sum game with a unique mixed strategy Nash-equilibrium. Rubinstein and Tversky (1993) introduced this game in which the hider team hides a treasure which is sought by the seeker. Whoever possesses the object at the end of the game wins a prize. The four possible locations are labeled $ABAA$. Faced with this strategic situation, players indeed hypothesise an intuitive reaction of one type of player and iterate best replies until they get to a prescription of play for themselves. The framing induces level-0 beliefs of non-random play, making it worthwhile to engage in level reasoning.

The communication transcripts are classified by two research assistants along the lines of a very general level-k model. This classification is able to provide bounds on most of the individuals' level of reasoning and also gives rich insights in the level-0 beliefs of the players. To make optimal use of it, this information is subsequently used as constraints in a maximum likelihood estimation. Using this experimental design for the 'hide and seek' game enables me to investigate the influence of salience on the level-0 belief as well as potential differences in the reasoning of the two types of players, hider and seeker. Furthermore, for the expected reactions to the salience it might be relevant which type is thought about in the level-0 belief.

As a first result on the nature of the level-0 belief, I find that more than half of the classified players have a B-frame in the sense that they perceive the differences in the locations through the different letters. Less than a fifth of the classified players partition the locations in middle- and endpoints. Other partitions are rare.

Whether the level-0 belief states an attraction or aversion to the B depends on which type of player is thought about. If the hider is subject of the level-0 belief, she is thought to avoid the salient B. If a seeker is though about, he is mostly expected to be attracted by the B. Interestingly, if the middle- vs. endpoints are differentiated, both hiders and seekers are believed to be attracted by the middle locations.

The average level of reasoning is found to be higher for seekers. However, the fraction of non-strategic level-0 players is the same for both types. The difference in average level is due to the fact that there are more level-1 hiders than level-1 seekers and more level-2 seekers than level-2 hiders. The level-distribution for seekers is flat for levels 0 until 2. Interestingly, the reason for this result can be found in the level-0 belief. It can be seen that more players start their reasoning by thinking about a seeker, implying that this is somehow easier or more intuitive.

Since the same subjects played both a 'beauty contest' game – as analysed in chapter 1 – and this 'hide and seek' game, it is possible to do a within-subject comparison of the level estimates. I find that there is no positive correlation across games neither between levels nor between upper or lower bounds. Controlling for the type of player in the 'hide and seek' game does not change this result.

Six of the 8 sessions of the 'hide and seek' game were played with a preceding round of a repeated 'beauty contest' game. Comparing the level estimates in those sessions with the ones in the sessions where only the 'hide and seek' was played, it can be seen that the prior play of the 'beauty contest' increases the level of reasoning in the subsequent 'hide and seek' game. This result raises the question to which extent previous experimental experience of subjects biases results in subsequent experiments.

The chapter is organised as follows. Section 2.2 introduces the level-k model in the context of non-neutral framing and reviews the related literature on the 'hide and seek' game. Section 2.3 presents the experimental design and the classification procedure, before section 2.4 gives the results that follow immediately from the classification. Finally, section 2.5 describes the estimation procedure and gives the resulting findings before the final section 2.6 concludes.

2.2 Focality and the level-k model

'Hide and seek' games have been studied by economists since the early days of game theory (von Neumann, 1953). The class of zero-sum games was the first for which the theory established precise predictions. Empirical tests of the minimax solution of von Neumann and Morgenstern (1944) or of the mixed strategy Nash equilibrium prediction had at some point to face the fact that game theory's abstraction from labelling led to poor predictions. O'Neill (1987) mentioned that the use of the 'ace' in his card experiment alongside the cards 'two' and 'three' "may have been a mistake" since "players were attracted by the powerful connotations of an ace" (p. 2108), playing it more frequently although the game structure treated the three cards equally. The 'hide and seek' game by Rubinstein and Tversky (1993) with the $ABAA$ framing made this notion of influence explicit and led to striking results that did not conform with the predictions of game theory.

The relevance of focal points has first been highlighted by Schelling (1960, p. 57), who mentioned that they provide a clue "for each person's expectation of what the other expects him to expect to be expected to do". In the study by Mehta, Starmer and Sugden (1994) the notion of salience was put under experimental scrutiny. While many factors play a role in defining what can be focal in a specific situation, the required "uniqueness in some conspicuous respect" (Lewis, 1969, p. 35) led Bacharach (1993) to theorise how focality can enter game theory.

In an initially unrelated field of study, Nagel (1995) and Stahl and Wilson (1995) introduced a model of level of reasoning in order to explain experimental data that did not conform with equilibrium, but which still exhibited patterns of strategic reasoning. In this model, a so-called 'level-0' reasoner plays non-strategically a random action. A 'level-1' reasoner plays a best response to this action, assuming that everybody else is a level-0 player. Iteratively, higher level reasoners play best responses to the play of lower level players. A level-2 player needs to form a belief about the relative frequency of level-0 and level-1 players in the population in order to best-respond. This population belief is non-trivial for all levels above level-1.

It can be seen that the level-0 action plays the important role of a starting point for the iterative best responses. The non-strategic nature of this kind of play led Bacharach and Stahl (1997; 2000) to incorporate Bacharach's ideas of focality into the level-0 of a level-k model of reasoning. In the basic theory of focality by Bacharach (1993), a frame consists of one or more classifiers of objects that the player perceives. Each classifier induces a partition of the action space. For example, the letter classifier for a set of actions $\mathcal{A} = \{A_1, B_2, A_3, A_4\}$ would induce a partition $P_l = \{\{A_1, A_3, A_4\}, \{B_2\}\}$.

Bringing this to the level-k model, the non-strategic, frame-induced play of a level-0 player results from uniform randomisation first over the cells in the partition and then over the objects in the cell. For the action space $\mathcal{A}$ in this example this would give the probabilities as in table 2.1. This way, the salience of the B can be captured in terms of probabilities. The intuitive reaction to the framing enters the level-0 and provides a starting point for level reasoning.

A_1	B_2	A_3	A_4
$\frac{1}{6}$	$\frac{1}{2}$	$\frac{1}{6}$	$\frac{1}{6}$

Table 2.1: Probabilities under $P_l = \{\{A_1, A_3, A_4\}, \{B_2\}\}$.

Apart from the innovations in the level-0 play, the level-k model of Bacharach and Stahl (2000) is based on the model of Stahl and Wilson (1995), featuring a probabilistic best response with probabilities of playing a particular action that increase in the expected payoffs from this choice. The population beliefs of high level players are potentially non-degenerate, implying that a level-k player does not only believe level-$(k-1)$ players to be around, but a mixture of lower level players. In Bacharach and Stahl (1997), a simplified version of the 'hide and seek' games of Rubinstein et al. (1996) is investigated theoretically from this theory's perspective. It should be noted that Bacharach and Stahl do not assume that the probabilities as given in table 2.1 apply across tasks. They rather assume the seeker to be attracted to the B while the hiders would intuitively avoid the B.

Another level-k model of framing, in the following proposed by Crawford and Iriberri (2007a), assumes that the level-0 play is symmetric across types of players. The underlying structure of the level-k model follows closely the models of Costa-Gomes et al. (2001; 2006) in that it assumes degenerate population beliefs and optimal responses. Furthermore, for econometric purposes, it is assumed that the level-0 player only exists in the heads of higher level players.

Like the earlier study presented in chapter 1, the current study aims to illuminate the adequateness of the assumptions made in these two models. Therefore, I propose a general model that does not make any symmetry or asymmetry assumption in the level-0 beliefs or level distribution and that does not restrict the population beliefs. The general framework only assumes that players iteratively best respond to the other players which are always believed to stop their iteration earlier than oneself. Of course, different kinds of reasoning, like equilibrium or sophisticated reasoning can also be expected and should be allowed for.[34] The nature of reasoning shall be investigated empirically in this study.

The study of Crawford and Iriberri (2007a) was the first to bring the level-k model with framing to the data, using the aggregate action data of studies by Rubinstein et al. (1993; 1996; 1999). While the action data shows a pattern that suggests some strategic non-equilibrium thinking, it is not clear that level-k is the appropriate model. The data used is rather coarse and does not allow to infer the underlying type of reasoning. To my knowledge, the original studies were not accompanied by questionnaires which

[34]Note that dominance reasoning is excluded from this list since one can reasonably expect this not to arise.

could have indicated the underlying kind of reasoning.

In other games, like the 'beauty contest', more evidence of level-k reasoning has been gathered. Due to the particular game structure, it is possible to associate some actions with particular levels (Nagel, 1995). More qualitative insights have been obtained through questionnaires that were handed out after the experiment, suggesting that level reasoning is applied (Nagel, 1994; Bosch-Domènech et al., 2002). The loose connection of the comments in the questionnaires with the undertaken reasoning, however, stood in the way to use this data more rigorously.

The mentioned connection between actions and levels allowed Costa-Gomes and Crawford (2006) to use action data from a repeated 'beauty contest' in order to investigate the presence of particular types of reasoners. Since actions in the 'hide and seek' game do not allow for conclusions on levels and since response patterns to focality might be fragile, fingerprints of repeated games would not help the analysis in this case.

The experimental design that was proposed in the first chapter for investigating the 'beauty contest' is versatile enough to be of help in this setting. It uses a team-setup, providing communication transcripts that allow for inferences on the individual player's type of reasoning. Given the non-informativeness of the action for inferring the level of reasoning, the key feature of the design is that it does not only look for the level of reasoning in the actions, but rather in the communication statements, which are a natural and reliable mirror of the underlying reasoning.[35]

2.3 Experiment design

The experiment was conducted in 8 sessions in the Experimental Economics Laboratory of the Department of Economics in Royal Holloway (University of London).

2.3.1 Overall structure

The 114 participants were mainly undergraduate students in Royal Holloway. 18-20 students were invited per session. Since some students did not show up, the highest possible even number of students was allowed in the laboratory. The 'hide and seek' game did not differ across the 8 sessions. However, 6 of the 8 sessions were organised such that a three-period 'beauty contest' game was played before the 'hide and seek' game. In all sessions, the number of participants was 12, 14 or 16.[36]

The 'hide and seek' game played was identical to the original games played in Rubinstein and Tversky (1993). The hiders had to hide an object in one of four locations, winning a prize in case the seekers did not find it. The seekers won the prize when they found the object. The object was a treasure and the four locations

[35]For a detailed discussion of related research methodologies, advantages and disadvantages of the design as well as its performance in a different game, see section 1.3.

[36]In the session with 14 participants, one of the seeker teams' decision was used for two hider teams.

were labeled $ABAA$.[37] The participants of the experiment were paid a show-up fee of £5. The winning team won a prize of £10 (£5 per team player).

Participants were randomly assigned into teams of 2 players. The two members were connected through the chat module of the experiment software.[38] The team setup was essential to the investigation of the individuals' reasoning. The same communication structure as in the first chapter was used here, see section 1.4 and figure 1.1.

2.3.2 Classification of communication transcripts

In order to make use of the communication transcripts in a rigorous way it is necessary to code the players' statements. I asked two research assistants to classify the messages along the lines of the general level-k model outlined in section 2.2, following a cautious approach.

The two Ph.D. students from the Department of Economics at LSE were given all players' messages and suggested decisions and the instructions that I wrote after having seen the communication transcripts myself.[39] The instructions that are reproduced in appendix B.5 were self-contained and not complemented by verbal comments. Questions of the research assistants were addressed via an e-mail list that included all three of us. The two RAs individually classified the messages and then met to reconcile their classifications and to produce one joint classification which they both agreed upon, if possible.

The three model parameters of interest were the level of reasoning, the level-0 belief and the equilibrium identification. The statements indicate that many players assume a particular reaction of one type of player and iterate best replies until they get a prescription of play for themselves. From the outset it might seem very difficult to disentangle an intuitive reaction to the salience of a player on the one side and the following steps of best responses on the other side. After all, the main point of both kinds of considerations will be the same, they will advocate the choice of a particular location due to attraction or aversion of the opponent to some locations. What comes across as an intuitively expected reaction in the message might be the result of a non-stated best-response consideration. However, the experimental design provides incentives for the subjects to be as convincing as possible, and therefore to state all steps of reasoning explicitly. The following examples should give an idea to which extent the statements risk to omit explicitly made steps of reasoning.

[37] Four rounds without feedback were played. In the second round, the labels were '1234'. In rounds 3 and 4 the object was a mine, reversing the roles of the players. Round 3 has the $ABAA$ labels and round 4 the '1234' ones. One of the four rounds was chosen randomly to determine of the payoffs. In this study, I will restrict attention to the first round. The later periods will be subject of further research in the future.

[38] The experiment was programmed and conducted with the software z-Tree (Fischbacher, 2007).

[39] As opposed to the 'beauty contest' game, there is no ex-ante association of actions with particular levels in the 'hide and seek' game. Therefore, it was innocuous to provide the action data at the same time as the messages.

Seeker: *"im guessing it's a bluff, so am picking randomly"* (Proposal: A_4)

Hider: *"i think hide item behind the b is dangerous. it is distinguish so the seeker team must open the b. so i choose the last a box."* (Proposal: A_4)

Seeker: *"people usually choose the ones in the middle..so i think the hider has put somewhere else other than in the middle?"* (Proposal: A_1)

Hider: *"i think that the other team will think that we have hidden it behind an 'a' block as there is more of them, so we wouldnt put it behind the obvious b block by itself"* (Proposal: B)

For the reasons indicated above and due to the ambiguity of common language and uncertainties in the composition of the messages, one cannot expect a precise identification of the level of reasoning from the communication. In order to still extract as much information as cautiously possible from the messages, I asked the RAs to state a lower bound and an upper bound for the level of reasoning.[40] The lower bound was defined to be the lowest level of reasoning that the statement clearly exhibits, in the sense that any lower level could be excluded on the basis of the statement. The upper bound gives the maximum level of reasoning that could possibly be interpreted into the statement. These definitions were chosen in order to exclude only levels that were either clearly surpassed or that were surely not stated in the message. To account for the subtleties connected with the classification of the upper bounds, which involve the difficult distinction between instinctive level-0 beliefs and explicit steps of reasoning, the econometric results in section 2.5 are checked for robustness by running estimations without using the information on the upper bounds.

If level reasoning is observed, the level-0 belief characterises the starting point of the argument. It consists of what type of player is thought about and what salient location this player is expected to be attracted or averse to. The most relevant location in a player's level-0 belief is the one that is believed to be most or least attractive to the hider or seeker that the player has in mind. She will emphasise this in the message and hence this prominent location will often be the major piece of information that is extractable in the classification. In terms of the model, the most attractive location is reflected by having the highest probability of being chosen among the 4 locations. From the statement, one can only expect to get an ordinal ranking of the locations in terms of their attractiveness.

The classifiers were asked to give such a ranking over the four locations using a 'more attractive than' relation. With the locations represented by their position from left to right as '1', '2', '3', and '4', a B attraction would be coded as $2 > 134$. This way of coding allows for a very rich representation of the individuals' level-0 beliefs

[40]If two differing classifications were provided, the minimum between the two RAs' classification would be used for the lower bound and the maximum for the upper bound.

which can afterwards be grouped into relevant categories according to the analysis's needs.

Note that a ranking was given explicitly for each level of reasoning between the lower and the upper bound. This is because a certain suggested decision can be justified by a pair of level and level-0 belief. To see this, imagine a seeker that sought at B and that is classified to be either level-1 or level-2. As a level-1 player she is classified to exhibit a level-0 belief of B attraction of a hider. Her play of the B will be justified differently if she is assumed to be level-2. In this case she would have thought about a seeker that is averse to the B in her level-0 belief, which made the hider best-respond by playing the B. It can be seen from this example that the level-0 belief changes with the level in a mechanical way. This happens since, for the same action, the level-0 belief for a hypothesised level-1 player is the best response to the level-0 belief of a level-2 player. Furthermore, the type of player that the belief is about alternates in the level. Seekers, for example, have a level-0 belief about hiders at odd levels 1, 3, 5, etc., and about seekers at even levels 2, 4, 6, etc. Table 2.8 exhibits this connection in a possibly more approachable fashion and shows how level and level-0 belief comove. Although the RAs gave a ranking for each possible level, the comovement was always exhibited in their ranking as described here.

The final parameter of interest is an indicator for equilibrium play. This was introduced to distinguish between random level-0 play and random equilibrium play. The indicator reflects that the player gives arguments for his random play by mentioning that any location is a best response to random play, which distinguishes an equilibrium player from a level-0 player.

For smaller details of the level-k model, I instructed the RAs to indicate reasoning characteristics in the form of optional comments. I explicitly asked them to comment on any form of a non-degenerate population belief. Similarly, I asked them to make a comment if a participant reached a level in which the reasoning comes to the same action prescription as level-1, i. e. it reaches a cycling phase. This would be a player who could potentially reach an argument for randomising over strategies of level-1 to level-5 by observing that all that matters is whether the opponent stops reasoning at just one level below him in this cycle.[41] Identifying this meta-coordination game is the closest that level reasoning can get to uniform randomisation as in the mixed strategy equilibrium. The randomisation over the strategies is, however, not necessarily uniform random.

The two research assistants had only a few minor discrepancies in their judgements and all data that is used was exclusively agreed upon by them.

[41] That is, if he is reasoning under a degenerate population belief as hypothesised in Crawford and Iriberri (2007a).

2.4 Results

This section presents the data on the suggested decisions and the results of the classification. Initially, the background of the students will be summarised.

2.4.1 Descriptive data

A total number of 114 students participated in the sessions, all of which were recruited by the host institution. Of the 21 economics students that participated, 16 were in their first year, 3 in the second year, 1 in the third year, and 1 was a PhD student. 19 students had received training in game theory and 12 had been confronted with the 'hide and seek' game before. The majority of participants had taken part in an experiment before. Table 2.2 and table 2.3 summarise some of the participants' background characteristics.

Table 2.2: Players' background.

Field of Study	*Economics*	*Other*	*NA*	
	21	91	2	
Degree	*Bachelors*	*Masters*	*Ph.D.*	*Other*
	83	15	6	10
Year in Degree	*1st*	*2nd*	*3rd*	*4th or more*
	75	20	12	7
Gender	*female*	*male*		
	56	58		
Age	*Mean*			
	21.19			

Table 2.3: Players' prior training.

	No	Yes
Game theory training	95	19
Prior experiment participation	32	82
Hide and seek experience	102	12

The action data is summarised in figures 2.1 and 2.2. It can be compared to the aggregate data used in Crawford and Iriberri (2007a) depicted in figures 2.3 and 2.4.

In the latter figures it can be seen that, irrespective of the type of the players, a majority plays action 3, the central A_3, while the other positions are chosen with roughly the same probability.[42] This regularity only holds for the seekers in my data. Most of the hiders in my study suggest the last position, the right end A_4, generally avoiding the B. In the following sections the communication data will be described, allowing for a look at the reasoning structures underlying these choices.

[42] Appendix B.2 shows that the disaggregated data for other treasure games with $ABAA$ labels exhibits roughly the same regularity.

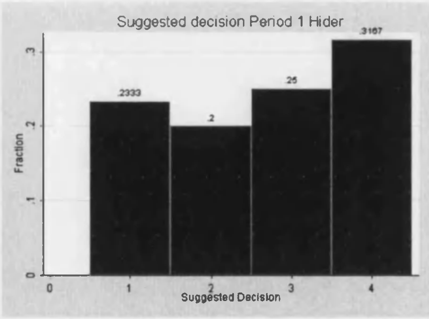

Figure 2.1: Hiders' suggested decisions (N=60).

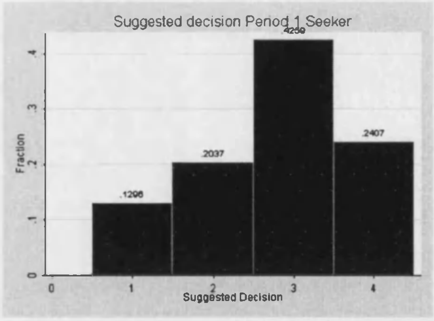

Figure 2.2: Seekers' suggested decisions (N=54).

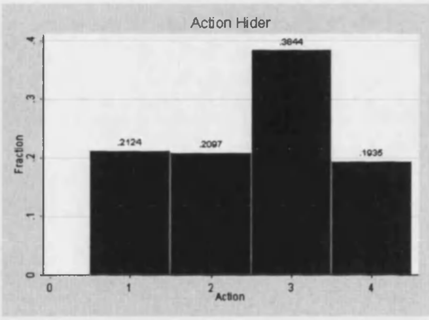

Figure 2.3: 624 Hiders, aggregate as in Crawford and Iriberri (2007a).

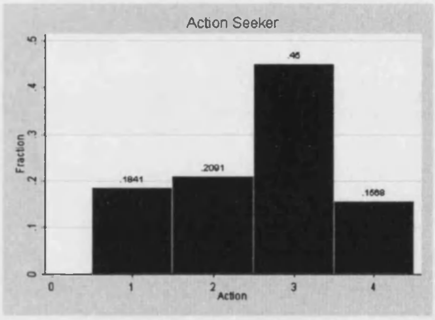

Figure 2.4: 560 Seekers, aggregate as in Crawford and Iriberri (2007a).

2.4.2 Communication information

The classification asked for three pieces of information: the level of reasoning, the level-0 belief and whether equilibrium reasoning was employed. The following paragraphs present the findings one by one.

Level of reasoning Table 2.4 presents the classification information in terms of the lower and upper bounds on the level of reasoning. For 96 of the 114 subjects, both a lower and an upper bound was stated. 2 hiders identified the equilibrium by making an argument for random play as the best response to random play. For one participant only a lower bound was indicated. The messages of the remaining 15 subjects did not allow for an identification of the level interval.

For 57 participants the level of reasoning was fully determined, for 36 it was possible to restrict the set of possible levels to 2 and only for 3 a total of 3 levels were considered to be in accordance with the message. Both the lower and the upper bounds' distributions are hump shaped, as one expects in a homogenous population. If the upper bound was identified, it was never set higher than 4. This implies that no participant reached the phase in which the argumentation would start cycling, as a level-5 reasoner comes to the same conclusion as a level-1 reasoner.[43]

Tables 2.5 and 2.6 show the information by task. It can be seen that both types'

[43] Furthermore, no written comment was given by the RAs that indicated the presence of cycling.

Table 2.4: Level classification results.

	Level Upper Bounds (⌀ 1.29)					
	0	1	2	3	4	Total
0	22	14	0	0	0	36
Level Lower Bounds 1		21	17	0	0	38
(⌀ 0.86) 2			12	5	3	20
3				1	0	1
4					1	1
Total	22	35	29	6	4	96

messages lend themselves to a classification with the same probability (≈ 85%). For both the lower and the upper bounds, the average is higher for the seekers. An interesting detail is that the lower and upper bounds are not hump-shaped for the seekers as opposed to the hiders. The fact that the seekers have two peaks at level-0 and level-2 and the hiders at level-1 suggests that the reasoning often starts by considering what the intuitive reaction of a seeker is. Both of these points shall be investigated in more detail in the estimation section.

Table 2.5: Level classification of hiders.

60 Hiders	*Level Upper Bounds* (⌀ 1.22)					
	0	1	2	3	4	Total
0	9	6	0	0	0	15
Level Lower Bounds 1		19	8	0	0	27
(⌀ 0.82) 2			2	2	2	6
3				1	0	1
4					0	0
Total	9	25	10	3	2	49

Table 2.6: Level classification of seekers.

54 Seekers	*Level Upper Bounds* (⌀ 1.38)					
	0	1	2	3	4	Total
0	13	8	0	0	0	21
Level Lower Bounds 1		2	9	0	0	11
(⌀ 0.91) 2			10	3	1	14
3				0	0	0
4					1	1
Total	13	10	19	3	2	47

Level-0 Belief Of the 96 participants with a classified level interval, for 74 it was possible to obtain a ranking of attraction for the starting point of the level reasoning, the level-0 belief. 22 subjects with a level classification could not be given a ranking of attraction since they were identified as level-0 reasoners that did not react explicitly to the framing in the message.

Once these rankings are available, one can group them according to key characteristics. Natural groupings would be beliefs that exhibit a B-frame, an endpoint-frame or a central A_3-attention.[44] For the sake of completeness, the list of relevant attention patterns is amended by A_1-attention and A_4-attention. Other frames outside this partition are found in only one case. Table 2.7 shows the overall and the type-specific frequencies of the various types of frames.

Table 2.7: Belief classification results.

Frame	B	M	A_3	A_1	A_4	$\varnothing$	$\sum$
Hiders	27	7	1	2	2	1	40
Seekers	21	7	4	1	1	0	34
Total	48	14	5	3	3	1	74

It can be seen that the vast majority of 48 players pays attention to the B_2-frame (75%). 14 (19%) divide the locations into middle- and endpoints. The A_3 frame is quite rare, it might be occurring because the player omitted to state the initial salience of B and the endpoints that made him consider the central A_3. All As are considered salient in very few cases only. One hider exhibited a frame that showed an equal attraction both to B and to A_4, which was not categorised ($\varnothing$). Two types of attentions are considered, attraction and aversion. Note that one can say from the classification whether an attraction or an aversion was exhibited. However, this would be connected to a certain level of reasoning, since for a given action, the justifying level-0 belief changes with the level as outlined in section 2.3.2. I will only be able to investigate patterns of attractions and aversions once the levels are estimated in section 2.5. In table 2.7 it can be seen that the relative frequency of the frames are indeed very similar across types.[45]

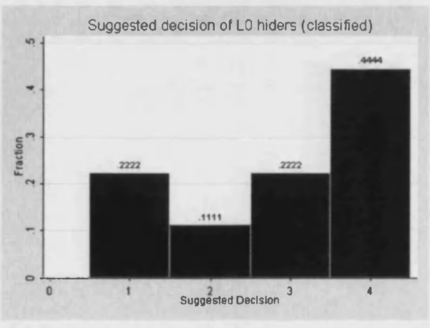

Figure 2.5: Suggested decisions of 9 level-0 hiders.

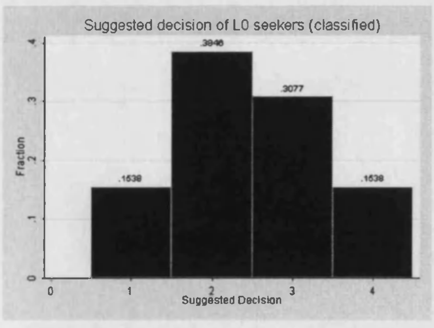

Figure 2.6: Suggested decisions of 13 level-0 seekers.

[44] To give an example, a B-frame would be the one where the second location was either the highest or the lowest ranked position. Whether it is highest or lowest determines the type of the attention, attraction or aversion.

[45] There is a small difference in that the seekers exhibit more A_3-frames. If one takes the stance that the A_3-frame is the opposite of a B-frame in the sense that the very first reaction to the B was not written down, the frequencies would match even better.

Level-0 Actions After considering the intuitive reactions to the framing of the locations, it is possible to check whether the level-0 players indeed follow the frames as in the level-0 belief. Figures 2.5 and 2.6 show that level-0 hiders seem to be attracted to A_4, avoiding the B, while seekers choose mostly B. However, the number of datapoints is limited to 9 hiders and 13 seekers, so that the result is not more than indicative. Note that for the hiders, A_4 is the mode of all players' suggested decisions. For the seekers, the level-0 mode with B is different to the mode of A_3 for the seekers. Whether the attraction/aversion patterns of the beliefs are consistent with these actions will be investigated in section 2.5.2.

2.5 Maximum Likelihood Estimation

The previous section showed that the communication transcripts themselves allow to obtain ample information about the players' reasoning. However, I can now go a step further and obtain a distribution of players' levels in terms of point estimates rather then classification intervals. This will finally allow me to indicate whether attraction or aversion prevails for the individual level-0 frames found in the classification and which type of player, hider or seeker, is this attraction or aversion attributed to. For this purpose, section 2.5.1 introduces an econometric model and section 2.5.2 presents the estimation results.

2.5.1 An Econometric Model

In the maximum likelihood estimation, I want to use all information from the classification about the individual's reasoning to estimate $\hat{\tau}^t$, the parameter of the type-specific Poisson level distribution $l^t(k;\tau^t)$. Note that I allow for two separate distribution parameters in order to investigate level distributions that differ between hiders and seekers.

In addition, $r^t_{c,a}$ is estimated, the probability that types $t \in \{\text{hider, seeker}\}$ are believed to exhibit attention $a \in \{\text{attraction, aversion}\}$ to frame $c \in \{B, M, A_3, A_1, A_4\}$.[46] Since they are considered by type and frame, it has to be that $r^t_{c,att} + r^t_{c,av} = 1 \;\forall\, t, c$. The available information is the sequence of actions $\{a_i\}_{i=1}^N$ and the sequence of frames $\{c_i\}_{i=1}^N$ from the classification. Also, the lower bound of level reasoning $\underline{k}$ and the upper bound of level reasoning $\overline{k}$ will be used. Since actions are independent, the log-likelihood function is:

$$L(\tau^t, r^t_{c,a}; \{a_i\}_{i=1}^N, \{c_i\}_{i=1}^N) \;=\; \sum_{i=1}^{N} \log \mathcal{F}(\tau^t, r^t_{c,a}; a_i, c_i),$$

[46]Since it is not possible to estimate without a given frame, the 1 observation with frame $\varnothing$ cannot enter this estimation.

k	0	1	2	3	4	constraint
Level 0 belief on	B av. S	B av. H	B at. S	B at. H	B av. S	
Level 0 belief on	B av. S	B av. H	B at. S	B at. H	B av. S	$lb = 1$, $ub = 2$

Table 2.8: The way level, action and frame determine the attention is reflected in $h(k, c_i, a_i)$.

where

$$\mathcal{F}(\tau^t, r^t_{c,a}; a_i, c_i) = \max_{\{(k,c,a)|k\in[\underline{k},\overline{k}]\}} l^t(k;\tau^t) \cdot r^t_{c,h(k,c_i,a_i)}.$$

The structure of the level-k model is taken into account as follows: The action a_i of a player with frame c_i determines the frame attention for the level. For example, a seeker playing A_3 and having a B frame exhibits $a = \{\text{aversion}\}$ if she was a level-0 player. The likelihood contribution from such a player would then be

$$l^s(0;\tau^s) \cdot r^s_{B,av}.$$

However, the same action could be justified with the seeker being level-1 reasoner and having a level-0 belief on a hider that is averse to the B:

$$l^s(1;\tau^s) \cdot r^h_{B,av}.$$

Similar to the estimation in the first chapter, the maximisation in the $\mathcal{F}$-function can be controversial. As before, an alternative approach is a mixture model which maintains the structure of having a certain probability of facing a level-k player. The log-likelihood function would be reflected by

$$L(\tau^t, r^t_{c,a}; \{a_i\}_{i=1}^N, \{c_i\}_{i=1}^N) = \sum_{i=1}^N \log\left(\sum_{k=\underline{k}}^{\overline{k}} l^t(k,\tau^t) \cdot r^t_{c,h(k,c_i,a_i)}\right).$$

Again, it is assumed that such an estimation would not lead to different results.

The first line of table 2.8 shows the pattern that emerges for higher levels. The second line shows how the imposition of the lower and upper bounds puts restricting constraints on the estimation.

2.5.2 Estimation Results

Since the individual frames c_i enter the estimation as inputs, only those 74 subjects' levels can be estimated that were given a level-0 frame in the classification. The one player whose level-0 belief could not be categorised in the standard frames was not considered in this estimation.

The estimation results presented in the following do not impose symmetry between hiders and seekers at any point. The level distributions are specific to the types of the players. The attraction/aversion probabilities are specific to the types that the level-0 belief is about. Symmetric constellations are then nested and the estimation hence allows to discriminate between various assumptions on symmetries in the literature. Furthermore, the estimation presented will impose both the lower and upper bound.[47]

Level distribution A first important estimate is the level parameter and the resulting distribution of levels. The average level is differing between types, with $\tau^h = 0.92$ and $\tau^s = 1.11$. Table 2.9 shows the resulting overall distribution of types of reasoning. It shows that it is hump-shaped as expected and similar to the findings in chapter 1 and (Costa-Gomes and Crawford, 2006, CGC).

Table 2.9: Estimated overall level distribution.

	'Hide and seek'		'Beauty contest'		CGC	
L0	30	(26.3%)	30	(35.7%)	–	
L1	40	(35.1%)	37	(44.0%)	27	(30.7%)
L2	20	(17.5%)	12	(14.3%)	17	(19.3%)
L3	4	(3.5%)	1	(1.2%)	1	(1.1%)
L4	1	(0.9%)	–		–	
D1	–		2	(2.4%)	1	(1.1%)
S2	–		1	(1.2%)	1[a]	(1.1%)
E	2	(1.8%)	1	(1.2%)	11	(12.5%)
n.c.	17	(14.9%)	–		30	(34.1%)
N	114		84		88	

[a] 'Sophisticated' without level indication.

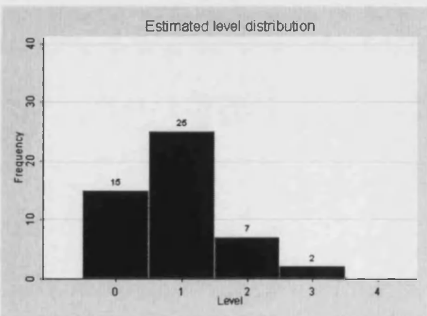

Figure 2.7: Hiders' estimated level distribution.

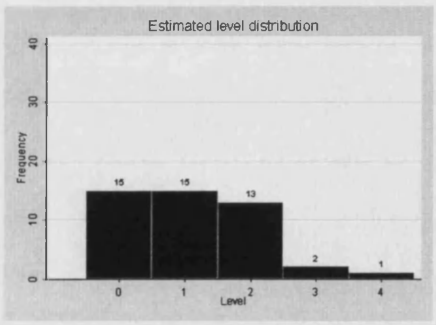

Figure 2.8: Seekers' estimated level distribution.

Rubinstein et al. (1996) hypothesise that the seeker has a psychological advantage because he should be regarded as 'responding' to the hider, having thus naturally an easier access to strategic considerations. By looking at the level estimates by types,

[47]Appendix B.3 presents results of an estimation that does not impose upper bounds.

Table 2.10: Who is subject of the level-0 belief?

	L0 belief		
L0 belief frame	on seekers	on hiders	Total
Total	44	29	73

I can check whether this hypothesis is reasonable. Figures 2.7 and 2.8 show the two distributions.

In a first step, it can be checked whether a higher fraction of seekers exhibits a positive level of reasoning. Interestingly, the number of level-0 players is 15 for both types and the fraction of positive level players is 0.69 (34 out of 49) for hiders and 0.67 (31 out of 46) for seekers. In that respect, the proportion of strategic vs. non-strategic players is the same across tasks.

The difference in the average level seems to results from another phenomenon that I touched upon when discussing the lower and the upper bounds from the classification. Here again, it can be seen that there are more level-1 hiders than level-1 seekers. Also, there are clearly more level-2 seekers than level-2 hiders. Note that the more frequent level-1 hiders share with the more frequent level-2 seekers that they start the level reasoning by thinking about the intuitive response of a seeker to the framing. Table 2.10 shows that there are indeed more players that start with a belief on a seeker than on a hider.

It is this phenomenon that seems to be behind the higher level of reasoning which is observed on average for seekers. Seekers that think about a seeker's intuitive action are of levels 2, 4, 6, etc., while hiders are only of levels 1, 3, 5 etc. At the same time, it is not quite clear, why the seeker is an easier starting point for the reasoning. The results suggest that there is something intuitive about starting to think about the last decision maker whose decision does not need to anticipate anybody else's decision.[48]

Level-0 belief The starting point of the reasoning is described by the estimates $r^t_{c,a}$, which – in addition to the frame – allow a statement on the attraction/aversion and also on the object of the belief, hider or seeker. Table 2.11 shows the results of the estimation.

The table shows that a level-0 belief on a hider is connected with B aversion, while it is predominantly connected with B attraction when about seekers. Although it can only be indicative, it is interesting to note that this is consistent with the level-0 action exhibited in figures 2.5 and 2.6, where hiders are indeed avoiding the B and seekers mostly play the B. For the frame that distinguishes middle- and endpoints, it can be seen that the level-0 belief on both types considers an attraction to the midpoints.[49]

[48]Note that in the extensive form this game would be such that the seeker is considered first when applying backward induction under subgame perfection. The result indicates that there is something intuitive about this way of thinking in sequential settings.

[49]This is in accordance with findings in the psychological literature, which are about a natural attraction to middle objects in various settings (Christenfeld, 1995; Attali and Bar-Hillel, 2003).

Table 2.11: Estimates $r^t_{c,a}$.

	t	
	Hider	Seeker
$r^t_{B_2,at}$	0.00	0.75
$r^t_{M,at}$	1.00	0.83
$r^t_{A_3,at}$	1.00	0.00
$r^t_{A_1,at}$	0.00	0.00
$r^t_{A_4,at}$	1.00	0.00

Table 2.12: Who is believed to react how?

	Level-0 belief		
Level-0 belief frame	on hiders	on seekers	Total
B_2, att	0	21	21
B_2, av	20	7	27
M, att	2	10	12
M, av	0	2	2
A_3, att	4	0	4
A_3, av	0	1	1
A_1, att	0	0	0
A_1, av	1	2	3
A_4, att	2	0	2
A_4, av	0	1	1
Total	29	44	73

The next table 2.12 shows how many observations lie behind these estimates. It can be seen that the last 3 frames are quite rare, so that these estimates should not be given significant weight. It is interesting to note, however, that the A_3-framed belief goes in the exact opposite direction of the B-framed belief. The A_3 attraction of hiders and the aversion of seekers fits in the picture of the B-framed belief if it was the counterpart as conjectured earlier.

After distinguishing the level-0 belief depending on which types were thought about, one can also ask whether the beliefs differ depending on who holds these beliefs. Table 2.13 shows that the beliefs held do not differ much across tasks. For the two sufficiently often observed frames B and M, it can be seen that the previously observed patterns hold for both types individually. This indicates that the level-0 belief differs depending on whether it is about a hider or a seeker. However, it does not matter which type holds this belief.

Note that the hiders start more often to think about seekers while seekers start thinking about both types equally often. This pattern might, however, emerge from a superimposition of two effects. One is the expected higher proportion of level-1 reasoners than level-2 reasoners in a standard student population. This favours seekers as the object of hiders' level-0 belief and hiders as the object of seekers' level-0 belief. The other effect is then the conjectured ease to start thinking about seekers.

Table 2.13: Who holds these beliefs?

	Hiders		Seekers	
	L0 belief		L0 belief	
L0 belief frame	on seekers	on hiders	on seekers	on hiders
B_2, att	12	0	9	0
B_2, av	4	11	3	9
M, att	5	0	5	2
M, av	2	0	0	0
A_3, att	0	0	0	4
A_3, av	1	0	0	0
A_1, att	0	0	0	0
A_1, av	2	0	0	1
A_4, att	0	1	0	1
A_4, av	1	0	0	0
Total	27	12	17	17

This effect could apply to both types of players to the same extent and lead to the pattern at hand.

Choices by level and task I said earlier that the 'hide and seek' game is characterised by a non-existing link between the action choices and the level of reasoning of the players. The estimation results allow me to plot the choices by the different levels and types of the players in figure 2.9.

There is no clear-cut pattern in either of the two figures. For hiders, there is a tendency that players choosing B are most likely level-2 players, while the choice of any A would indicate level-0 or level-1 play. For seekers, the picture is even less indicative, playing A_3 hinting to level-1 and B to level-0 reasoning.

2.5.3 Levels by treatment

Six of the 8 sessions of the 'hide and seek' game where conducted after the participants had played three rounds of a 'beauty contest' with the same team decision structure and different teams, see chapter 1. While the absolute level of reasoning is not the focus of this study of the 'hide and seek' game, the data on the two sessions without preceding 'beauty contest' can be used to check for any training effect regarding strategic thinking.

It can be seen that the average level of reasoning is indeed higher for the 'hide and seek' games if a 'beauty contest' treatment had been played before. This result might raise the question whether the previous participation of subjects in other experiments should be taken more seriously than it currently is.

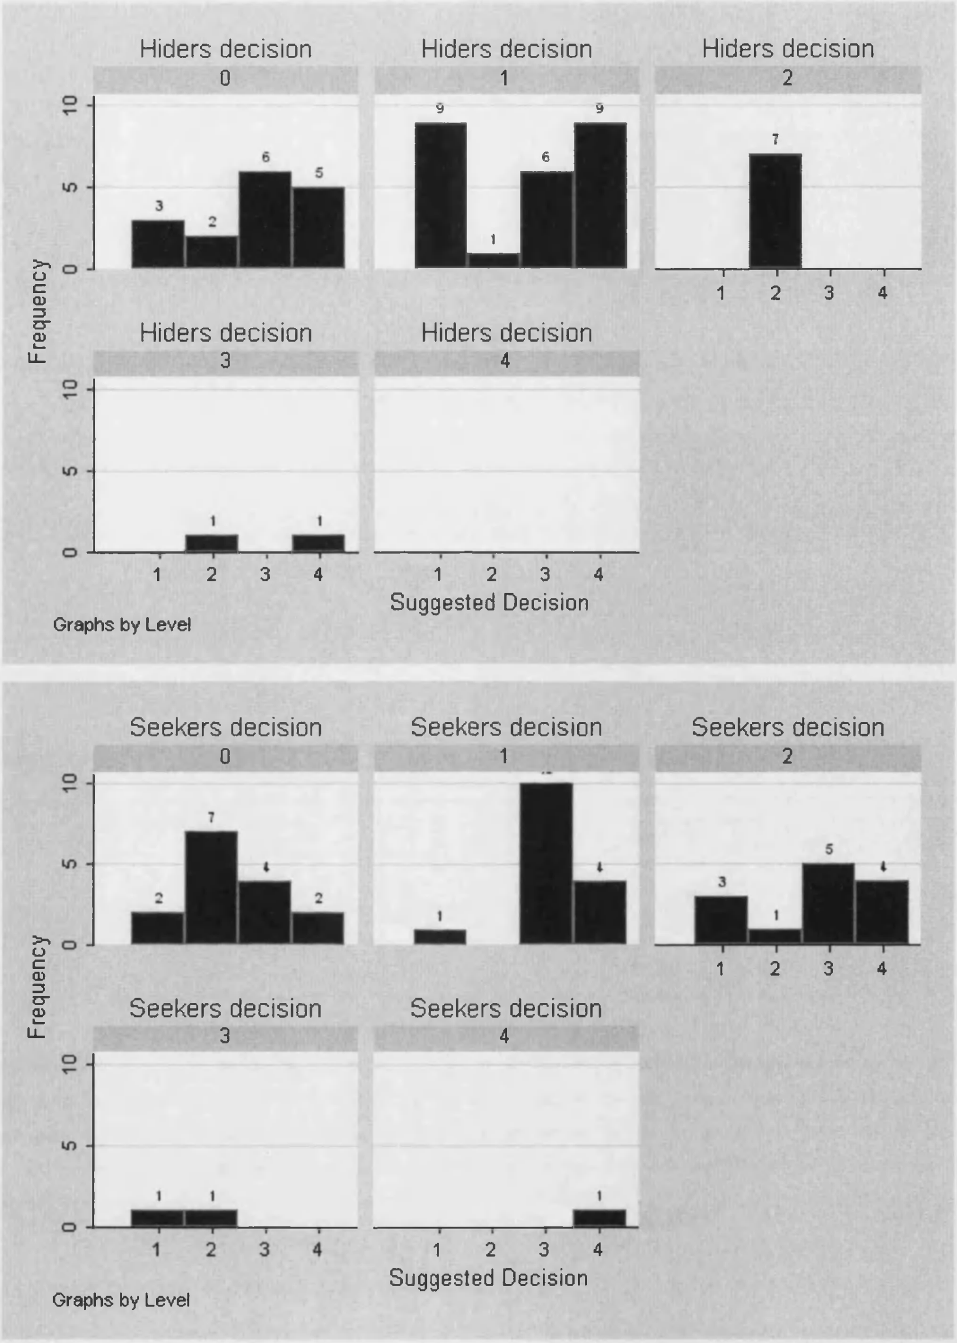

Figure 2.9: Hiders' and seekers' choices by level.

Table 2.14: Summary statistics

	Treatment	Mean	Std. Dev.	N
Hiders	BC/HS	1.00	0.83	36
	HS	0.64	0.63	14
Seekers	BC/HS	1.14	1.06	37
	HS	1.00	0.701	9
Overall	BC/HS	1.07	0.95	73
	HS	0.78	0.67	23

2.5.4 Within subject level comparison between games

For the mentioned 'beauty contest' treatments that were played before the 'hide and seek' treatments, the levels were estimated for individual players in chapter 1. This allows me to do a within-subject comparison of level estimates between the two games.[50] Table 2.15 presents the data of 70 subjects, for which level estimates were found in both games.[51]

It shows that 26 subjects ($\approx$ 37%) are estimated to have the same level of reasoning in both games. 17 have one level more in the 'beauty contest' and 14 have one more in the 'hide and seek'. 13 players in total differ by more than two levels between the games. The correlation between the two levels is slightly negative with -0.028.

In the light of discussions in the literature, these results are quite surprising. Often, a notion of a type of player emerges, which seems to imply that the level of reasoning should be typical for an individual. The results show that the levels of reasoning are highly variable across games and do not even show a positive correlation. At the same time, the shape and the mean of the level distribution of the population seems to be very stable across games. This suggests that the competence to reason strategically is distributed very similarly in the population for a given game, but does change a lot within a subject across games.

Table 2.15: Inter-game level estimate comparison.

		Level 'beauty contest'				
		0	1	2	3	Total
Level 'hide and seek'	0	9	10	3	0	22
	1	7	14	6	0	27
	2	6	7	3	1	17
	3	2	1	0	0	3
	4	1	0	0	0	1
	Total	25	32	12	1	70

2.6 Conclusion

Using a novel experiment design, this chapter has investigated individual reasoning in a 'hide and seek' game. The qualitatively rich data obtained from team communication transcripts indicates that it is indeed adequate to capture non-equilibrium strategic thinking with a level of reasoning model.

The 'hide and seek' game allowed me to look into the influence of focality on the level-0 beliefs. I found that the level-0 belief is mostly shaped by the conspicuous B in the $ABAA$-framing. Hiders are believed to avoid this B while seekers are believed to be attracted by it.

[50] A very similar picture emerges when comparing the lower and upper bounds within subjects and between games. Also, a distinction by task in the 'hide and seek' game does not change the results.

[51] The 4 players that exhibited equilibrium reasoning in either of the games (2 distinct players in the two games) were dropped from the sample.

The two roles that players can take in the game allow to investigate influences of type heterogeneity on the reasoning. While the reasoning seems to be on higher levels on average for seekers, this is not a simple shift of the distribution, but induced by the predominant start of strategic thinking on the seeker's action. This is to some extent reminiscent of backward induction, suggesting that it is somehow easier to think intuitively about a player whose decision is not anticipating anybody else's.

The data allows to undertake a within-subject comparison of levels of reasoning, which leads to the result that the individuals' levels of reasoning are not positively correlated between games. According to this result, the notion of a 'type' of player in terms of level of reasoning should not be evoked across games. Interestingly, the aggregate distribution of levels seems very stable, both between these games and compared to results in the literature. But individually there is apparently a lot of mobility in levels between games.

In addition to the original study in the first chapter, this chapter underpins that the experiment design with intra-team communication is a powerful tool in the investigation of individual reasoning. It provides a versatile structure that can be applied very generally and allows to monitor the understanding and the underlying reasoning in one-shot games.

B Appendix

B.1 Final decisions

For the final decision, the pattern for the seekers is virtually unchanged, so that no major influence from the team communication is apparent. For the hider, the most chosen location is now A_1 instead of A_4, although this does not induce any conjecture as to what is at the root of this change.

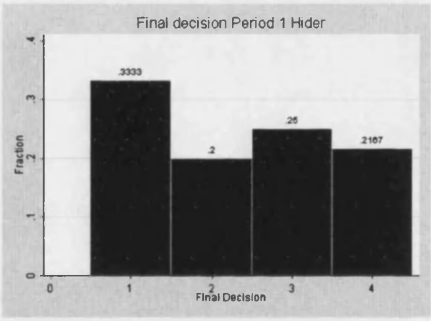

Figure 2.10: Hiders' final decisions (N=60).

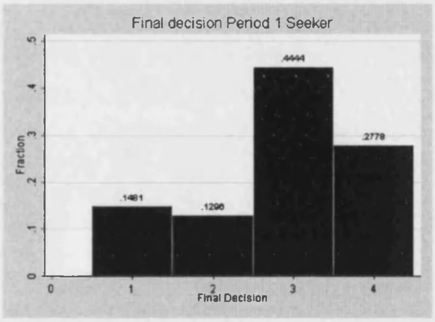

Figure 2.11: Seekers' final decisions (N=54).

B.2 Other *ABAA* treasure games

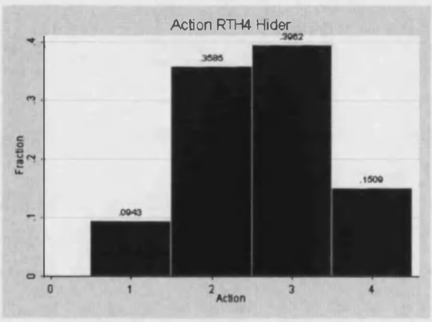

Figure 2.12: 53 Hiders in Rubinstein et al. (1996).

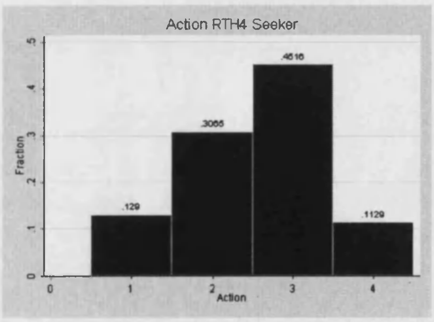

Figure 2.13: 62 Seekers in Rubinstein et al. (1996).

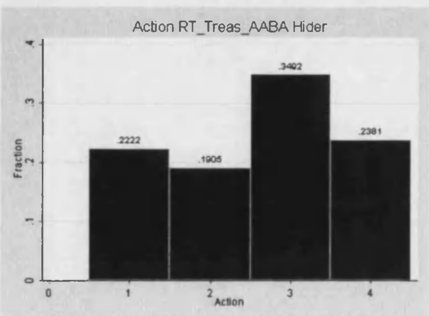

Figure 2.14: 189 Hiders in Rubinstein and Tversky (1993).

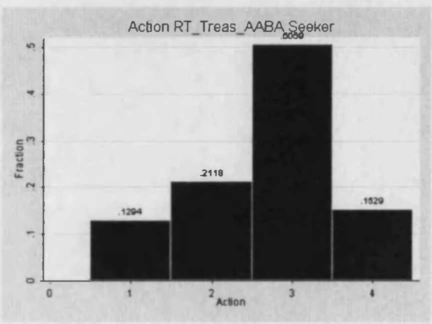

Figure 2.15: 85 Seekers in Rubinstein and Tversky (1993).

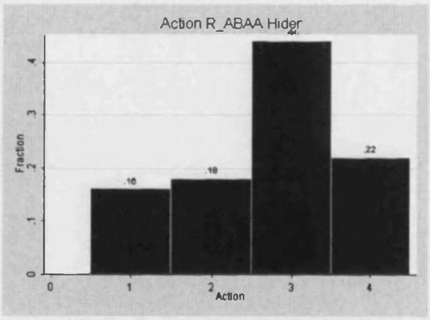

Figure 2.16: 50 Hiders in Rubinstein (1999).

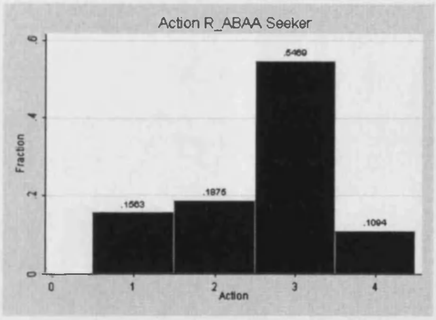

Figure 2.17: 64 Seekers in Rubinstein (1999).

B.3 Estimation without upper bounds

The estimation results that refrains from the use of the upper bounds from the communication are very similar to the results presented earlier in this chapter. The main difference is that the levels are estimated to be higher for both types of players, $\tau^h = 1.18$ for hiders and $\tau^s = 1.20$ for seekers. One seeker is now even estimated to be level-5. The differences between hiders and seekers in the level estimate are much smaller, but the level-1 and level-2 patterns in the level histograms that were highlighted in the text are still present.

The level-0 belief estimates are also qualitatively similar to the ones in the estimation with imposed upper bounds. The asymmetry between types terms of B-attraction and B-aversion is even more pronounced, like the common attraction to the middle-points. The estimates on the other beliefs are entirely unchanged.

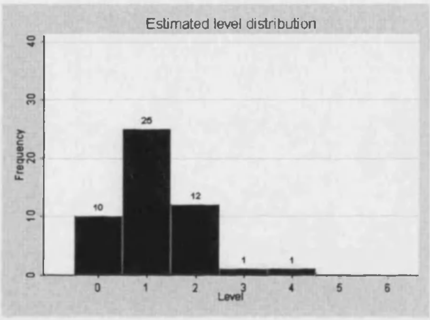

Figure 2.18: Hiders' estimated level distribution with $\tau^h = 1.18$.

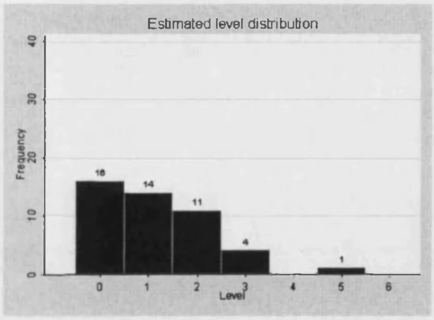

Figure 2.19: Seekers' estimated level distribution with $\tau^s = 1.20$.

Table 2.16: Estimates $r^t_{c,a}$ (no upper bounds).

	t	
	Hider	Seeker
$r^t_{B_2,at}$	0.00	1.00
$r^t_{M,at}$	1.00	1.00
$r^t_{A_3,at}$	1.00	0.00
$r^t_{A_1,at}$	0.00	0.00
$r^t_{A_4,at}$	1.00	0.00

B.4 Experiment instructions

Welcome to the experiment!

Introduction

You are about to participate in an experiment in team decision making. The experiment is funded by the Michio Morishima fund, the London School of Economics and the German Society of Experimental Economic Research. Please follow the instructions carefully.

In addition to the participation fee of £5, you may earn an additional amount of money. Your decisions and the decisions of the other participants determine the additional amount. You will be instructed in detail how your earnings depend on your and the others' decisions. All that you earn is yours to keep, and will be paid to you in private, in cash, after today's session.

It is important to us that you remain silent and do not look at other people's screens. If you have any questions or need assistance of any kind, please raise your hand, and an experimenter will come to you. If you talk, exclaim out loud, etc., you will be asked to leave. Thank you.

Since this is a team experiment, you will at various times be matched randomly with another participant in this room, to form a team that plays as one entity. Your team's earnings will be shared equally between you and your team partner.

The experiment consists of four rounds and the way you interact as a team to take decisions will be the same throughout the experiment.

Now, let us explain how your ***Team's Action*** is determined. In fact, both your team partner and you will enter a ***Final Decision*** individually and the computer will choose randomly which one of your two final decisions counts as your team's action. The probability that your team partner's final decision is chosen is equal to the probability that your final decision will be chosen (i.e. your chances are 50:50). However, you have the possibility to influence your partner's final decision in the following way: Before you enter your final decision, you can propose to your partner a ***Suggested Decision*** and send him one and only one text ***Message***. *Note that this message is your only chance to convince your partner of the reasoning behind your suggested decision. Therefore, use the message to explain your suggested decision to your team partner.* After you finish entering your suggested decision and your message, these will be shown to your team partner. She/he will then make her/his final decision. Similarly, you will receive your partner's suggested decision and message. You will then make your final decision. As outlined above, once you both enter your final decision, the computer chooses randomly one of your final decisions as your team's action.

If you have any questions at this point, please raise your hand. In order for you to get familiar with the messaging system, you will now try it out in a **Test Period**. Please turn the page for further instructions.

Test period

A participant in this room is now randomly chosen to be your team partner. The **Test Period** has two rounds, with one question to answer in each round. Since this is only a test, your earnings will not depend on any decision taken now. In both test rounds you will need to answer a question about the year of an historic event. The team that is closest to the correct year wins.

As described, you will be able to send one ***Suggested Decision*** with your proposed year and an explaining ***Message***. After having read your partner's suggested decision and message, you will enter your ***Final Decision***. As described earlier, either your or your partner's final decision will be chosen randomly to be your ***Team's Action***.

The messenger allows ***Messages*** of any size. However, you have to enter the message line by line since the input space is only one line. Within this line you can delete by using the usual "Backspace" button of your keyboard. By pressing "Enter" on the keyboard, you add the written sentence to the message. Please note that only added sentences will be sent and seen by your partner. *The words in the blue input line will **not** be sent.* You can always delete previously added sentences by clicking the "Clear Input" button. The number of lines you send is not limited. You can therefore send messages of any length. You finally send the message to your partner by clicking the "Send Message" button.

When you are ready, please click the "Ready" button to start the **Test Period**.

Start

You are about to start the experiment. You are now randomly matched with a new partner. For each of the next four rounds you will be matched with a new team partner, i.e. in each of the following rounds you will play with a different person.

In each of the four rounds your team will play against only one other team, a different team each round. From the four rounds, one round is chosen randomly and will be considered for determining the payoff. If your team wins this selected round, your team will earn £10 (£5 per team player). There will be *no feedback* after the individual rounds, you will be informed about your success at the end of the experiment.

Your task is the following:

In the beginning, the computer will tell you whether your role throughout the four rounds is "***Hider***" or "***Seeker***".

If you are ***Hider***, your task is to hide an object behind one of four items. In rounds 1 and 2, the object is a ***Treasure***. In rounds 3 and 4, the object is a ***Mine***. The hider team wins the round if the treasure was not found by the seeker or if the mine was found by the seeker. The seeker does not observe where you hide the object. The seeker will look behind one item in each round, not more and not less.

If you are ***Seeker***, your task is to find the treasure in rounds 1 and 2 and to avoid the mine in rounds 3 and 4. The seeker team wins the round if it chooses the particular item behind which the treasure ***was*** hidden or if it chooses an item behind which the mine ***was not*** hidden.

Remember, you can send a ***Suggested Decision*** and an explaining ***Message*** to your team partner. *(And note again that the words in the blue input line will **not** be sent. Press "Enter" to add them to the message.)* From your two ***Final Decisions*** the computer again chooses the ***Team's Action***.

When you click the "Ready" button, you will start the first round of the experiment.

B.5 Classification instructions

Classification document

Hide and Seek game - Period 1

In the following I will describe the classification process for the analysis of the experiment. It is assumed that you are familiar with the level-k model as it has been introduced by Nagel (1995) or represented by Camerer et al. (2004). The model here is extended to incorporate salience in the level-0 belief according to Bacharach and Stahl (2000). In order to clarify potential questions of terminology and introduce the main features of the model, the appendix reproduces the main features of the model in the terminology used in this document.

After your individual classification, you will meet with your co-classifier to reconcile your classification. In this process, try to agree on common classifications if possible and note them in the third sheet provided. If an agreement is not possible and both of you keep your initial individual classification, simply note nothing in the third sheet. If you have questions about the procedure at any point, please write an email to me and I will clarify any point in an email to both of you.

Please read this document and the instructions for the experiment entirely in order to get an overview and then start the classification based on the player's sent message and action proposal. Note that the framing of the four possible locations is 'ABAA' in periods 1 and 3 and '1234' in periods 2 and 4. A player can be of two types, hider or seeker. It is useful to go through the process first for all hiders and then for all seekers. This way, you keep the perspective of one type of player and do not get confused.

Please read the messages of each player, taking into account his action, and note for each player every possible level of reasoning. These should lie in an interval between the lower and upper bound of the level reasoning, as specified later in detail. Below you find detailed instructions for classifying each player. It is important that you limit yourself to making inferences only from what can clearly be derived from the message stated, i.e. do not try to think about what the player might have thought.

IMPORTANT: When you think that the information does not clearly lend itself to any inference, simply do not note any classification. Consequently, do not note anything if no statement has been made! Please note only those classifications for which you are certain. Make use of the comments space if you are not certain but still want to indicate a feature of the reasoning. Similarly, please comment if the statement exhibits some argument that does not fit the level-k model as I present it here.

Levels

For the lower bound on the level of reasoning, you should ask yourself: "What is the minimum level of reasoning that this statement clearly exhibits?" Once noted, you should be able to say to yourself: "It seems impossible that the players' level of reasoning is below this number!" Here I ask you to be very cautious with the classification, not giving away high levels easily.

The upper bounds should give the maximum level of reasoning that could be interpreted into the statement. Therefore, you should ask yourself: "What is the highest level of reasoning that can be underlying this statement?" Once noted, you should be able to say: "Although maybe not clearly communicated, this statement could be an expression of this level. If the player reasoned higher than this number, this was not expressed in the statement!"

For both lower and upper bound, please refer to the following characterisation of the different levels.

Note that there are two necessary conditions for a player to exhibit a level greater than 0. First, the player has to be responsive to the salience of the games' framing. Secondly, the player has to be strategic in best-responding to his level-0 belief, which is shaped by salience. If he did not react to salience, he would have no reason to chose one over the other object, resulting in random play. Interestingly, this random play is observationally equivalent to equilibrium behaviour. Therefore, the level-0 players can be those that react to salience and do not play strategically, or they can ignore the framing and hence play randomly. As far as possible, I want to distinguish uniform random level-0 play and equilibrium play. However, regarding the level classification, you can classify every random play as level-0 play. The equilibrium play is taken into account by a specific dummy.

Level 0 The player does not exhibit any strategic reasoning whatsoever. Different versions of this might be randomly chosen or purely guessed actions, misunderstanding of the game structure or other non-strategic 'reasons' for picking a location, e.g. by taste or salience. It is important that no best-responding to the other's play occurs. There could be considerations of what others might play, but without best responding to it. Examples[52]: "Well, it's a pure guess", "There are no arguments. Simply choose any."

Level 1 This player best responds to some belief (in the treasure game, a hider mismatches the belief of the other's action, a seeker matches his belief of the other's action). However, he does not realise that others will be strategic as well. Example: "They are probably picking B, so we do as well", "Its at the end and people would naturally go for the middle, no?"

Level 2 This player not only shows the basic strategic consideration of playing best response (matching/mismatching), but also realises that other players best respond as well according to the belief they entertain. A level-2 player clearly contemplates how the other player might best respond to his frame. The player

[52]All examples have been made up for illustrative purposes.

plays a best response to this hypothesised consideration. Example: "They may think most will look in the middle two, therefore choosing one in the end. I therefore choose the first one."

Level 3 This player realises that others could be level-2 and reacts by best responding to the associated expected play. Put differently, he realises that others realise that others best respond to their initial belief. Therefore, a level-3 player clearly states that his opponent expects that he (the level-3 player at question) best-responds to a certain belief.

Level 4, 5, ... The process goes on in a similar fashion. The reasoning enters a cycle, in the sense that level-5 will come to the same best response as level-1. Please indicate in the comment section if the player reaches this cycling phase and recognises this pattern.

Level-0 belief

If level reasoning is observed in the statement, there has to be a starting point in the argument which states an attraction or aversion to one or more of the locations. This is then not derived by strategic reasons, but is an intuitive reaction to the framing of the locations. Otherwise, level reasoning would not occur.

Please indicate the underlying level-0 belief that is connected with each possible level of reasoning. Note that the level-0 belief of a person reasoning on an odd level, i.e. level 1, 3, 5, etc. is always with respect to how a player of the opposite side intuitively reacts to the framing. The belief of a person reasoning on an even level, i.e. level 2, 4, 6, etc. is always with respect to what the opposite type believes about the own type's intuitive reaction. You will have this in mind and can note it for completeness in the "H/S"-box, but since it follows from your level indication, it is not essential.

Usually, the most information one can get out of the communication is the most and least attractive location respectively. For example, as a seeker, I might communicate that the hider is most attracted to the B and then play central A. This indicates that the hider was believed to choose central A with the (weakly) lowest probability and B with highest probability. It follows that it will usually not be possible to rank *all* the locations by their attractiveness.

To reflect the exhibited level-0 beliefs please denote every level-0 belief by ranking the four locations with a 'more attractive than' relation. The locations are coded according to their position as '1', '2', '3', and '4'. In the example of the previous paragraph, the resulting statement is $2 > 14 > 3$, since B in the 2nd position is the most attractive location for the hider and central A, the '3', the least attractive. Not putting anything between numbers indicates that their level of attractiveness cannot be distinguished, as in this case with the two A's at the beginning (position

'1') and at the end (position '4'). Of course, depending on whether you get more or less information out of the communication, your statement can be $14 > 23$ or $2 > 4 > 3 > 1$. This notation should be flexible enough to encode every piece of information on the level-0 belief that is present in the communication.

Imagine a seeker that you classify to be level-1 or level-2. If you determined his level-0 belief to be $2 > 134$ (H) in the case he is level-1, then the level-0 belief should follow to be $134 > 2$ (S) in the case he is level-2.[53] For all cases where nothing in the communication speaks against this, feel free only to note the level-0 belief for the first conjectured level. A statement on each possible level's level-0 belief is only necessary if you think it does not follow in this mechanical way.

In order to keep the overview over the best responses that are connected with certain types at certain levels, I will make a small Excel-sheet available, which calculates automatically the best responses as a function of a specified level-0 belief. This will help you to get a feeling for the mapping of communication and action into the parameters.

Dummy

Equilibrium play In the description of the levels, I said that the level-0 player that does not react to the salience is considered to play randomly due to a lack of arguments for a specific option. Similarly, a player that plays according to the Nash equilibrium will have no argument for or against a specific option, therefore exhibiting the exact same behaviour. In order to distinguish the two where possible, please indicate here whether the player gives convincing arguments for his random play by mentioning that any location is a best response to random play. It is important that the fully random play of the others is considered and used in the best response argument. Put '1' if the player does so. Otherwise, the player will fit the description of level-0 players and the dummy should take the value '0'. Do not tick anything if the player is clearly neither one nor the other type or if the statement does not allow for a classification along these lines.

Concluding remarks

You might have noticed that there will be no classification of the population belief. This is because the opponents as a team are a single entity and a non-degenerate population belief makes only sense in a probabilistic interpretation. Also, the action will always[54] maximally reflect the mode of the probabilistic population belief, making it observationally equivalent to a degenerate population belief. This is why the population belief has not been discussed in the context of this game in the literature.

[53] This follows because for this last level-0 belief for a seeker, the previous level-0 belief of the hider constitutes a best response. Given the action of the player at hand, this is the way the levels and level-0 beliefs co-move.

[54] I abstract from considerations due to the cycling of behaviour from level 5 onwards.

At the current point of the study I let these details remain in the background. Still, please indicate in a comment if a player exhibits any non-degenerate population belief.

Compared to the Beauty Contest, it might at times be difficult to distinguish what is a level-0 belief and what is derived through level reasoning. Try to stick to what is written down and look for clearly stated arguments of reasoning.

Appendix: Model and terminology

The level-k model of bounded rationality assumes that players only think through a certain number (k) of best responses.[55] The model has four main ingredients:

Population distribution This distribution reflects the proportion of types with a certain level $k \in \mathbb{N}_0 = \{0, 1, 2, 3, 4, 5, \ldots\}$.

Level-0 distribution By definition, a level-0 player does not best respond. Hence, his actions are random to the game and distributed randomly over the action space. In our case, the action space is $\mathcal{A} = \{\{1\}, \{2\}, \{3\}, \{4\}\}$ and contains the four possible locations to hide or seek the object. The model incorporates salience by assuming higher probabilities in the level-0 distribution for actions that are salient. In our case, the level-0 distribution would not assign a uniform probability of 0.25 to each possible action, but $p > 0.25$ to the salient one and $q_i < p$ for the remaining actions.

Level-0 belief In the model, the best responses of players with $k > 0$ are anchored in what they believe the level-0 players play. Their level-0 belief might not be consistent with the level-0 distribution. For best responding, all that matters is the expected payoff from choosing a particular location. One would therefore seek (hide) the treasure where the probability is highest (lowest), that the opponent chooses the same location.

Population belief Players do not expect other players to be of the same or a higher level of reasoning. For a level-k player, the population belief is therefore defined on the set of levels strictly below k. It follows that level-0 players have no defined belief, level-1 players have a trivial belief with full probability mass on $\{0\}$, level-2 players have a well defined belief on $\{\{0\}, \{1\}\}$. From level 3 higher order beliefs are relevant as level-3 players have to form a belief about level-2's beliefs.

[55] See the paper by Camerer et al. (2004) for a more detailed account of one version of the model.

3 A model of heterogeneous career concerns

3.1 Introduction

Career concerns have been studied by economists since Holmström (1999, first published 1982) formalised the idea of Fama (1980) according to which concerns about reputation in the labour market suffice to mitigate moral hazard problems. Holmström shows that costly effort is exerted because it increases the agent's expectation of future wages by increasing the market's belief of the worker's productivity.

In recent years, economists have made extensive use of this idea and studied not only reputation for productivity[56] but also for expertise[57] and ideology[58]. The original model of Holmström considers one-dimensional productivity ("talent"), output and effort. In a generalisation, Dewatripont, Jewitt and Tirole (1999a; 1999b) extend this model to multi-dimensional output and effort vectors and to a general stochastic relationship between output, effort and talent.

Apart from an appropriate use of past information and a sufficient weight of the wage revision process, one major prerequisite for effort to arise is the presence of uncertainty about the agent's talent. In Holmström's model the effort is monotonically increasing in uncertainty. One of Holmström's extensions considers stochastic talent in order to avoid the situation of full revelation, allowing for stable effort exertion.

In this chapter I explore implications of heterogeneity among the agents in addition to varying talent. The agents differ in the magnitude of career concerns, which leads them to exert different levels of effort in a given situation. If principals are uncertain about this aspect of the agent's type, I find that depending on the circumstances, this additional uncertainty might decrease total effort exertion over time.

Looking at the real world it seems to be very realistic to think about heterogeneous levels of career concerns. People have different aspirations, different cost of effort, value more or less their income in the future or might exhibit optimism. Furthermore, it does seem that the level of effort varies considerably across individuals.

The most closely related paper to this study is by Kőszegi and Li (2008). In their analysis the heterogeneity arises due to varying 'drive' or 'ambition' of the agent. They introduce heterogeneity by assuming a differing marginal utility of income. While they use a three-period model to depict the implications of the heterogeneity, much of the general intuition about heterogeneity can be found in their analysis. The analysis with infinite periods in section 3.3 is based on their theoretical setup. To my knowledge it is the only work that considers the implications of additional heterogeneity.

[56]See, for example: Tirole (1994).

[57]See, for example: Scharfstein and Stein (1990), Prendergast and Stole (1996), Ottaviani and Sorensen (2006), Levy (2007), Prat (2005).

[58]See, for example: Coate and Morris (1995), Morris (2001), Maskin and Tirole (2004).

Generally, heterogeneous career concerns can be introduced in other ways. First, the level of concern for the future is reflected in the time discount rate δ. People that value their present situation more have a higher δ and hence exert less effort for the purpose of reputation building.

The second, relatively obvious possibility is a differing cost of effort. This point is not specific to the career concerns model and would lead to varying effort levels under explicit incentives as well. Still, this interpretation is fully in line with the analysis of different levels of career concerns.

Third, the set of relevant future employers can be different so that varying firm productivities lead to different income prospects. For example, agents that are willing to go abroad have potentially higher career concerns due to higher payment expectations. Dewatripont, Jewitt and Tirole (1999b, p. 216) consider this to be an interesting way to think about incentives in government agencies. The study by Wilson (1989) suggests the relevance of incentives from sources external to the agency. In previous work I considered the implications of facing multiple principals in career concerns under certainty (Penczynski, 2006).

I show in this study that there are setups in which the introduction of uncertain type traits strictly decreases the effort exertion of all types of agents. This follows from an effect which was coined 'backward attribution' by Kőszegi and Li (2008). Here, the muting of effort follows from the attribution of past output to effort rather than talent in the case of high present effort (and output), implying an ex post downward correction of the inferred talent. It can be shown that this effect is responsible for a general effort reduction in any two-period career concerns model.

In the analysis of setups of more than two periods, an additional 'forward attribution' enters the stage since it becomes desirable to signal high career concerns. A high effort level in the beginning 'promises' high effort in remaining periods due to high expected career concerns, potentially increasing the future wage due to a higher expected effort level. Of course, in a setup with more than two periods, both forward and backward attribution are present, making the analysis of effort exertion in any period depend on the past, present and future effort allocations, information updates and incentives.

In order to present the main results of such a consideration, the chapter is structured as follows. Section 3.2 uses the framework of Dewatripont et al. (1999a; 1999b) in order to consider the two-period setup in full generality. Section 3.3 uses the setup of Kőszegi and Li (2008) to analyse the multi-period situation. Extensions and limitations are discussed in the concluding section 3.4.

3.2 The general two-period model

Although I mentioned equivalent ways of thinking about the heterogeneity in a career concerns model, I will stick to the interpretation of facing multiple principals or labour markets in upcoming periods. This view assumes that agents differ with respect to their set of principals or the subjective probability they attach to elements in this set. While I refrain from modelling the underlying considerations that give rise to this heterogeneity, I suppose it is very realistic to assume that each agent has its own view about his future, a view that might be influenced by his ambition, private priorities, optimism and so forth. Due to the multitude of factors it should be the rule rather than exception that agents differ in the magnitude of their career concerns.

Making use of this interpretation, it can be said that so far results in career concerns models have been obtained under common knowledge of the agent's probability distribution ω_i over the set of principals $\mathcal{P}$. Under these circumstances the equilibrium effort is easily calculable since there is no unknown force that influences the marginal incentives of the agents. In a model without observation error on output, the talent would be fully revealed after the first period.

In the following I will look at a situation in which principals do not perfectly know the provenance of the agents implicit incentives.

Consider two labor markets that pay differently for the same individual talent θ. One might think of differences in real terms of payments or in productivity levels (e. g. Kremer, 1993), which imply a mapping from underlying talent to productivity that is heterogeneous across markets. The agents hold a certain belief of ending up in either one of the markets. Furthermore, I assume that there is a "domestic" market that pays the marginal product $k^d \cdot E(\theta|y)$ and a "foreign" market which pays $k^f \cdot E(\theta|y)$ with $k^f > k^d$ after observing the output y.

Initially, I consider two types of agents that are differing in the probability that they assign to ending up in either of the two markets. These probabilities and the productivity differences result in expected future wages of $k^0 \cdot E(y_2|y_1)$ for the lower type, i. e. the agent that attaches a higher probability to the low productivity market. The high type expects a future wage of $k^1 \cdot E(y_2|y_1)$.

While both the markets and the agent have no information about the talent apart from the joint density $f(\theta, y|a)$, the agent knows his probability distribution ω_i over the set of principals, but the market has only a prior belief about the type[59], $Pr(\text{type } 1) = q$ and $Pr(\text{type } 0) = 1 - q$.

The further setup in this section is similar to Dewatripont et al. (1999a; 1999b). The agent exerts effort a_t in period t and the inference is based on $f(\theta, y|a)$, the known joint density function of talent θ and output y conditional on the effort a. The cost of effort is denoted $c(a)$.

[59]To be precise, it has to be said that under this interpretation the principal not only infers the talent from the observed output, but also from the fact that the agent is now present in his market. Therefore, we need to have a situation in mind in which no agent excludes any market. This would provide further information to the principal of the excluded market about the population of agents he faces. Strictly, I need to interpret the population distribution which is reflected by the probability q of facing a high type to be specific to the principal.

The model under certainty about the type The heterogeneity of the agents is reflected by the factor k^t. They choose the effort a in order to maximise the expected payoff

$$\max_a k^t \cdot E_y[E(\theta|y,a^*)] - c(a),$$

for which the first order conditions are

$$k^t \iint \theta f(\theta,y|a^*) \frac{\hat{f}_a(y|a^*)}{\hat{f}(y|a^*)}\, dy\, d\theta = c_a(a^*),$$

where $\hat{f}(y|a) = \int f(\theta,y|a)\, d\theta$.
Consequently, the marginal incentives as derived in Dewatripont et al. (1999a, p. 186) will change to the following expressions

$$k^t \cdot Cov\left(\theta, \frac{\hat{f}_a}{\hat{f}}\right) = c_a(a^*).$$

These first order conditions determine the level of effort exerted in the case where the magnitude of career concerns, reflected by k^t, is common knowledge.

The model under uncertainty about the type Even without common knowledge of the magnitude of career concerns, in equilibrium the principals' belief about the equilibrium effort of the two types of agents will be correct. However, as the type is not fully revealed by observation of the output, the principals can only remunerate on the basis of an inference that uses an expected effort level

$$\hat{a}^* = Pr(\text{type } k^1|y) \cdot a^{1*} + Pr(\text{type } k^0|y) \cdot a^{0*}.$$

The principals use the information about the output to update their prior belief about the type that they face according to Bayes' rule

$$Pr(\text{type } k|y) = \frac{Pr(y|\text{type } k) \cdot Pr(\text{type } k)}{Pr(y)}.$$

With the new level of uncertainty the objective function of the agent of type k becomes

$$\max_a k \cdot E_y\left[E_k\left[E(\theta|y,\hat{a}^*)\right]\right] - c(a), \tag{3.1}$$

which becomes

$$\int_y k\left[\sum_{k=k^0}^{k^1} Pr(k|y)\left(\int_\theta \theta \frac{f(\theta,y|a^{k*})}{\hat{f}(y|a^{k*})}\, d\theta\right)\right] \hat{f}(y|a)\, dy - c(a). \tag{3.2}$$

The new element in this expression is the sum which reflects the expectation with respect to the equilibrium effort level of the types.

For the subsequent results and for simplicity, we need to formulate two assumptions. Firstly, it is assumed that the expression in brackets

$$\left[\sum_{k=k^0}^{k^1} Pr(k|y)\left(\int_\theta \theta \frac{f(\theta,y|a^{k*})}{\hat{f}(y|a^{k*})}\,d\theta\right)\right]$$

is equal to

$$\int_\theta \theta \frac{f(\theta,y|\hat{a}^*)}{\hat{f}(y|\hat{a}^*)}\,d\theta \tag{3.3}$$

with $\hat{a}^* = \sum_{k=k^0}^{k^1} Pr(k|y)a^{k*}$. That is, it is assumed to be irrelevant whether the principals pays according to the marginal product computed with the expected effort level, or whether they pay according to the expected marginal product on the basis of each type's equilibrium effort level. Of course, this assumption is only correct if the relationship between effort level and inferred marginal product is linear, which is not necessarily the case. However, this light assumption makes computations tractable in the following.

Secondly, both the posterior density functions $\frac{f(\theta,y|a)}{\hat{f}(y|a)}$ and $\hat{f}(y|a)$ have the strict Monotone Likelihood Ratio Property (MLRP) with respect to effort a. This implies that the posterior distribution conditional on effort a^0 first-order stochastically dominates the distribution conditional on a^1 if $a^1 > a^0$ (See Milgrom, 1981, Proposition 2). For the model at hand this implies that the inferred talent decreases with the exerted effort as

$$\int_\theta f(\theta)\frac{f(\theta,y|a^0)}{\hat{f}(y|a^0)}\,d\theta > \int_\theta f(\theta)\frac{f(\theta,y|a^1)}{\hat{f}(y|a^1)}\,d\theta.$$

These assumptions are not strong either, since it is plausible that an increased effort increases output, hence the strict MLRP on $\hat{f}(y|a)$. For a given level of output, say the observed level, an increase in effort will imply a reduced importance of talent, hence the lower inference with respect to talent.

Proposition 1 *In the two-period Holmström model of career concerns the uncertainty about the magnitude of the agents' career concerns reduces the equilibrium effort of all types of agents.*

Proof. To see the change in the analysis to the case of common knowledge, compare the rewritten integral

$$\int_y k\left(Pr(\text{type } k^0|y)\cdot E(\theta|y,a^{0*}) + Pr(\text{type } k^1|y)\cdot E(\theta|y,a^{1*})\right)\hat{f}(y|a)\,dy \tag{3.4}$$

with the expression of the marginal incentives in the case of certainty:

$$\int_y kE(\theta|y,a^*)\hat{f}(y|a)\,dy.$$

The main difference between these two expressions is that the observed output impacts on the expected talent in the case of certainty, while it impacts both on the expected talent and on the probability of facing a certain agent of type k in the other case.

Note that $E(\theta|y,a)$ decreases in a for a given y by the assumed strict MLRP. Therefore, in our case of two types and for an observed output y, we get

$$E(\theta|y,a^{0*}) > E(\theta|y,a^{1*}).$$

Furthermore, we can state that $\sum_{k=k^0}^{k^1} Pr(k|y) = 1$. Therefore,

$$\frac{\partial}{\partial y} Pr(\text{type } 0|y) = -\frac{\partial}{\partial y} Pr(\text{type } 1|y)$$

A small change in exerted effort leads to a higher expected output as the density function now gives relatively more weight to higher output levels by the MLRP. In order to make a statement about the implications for the whole integral as in equation 3.4 the following part of the integrand has to be analysed for the situation of y increasing marginally.

$$Pr(\text{type } 0|y) \cdot E(\theta|y,a^{0*}) + Pr(\text{type } 1|y) \cdot E(\theta|y,a^{1*}).$$

The derivative of this integrand with respect to y gives

$$\begin{aligned}
\frac{\partial}{\partial y}() &= \frac{\partial}{\partial y} Pr(0|y) E(\theta|y,a^{0*}) + Pr(0|y)\frac{\partial}{\partial y} E(\theta|y,a^{0*}) + \\
&\quad \frac{\partial}{\partial y} Pr(1|y) E(\theta|y,a^{1*}) + Pr(1|y)\frac{\partial}{\partial y} E(\theta|y,a^{1*}) \\
&= \underbrace{\frac{\partial}{\partial y} Pr(0|y)}_{<0} \underbrace{\left[E(\theta|y,a^{0*}) - E(\theta|y,a^{1*})\right]}_{>0} + \frac{\partial}{\partial y} E(\theta|y,\hat{a}^*) \qquad (3.5)
\end{aligned}$$

By our assumption in equation (3.3), the last term here reflects the happenings under certainty, a higher output for a given effort leading to an increased expectation of talent. Integrated over all possible y this will lead to the positive marginal benefit of exerting effort.

Due to the uncertainty with respect to the types, there is an additional negative term which results from the change in the probability of being associated with one of the two types multiplied by the impact on the expected talent. It follows that the overall integrand is less strongly increasing in the output level, implying that the integral as a whole is reduced.

Put differently, the marginal incentive that in the case of certainty we used to express in terms of $Cov\left(\theta, \frac{\hat{f}_a}{f}\right)$ will now be lower because the link between θ and y has been weakened. The first-order conditions will determine an optimal level of effort that is lower for both types of agents involved. To see the proof for finitely many types of agents please refer to the appendix C.1. ■

3.2.1 Is it possible to have a complete offset of incentives?

The previous section could prove that the implicit incentives are muted under uncertainty. It would be interesting to know, whether this uncertainty could make all incentives vanish?

In order to answer this question in a rather heuristical fashion, we need to look again at the expression in equation 3.5. In order to offset the complete impact of a change of output on the expected talent $\frac{\partial}{\partial y}E(\theta|y, \hat{a}^*)$, we can see that both the impact on the probability $\frac{\partial}{\partial y}Pr(0|y)$ and the difference in evaluation the expected talent depending on the effort have to be large. Note that the maximum value of $\frac{\partial}{\partial y}Pr(0|y)$ is 1.

For a certain value of output y, a complete offset is plausible. If the probability $Pr(0|y)$ changed by 0.9 and the according difference $E(\theta|y, a^{0*}) - E(\theta|y, a^{1*})$ is high while the impact on the expected talent $\frac{\partial}{\partial y}E(\theta|y, \hat{a}^*)$ is low, the increase of output would (locally) do harm and lead to a loss in reputation.

However, as we are discussing the integrand of an integral, an offset implies that a loss occurs or no reputational benefits obtains when integrating over the whole range of output levels. The range is existing due to the uncertainty of the agent regarding his talent. So there is no specific level of output foreseeable.

While the whole range of $E(\theta|y, \hat{a}^*)$ has an absolute value of $\max\theta - \min\theta$ assuming that the density function has a large enough support, the expression $Pr(0|y)$ changes only by 1 over the whole spectrum of y. Multiplied with the difference in expected talent due to heterogeneous efforts, one might roughly get a value of even this difference. From there it seems plausible that a certain level of incentives always remains and that complete offset is not possible. Our example in the following will illustrate this.

3.2.2 A simple example

While the previous section treated the very general case, in the following I will provide a more tangible example in which I will be able to solve for the equilibrium effort level and show explicitly the muting of the incentives.

I will consider the production function $y = \theta + a$ where talent and effort are perfect substitutes. The agent's talent is distributed uniformly on $[0, 1]$, a value which is unknown to both the principals and the agents ex ante.

The agent can be of two different types. The principal attaches a probability of q on the agent being of type k^1 who is expecting $k^1 \cdot E(\theta|y_1)$ as future wage. With probability $1 - q$ the agent expects $k^0 \cdot E(\theta|y_1)$ as future wage and is of type k^0. Without loss of generality we can normalise k^0 to 1 and k^1 to $1 + c$.

In order to keep the model tractable, I consider a two-period model and a disutility of effort of $g(a) = \frac{1}{2}a^2$. The timing is like in the original model of Holmström (1999). In the notation, superscripts are referring to types, subscripts to time periods.

Model under certainty Let me briefly describe the situation without uncertainty. We know that in the second period the lack of career concerns leads to no effort exertion $a_2 = 0$. Therefore the agent solves

$$\max_{a_1} W = -g(a_1) + \delta \left[k^1 \cdot E(\theta|y_1)\right],$$

where $E(\theta|y_1) = E(y_1) - a_1^*$. From the first order conditions it follows

$$g'(a_1) = \tilde{a}_1^{1*} = \delta \cdot k^1 = (1+c)\delta$$

By similar reasoning the agent not facing the additional principal will end up with $g'(a_1) = \tilde{a}_1^{0*} = \delta \cdot k^0 = \delta$.

Model with uncertainty Following the previous paragraph we can state that the expected effort level is $a_1^* = q \cdot a_1^{1*} + (1-q) \cdot a_1^{0*}$. However, this is not the full specification of the equilibrium effort because the principals observe y_1 before making an inference about the agent's talent and paying the second period wage. The observation of y_1 allows for a better inference of the true talent and enables an updating of the prior probabilities. In our case of the uniform distribution, for some y_1 the talent can be inferred with certainty if only one type can produce this amount of output given the equilibrium effort levels of the two types. Therefore, for some y_1 it is not optimal to pay the wage according to an inference based on the expected effort as formulated above. Depending on the region of y_1 the principals then have three ways of calculating the inferred effort .

- If $y_1 < 0 + a_1^{1*} \equiv \underline{y}$, then $Pr(\text{type } 0|y_1) = 1$ and the principals would rationally be willing to pay up to $y_1 - a_1^{0*} > y_1 - q \cdot a_1^{1*} - (1-q) \cdot a_1^{0*}$. So for low enough y_1 the agent of type 0 benefits from being identified as exerting less effort, which then increases his inferred talent.

- If $y_1 > 1 + a_1^{0*} \equiv \bar{y}$ then $Pr(\text{type } 1|y_1) = 1$ and the principals would rationally be willing to pay only up to $y_1 - a_1^{1*} < y_1 - q \cdot a_1^{1*} - (1-q) \cdot a_1^{0*}$. For high y_1 the agent of type 1 is identified as a more 'concerned' type exerting higher effort, leading to a lower inferred talent.

- For $0 + a_1^{1*} < y_1 < 1 + a_1^{0*}$ the observation of the output does not convey any information about the agent's type, therefore the principals would pay $y_1 - q \cdot a_1^{1*} - (1-q) \cdot a_1^{0*}$ according to the expected effort level. Note that in this case the inferred talent is biased. Some of the effort of type 1 is interpreted as talent and some of the talent of type 0 is interpreted as effort.

This endogenous partition itself has an impact on the equilibrium effort of the agent because it influences the probability of falling in one of the two categories.

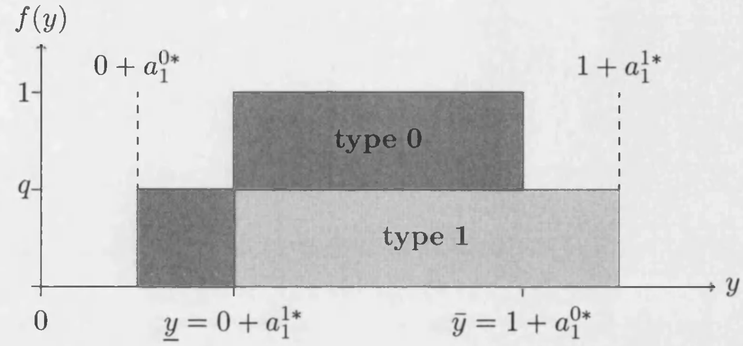

Figure 3.1: The output distribution for the two types.

Proposition 2 *In the two-period career concerns model with uniform talent distribution the equilibrium effort levels are given by*

$$\begin{aligned} a_1^{1*} &= \delta(1+c)\left(1+(1-q)\frac{c}{2q-1-\frac{1}{\delta}-(1-q)c}\right), \\ a_1^{0*} &= \delta\left(1+\frac{qc}{2q-1-\frac{1}{\delta}-(1-q)c}\right). \end{aligned}$$

Both levels are strictly smaller than their respective counterpart under certainty.

Proof. The agent's maximisation problem becomes

$$\begin{aligned} \max_{a_1} W^1 &= -g(a_1) + \\ & \delta(1+c)\left((y_1 - a_1^{1*}) \cdot Pr(y_1 > \bar{y}) + (y_1 - q \cdot a_1^{1*} - (1-q) \cdot a_1^{0*}) \cdot Pr(y_1 < \bar{y})\right), \end{aligned}$$

where

$$Pr(y_1 > \bar{y}) = Pr(\theta + a_1 > 1 + a_1^{0*}) = Pr(\theta > 1 + a_1^{0*} - a_1) = 1 - (1 + a_1^{0*} - a_1) = -a_1^{0*} + a_1.$$

Hence $\frac{\partial Pr(y_1 > \bar{y})}{\partial a_1} = 1$ and $\frac{\partial Pr(y_1 < \bar{y})}{\partial a_1} = -1$.

The first order condition therefore turns out to be

$$-g'(a_1) + \delta(1+c)\left(\left[Pr(y_1 > \bar{y}) + (y_1 - a_1^{0*})\right] + \left[-(y_1 - q a_1^{1*} - (1-q)a_1^{0*}) + Pr(y_1 < \bar{y})\right]\right) = 0$$

which can be solved for $a_1^{1*} = \delta(1+c)\left(1 + (1-q)(a_1^{0*} - a_1^{1*})\right)$. The same reasoning for type 0 gives $a_1^{0*} = \delta(1 + q(a_1^{0*} - a_1^{1*}))$. Solving for the equilibrium levels of effort proofs the claim.

Furthermore, it can be shown that

$$a_1^{1*} = \delta(1+c)\left(1+(1-q)\frac{c}{2q-1-\frac{1}{\delta}-(1-q)c}\right) \quad < \quad \delta(1+c),$$
$$a_1^{0*} = \delta\left(1+\frac{pc}{2q-1-\frac{1}{\delta}-(1-q)c}\right) \quad < \quad \delta.$$

■

To see more clearly what happens consider the case of $\delta = 1$, such that

$$\begin{aligned} a_1^{1*} &= 1+\frac{c}{2+c} \\ a_1^{0*} &= 1-\frac{q}{1-q}\frac{c}{2+c} \end{aligned}$$

Now the implications of proposition 1 can be shown easily. The incentives are muted for both types of agents and lead to lower effort levels.

This result suggests that in spite of the presence of an extra motivation of one type of agents, the incentives are muted due to the uncertainty. In particular, the type with higher career concerns will have muted incentives because at a certain level of output he will be revealed to be of this type, allowing for a perfect revelation of his talent. This is in contrast with a positively biased belief about his talent in the case where the principal/market can not infer which type of agent he faces. Increasing effort means for the motivated agent to decrease the probability of having this bias working for him.

For the less concerned agent this argument works in the other direction, but impacts his equilibrium effort in the same direction. For low enough outputs he is known to be of his type, which allows for a perfect revelation of his true talent. A high effort, however, increases the likelihood of facing the negative bias in the inferred talent, therefore reducing the incentives to exert effort. Interestingly, the incentives are now lower than without the presence of motivated types, in other words efficiency is reduced due to the uncertainty. Vice versa, it can be said that the possible presence of certain types of agents and the principals limited information about this decreases the effort of others.

Note that in Holmström (1999) the uncertainty about the talent is a driving factor of effort exertion because this way there remains scope to influence the principals beliefs. Since the updating reduces incentives over time, Holmström introduces a model with time-varying talent so that the result of decreasing and vanishing effort exertion is altered. In contrast, here the uncertainty about the origin of incentives mutes the effort.

3.2.3 An observation on the wage bias

While the analysis so far focussed on the exerted effort, it is worth directing attention for a moment towards the wage that is paid to the agents. In the framework of the simple example in the previous section, it can be seen that in the case of full revelation of the type through output ($y_1 < \underline{y}$ or $y_1 > \bar{y}$) the wage is appropriate for every agent since the principal knows the exact talent of the agent.

However, in the case of unchanged uncertainty, the wage payment will be made according to the expected effort level, thereby biasing the wage for the agent with low career concerns downwards and vice versa ($y_1 - q \cdot a_1^{1*} - (1-q) \cdot a_1^{0*}$). The existence of higher types in the dimension of the added heterogeneity therefore imposes an externality on the lower types.

This observation hints towards a general bias against agents with a higher cost of effort, a lower patience or a lower 'drive'. In fact, considering optimism as the dimension of added heterogeneity, this observation might help explain why a stable share of society display irrationally optimistic beliefs about the future.

Assume that c reflects the bias in expectation of future economic outlooks of the optimists (type 1). The other agents do have rational expectations (type 0). This leads to the following wages w and welfare W for the two types

$$
\begin{aligned}
w_2^{1*} &= \theta + (1-q)c\delta, \\
w_2^{0*} &= \theta - qc\delta, \\
W^1 &= -\frac{1}{2}(\delta(1+c))^2 + \delta(\theta + (1-q)c\delta) \\
W^0 &= -\frac{1}{2}\delta^2 + \delta(\theta - qc\delta)
\end{aligned}
$$

It can be shown that $W^1 - W^0 < 0$, i. e. that the optimists' bias does not lead to a preferable situation compared to the realists. However, since the two types should get the same wage, the difference in welfare should be the full cost of the additional effort $\Delta = \frac{1}{2}\delta^2 c(2+c)$. Instead, due to the bias it is reduced to $\frac{1}{2}\delta^2 c(2+c) - c\delta^2 = \Delta - c\delta^2 = \frac{1}{2}\delta^2 c^2$.

It can be seen that this bias mitigates the consequences of the bias in the expectations of future outcomes of the optimists. As they do not feel the full strength of the consequences it might not be obvious that keeping this bias is welfare reducing.

In an empirical analysis about the existence of optimism in society Puri and Robinson (2007) find out that many people are relatively optimistic with respect to their life-expectancy and future economic situation. Furthermore, their analysis shows that socioeconomic choices of optimists differ significantly across different areas like retirement, investment and saving decisions, remarriage etc.

One particular finding is that optimists use to work more hours. Referring to a common interpretation of worked hours as spent effort, we can relate this finding to the result of the presented model. The fact that optimists assume a brighter future implies that for them their present actions echo louder in the future since there is more 'at stake'. Hence, for them there is a good reason to put more effort than others, especially since they will not feel the full consequence of their incorrect beliefs.

3.3 Additional heterogeneity in the infinite periods model

When additional heterogeneity is considered in a multi-period setting, both the 'backward attribution' and the 'forward attribution' play a role in determining the incentives at a given point in time. In the three-period model of Kőszegi and Li (2008) the first period is characterised by a forward attribution that is higher than the backward attribution, leading to a higher effort level than under certainty. This raises the question about the general relationship between these two effects in a multi-period model.

The interplay between the two effects will be determined by the information available and hence by the updating process as sketched in figure 3.2. At the same time, the resulting effort influences the updating process, making the process depending on both past and future actions.

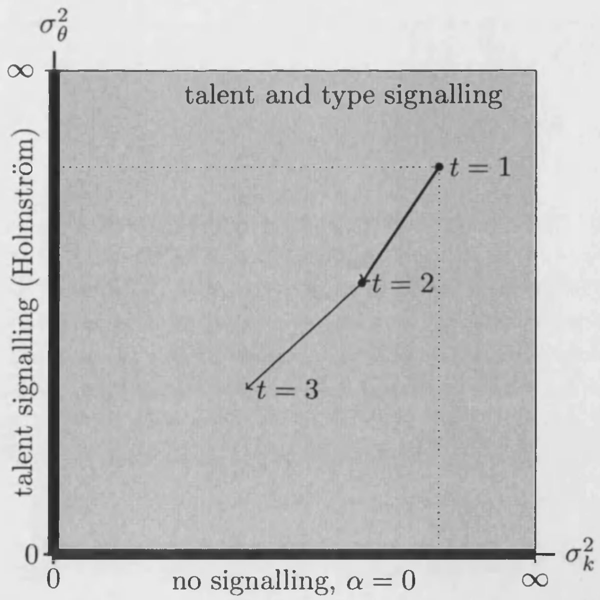

Figure 3.2: Talent and type updating over time.

3.3.1 The model setup

Consider an infinite period model at time t. In order to obtain a tractable reputation model, but still illustrate the main trade-off, I will consider an agent that is concerned about the next period's wage only. This implies that only his immediate reputation is relevant to him. Furthermore, in order to obtain a tractable model of the updating process I will specify the talent distribution to be normal.

The setup follows mainly Kőszegi and Li (2008). The output is given by $y_t = \theta + a_t + \varepsilon_t$, where $\theta \sim N(0, \sigma_\theta^2)$ and $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$. The heterogeneity is introduced through the factor $k \sim N(\mu_k, \sigma_k^2)$, which reflects the magnitude of the career concerns like in the two-period model before.

I consider equilibria that are linear rational expectation equilibria with $a_t = \alpha_t \cdot k$. There is a quadratic cost on effort exertion $c(a_t) = \frac{1}{2}a_t^2$. While both agents and principals are not informed about the talent, the type k realisation is known to the agent but not to the principal.

Due to perfect competition for the agent's labor force, his wage in each period is equal to the expectations of the principal regarding his output $E(y_{t+1}|y_t)$. The priors concerning θ and k are assumed to incorporate the full information available at time t. Abstracting from discounting, the agent's maximisation problem looks as follows

$$\max_{a_t} -c(a_t) + kE_{y_t}E(y_{t+1}|y_t).$$

The corresponding first order conditions imply that $a_t^* = k\frac{\partial E_{y_t}E(y_{t+1}|y_t)}{\partial a_t}$. From the linear rational expectation equilibrium it follows that $\alpha_t = \frac{\partial E_{y_t} w_{t+1}}{\partial y_t}$, since due to the additivity of the output, the derivative with respect to a_t is the same as with respect to output y. In order to calculate this partial derivative, note that $w_{t+1} = E(y_{t+1}|y_t) = E(\theta + a_{t+1}|y_t) = E(E(\theta|k, y_t)|y_t) + E(a_{t+1}|y_t)$, where $E(a_{t+1}|y_t) = \alpha_{t+1}E(k|y_t)$. Overall, the expression for the wage becomes:

$$w_{t+1}E(y_{t+1}|y_t) = E(E(\theta|k, y_t)|y_t) + \alpha_{t+1}E(k|y_t).$$

Proposition 3 *In the infinite-period normal career concern model with a one-period reputation horizon, the equilibrium effort factor α_t can be stated implicitly as a function of α_{t+1} as follows.*

$$\frac{\partial E_{y_t} w_{t+1}}{\partial y_t} = \alpha_t = \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_\varepsilon^2} - \underbrace{\frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_\varepsilon^2}\alpha_t \frac{\partial E(k|y_t)}{\partial y_t}}_{\text{backward attribution}} + \underbrace{\alpha_{t+1}\frac{\partial E(k|y_t)}{\partial y_t}}_{\text{forward attribution}} \tag{3.6}$$

$$= \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_\varepsilon^2} - \frac{\alpha_t^2\sigma_k^2\sigma_\theta^2}{(\sigma_\theta^2 + \alpha_t^2\sigma_k^2 + \sigma_\varepsilon^2)(\sigma_\theta^2 + \sigma_\varepsilon^2)} + \frac{\alpha_t\sigma_k^2\alpha_{t+1}}{\sigma_\theta^2 + \alpha_t^2\sigma_k^2 + \sigma_\varepsilon^2} \tag{3.7}$$

$$= \frac{\sigma_\theta^2 + \alpha_t\alpha_{t+1}\sigma_k^2}{\sigma_\theta^2 + \alpha_t^2\sigma_k^2 + \sigma_\varepsilon^2} \tag{3.8}$$

Proof. See appendix C.2. ■

3.3.2 Characteristics of the effort level in period t

Proposition 4 *The effort level $\alpha_t{\cdot}k$ is 0 only if k is 0. As long as we have uncertainty about the talent θ α_t is never equal to 0.*

Proof. Using $\alpha_t = 0$ in equation 3.8, we get

$$0 = \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_\varepsilon^2},$$

a contradiction as long as we have uncertainty about the talent θ. ∎

Proposition 5 *α_t can only be negative if $\frac{\partial E(k|y_t)}{\partial y_t} < 0$, i. e. if the type update increases with lower output, which is equivalent to a negative forward attribution.*

Proof. Requiring the right hand side of equation 3.7 to be negative, we observe that the backward attribution is always negative, even with a negative α_t, but never outweighs the first part from the talent signalling. To reach a level below 0, the forward attribution has to be negative. Transforming the negativity condition, we get

$$\frac{\sigma_\theta^2}{\sigma_k^2} < |\alpha_t|\alpha_{t+1},$$

implying that α_{t+1} needs to be positive for this condition to be satisfied, since the fraction on the left hand side is always positive. Hence, the only way to obtain a negative forward attribution is $\frac{\partial E(k|y_t)}{\partial y_t} < 0$. ∎

Proposition 6 *Given initial uncertainty, there does not exist an equilibrium with $\alpha_t \leq 0$.*

Proof. Economically, it does not make sense to consider a situation in which $\frac{\partial E(k|y_t)}{\partial y_t} < 0$. An agent that would want to signal his talent through a high output, would not destroy output to signal drive. Hence, assuring that $\frac{\partial E(k|y_t)}{\partial y_t} < 0$ does not make economical sense, the first order conditions cannot be met with any $\alpha_t \leq 0$. Therefore, any equilibrium must exhibit $\alpha_t > 0$. ∎

3.3.3 Relating effort levels in period *t* and *t+1*

In the career concerns model as I propose it here, there is still the important notion of uncertainty that drives the effort exertion of the agent. It can be seen that $\alpha_t = 0$ if and only if both $\sigma_\theta^2 = 0$ and $\sigma_k^2 = 0$. Recalling that any equilibrium in which uncertainty exists exhibits a positive α, it is clear that the dynamic path of effort exertion is tightly connected with the information flow and updating of talent and type estimates.

Proposition 7 *With every observation of realised output, the talent and type estimates become more precise. Therefore, as $t \to \infty$, all uncertainty vanishes and $\alpha_t = 0$.*

Proof. In equations 3.10 and 3.12 (page 110) it can be seen that both $\bar{\sigma}_k^2 < \sigma_k^2$ and $\bar{\sigma}_\theta^2 < \sigma_\theta^2$, where $\bar{\sigma}$ reflects the new variance after updating with one period's output information. The updating can only result in an unchanged variance if the variances are 0 already. Hence, over time the uncertainty must vanish, implying that $\alpha_t = 0$. ■

Lemma 8 *It will never be that $\alpha_t = \alpha_{t+1}$ when $\alpha_{t+2} \leq \alpha_t$.*

Proof. Suppose yes and $\alpha_t = \alpha_{t+1}$, then equation 3.8 implies

$$\frac{\bar{\sigma}_\theta^2 + \alpha_t \alpha_{t+2} \bar{\sigma}_k^2}{\bar{\sigma}_\theta^2 + \alpha_t^2 \bar{\sigma}_k^2 + \sigma_\varepsilon^2} = \frac{\sigma_\theta^2 + \alpha_t \alpha_t \sigma_k^2}{\sigma_\theta^2 + \alpha_t^2 \sigma_k^2 + \sigma_\varepsilon^2}.$$

Referring to equations 3.10 and 3.12 again, it can be seen that

$$\frac{\bar{\sigma}_\theta^2 + \alpha_t \alpha_{t+2} \bar{\sigma}_k^2}{\bar{\sigma}_\theta^2 + \alpha_t^2 \bar{\sigma}_k^2 + \sigma_\varepsilon^2} < \frac{\sigma_\theta^2 + \alpha_t \alpha_t \sigma_k^2}{\sigma_\theta^2 + \alpha_t^2 \sigma_k^2 + \sigma_\varepsilon^2},$$

contradicting the initial assumption. ■

Proposition 9 *For all periods t it will be true that $\alpha_t \geq \alpha_{t+1}$.*

Proof. Suppose not and $\alpha_t < \alpha_{t+1}$, implying that

$$\frac{\bar{\sigma}_\theta^2 + \alpha_{t+1} \alpha_{t+2} \bar{\sigma}_k^2}{\bar{\sigma}_\theta^2 + \alpha_{t+1}^2 \bar{\sigma}_k^2 + \sigma_\varepsilon^2} > \frac{\sigma_\theta^2 + \alpha_t \alpha_{t+1} \sigma_k^2}{\sigma_\theta^2 + \alpha_t^2 \sigma_k^2 + \sigma_\varepsilon^2}.$$

With proposition 7 this must be in particular true for a situation in which both $\alpha_t \leq \alpha_{t+1}$ and $\alpha_{t+2} \leq \alpha_{t+1}$, since α_t will go to zero at some point. Multiplying the equation by the two denominators, the condition becomes

$$\begin{aligned}\bar{\sigma}_\theta^2 \sigma_k^2 (\alpha_t^2 - \alpha_t \alpha_{t+1}) + \sigma_\theta^2 \bar{\sigma}_k^2 (\alpha_{t+1}\alpha_{t+2} - \alpha_{t+1}^2) + \bar{\sigma}_k^2 \sigma_k^2 (\alpha_t^2 \alpha_{t+1} \alpha_{t+2} - \alpha_t \alpha_{t+1}^3) > \\ \sigma_\varepsilon^2 (\sigma_\theta^2 - \bar{\sigma}_\theta^2) + \alpha_{t+1} \sigma_\varepsilon^2 (\alpha_t \sigma_k^2 - \alpha_{t+2} \bar{\sigma}_k^2),\end{aligned}$$

Using the relationship between σ_θ^2 and $\bar{\sigma}_\theta^2$ from equation 3.12, this is a contradiction since the left hand side is negative and the right hand side positive. ■

Conjecture 10 *There exists a linear equilibrium in which $\alpha_t \geq 0$ $\forall t$ and $\alpha_t \geq \alpha_{t+1}$ $\forall t$.*

The argument evolves around the relevant equation 3.8

$$\alpha_t = \frac{\sigma_{t,\theta}^2 + \alpha_t \alpha_{t+1} \sigma_{t,k}^2}{\sigma_{t,\theta}^2 + \alpha_t^2 \sigma_{t,k}^2 + \sigma_\varepsilon^2}.$$

For a sufficiently large positive constant K, consider the set

$$S = \{(\alpha_1, \alpha_2, \alpha_3, \ldots) | K \geq \alpha_1 \geq \alpha_2 \geq \alpha_3 \geq \ldots \geq 0\}$$

in $\mathbb{R}^\infty$. The map $f : S \to \mathbb{R}^\infty$ is defined componentwise, placing in t-th position the right-hand side of equation 3.8.

As opposed to earlier in this chapter, now the variances of talent and type have to be indexed by the period t. The law of motion for $\sigma^2_{t,\theta}$ is given, analogous to equation 3.12, to be

$$\sigma^2_{t+1,\theta} = \frac{\sigma^2_{t,\theta}\sigma^2_{\varepsilon}}{\sigma^2_{t,\theta} + \sigma^2_{\varepsilon}},$$

which is independent of the sequence of effort choices $\{\alpha_t\}$. The law of motion for $\sigma^2_{t,k}$ is given in equation 3.10 to be

$$\begin{aligned} \sigma^2_{t+1,k} &= \sigma^2_{t,k} - \frac{(\alpha_t \sigma^2_{t,k})^2}{\sigma^2_{t,\theta} + \alpha_t^2 \sigma^2_{t,k} + \sigma^2_{\varepsilon}} < \sigma^2_{t,k} \\ &= \frac{\sigma^2_{t,k}(\sigma^2_{t,\theta} + \sigma^2_{\varepsilon})}{\sigma^2_{t,\theta} + \alpha_t^2 \sigma^2_{t,k} + \sigma^2_{\varepsilon}}, \end{aligned}$$

which is decreasing in α_t. Both processes start with initial finite values $\sigma^2_{1,k}$ and $\sigma^2_{1,\theta}$. Through the dependence of $\sigma^2_{t+1,k}$ on α_t, all previous effort choices are relevant, leading to an expression $f_t = f_t(\alpha_1, \alpha_2, \ldots, \alpha_t, \alpha_{t+1})$.

I conjecture that it is possible to check that the map f is inward as defined in Halpern and Bergman (1968)[60] by checking that, at the boundaries of S, the map f points into the subspace that includes S. This should be assured by the following conditions which can be proven to hold:

1. One can choose K such that it is true that $f_1(K, \alpha_2) < K$ for any $\alpha_2 \leq K$.

2. For finite t, whenever $\alpha_t = \alpha_{t+1}$, it is true that $f_t(\alpha_1, \ldots, \alpha_{t+1}) > f_{t+1}(\alpha_1, \ldots, \alpha_{t+2})$ for any given $(\alpha_1, \ldots, \alpha_t) \in S^t$, where S^r denotes the space spanned by the r first dimensions of S.

3. For any finite t it is true that $f_t(\alpha_1, \ldots, \alpha_{t+1}) > 0$.

4. f is continuous.

Condition 1 can be proven to hold by looking at

$$f_1(K, \alpha_2) = \frac{\sigma^2_{1,\theta} + K\alpha_2\sigma^2_{1,k}}{\sigma^2_{1,\theta} + K^2\sigma^2_{1,k} + \sigma^2_{\varepsilon}},$$

which is always less than one for $\sigma^2_{\varepsilon} > 0$ and $\alpha_2 \leq K$. Any $K \geq 1$ will satisfy this condition.

[60]A map $F : S \rightarrow X$ is called inward if for all points $s \in S$, $F(s)$ belongs to the inward set. For a given $s \in S$, the inward set is the "union of all rays originating at s and drawn so as to pass through some other point r of S." (p. 353).

For condition 2 to hold under $\alpha_t = \alpha_{t+1}$, it can be shown that

$$\frac{\sigma^2_{t,\theta} + \alpha_t^2 \sigma^2_{t,k}}{\sigma^2_{t,\theta} + \alpha_t^2 \sigma^2_{t,k} + \sigma^2_\varepsilon} > \frac{\sigma^2_{t+1,\theta} + \alpha_t \alpha_{t+2} \sigma^2_{t+1,k}}{\sigma^2_{t+1,\theta} + \alpha_t^2 \sigma^2_{t+1,k} + \sigma^2_\varepsilon}.$$

This can be seen by multiplying out and rewriting the inequality without the expressions that cancel out as

$$\begin{aligned}&\sigma^2_{t,\theta} \cdot \alpha_t^2 \sigma^2_{t+1,k} + \sigma^2_{t,\theta} \cdot \sigma^2_\varepsilon + \alpha_t^2 \sigma^2_{t,k} \cdot \alpha_t^2 \sigma^2_{t+1,k} + \alpha_t^2 \sigma^2_{t,k} \cdot \sigma^2_\varepsilon > \\ &\sigma^2_{t,\theta} \cdot \alpha_t \alpha_{t+2} \sigma^2_{t+1,k} + \sigma^2_{t+1,\theta} \cdot \sigma^2_\varepsilon + \alpha_t^2 \sigma^2_{t,k} \cdot \alpha_t \alpha_{t+2} \sigma^2_{t+1,k} + \alpha_t \alpha_{t+2} \sigma^2_{t+1,k} \cdot \sigma^2_\varepsilon.\end{aligned}$$

Combining the facts that $\sigma^2_{t,\theta} > \sigma^2_{t+1,\theta}$, $\sigma^2_{t,k} > \sigma^2_{t+1,k}$, and that $\alpha_t^2 > \alpha_t \alpha_{t+2}$, this inequality will hold for any $\sigma^2_{t+1,k}$ and – since this is the only point where earlier α_r, $r < t$ come in – therefore also for any $(\alpha_1, \ldots, \alpha_t) \in S^t$. Note that this argument includes boundary locations in further dimensions.

For condition 3, note that $f_t(\alpha_1, \ldots, \alpha_{t+1})$ is positive for finite t and converges to 0 only as $t \to \infty$, due to the trajectory of $\sigma^2_{t,\theta}$.

Condition 4 is satisfied because every f_t depends continuously on $\alpha_1, \ldots, \alpha_{t+1}$. The definition of the product topology implies therefore that f is continuous as well.

It is conjectured that the first three conditions make sure iteratively that the mapping f is inward. Combined with the forth condition, by the Halpern-Bergman Theorem (Halpern and Bergman, 1968, Lemma 3.1), the map f has a fixed point. This fixed point is a linear rational expectations equilibrium.

One can think about the uncertainty at time t as defining a kind of potential for a possible effort exertion over time. In the original model in Holmström (1999) this potential is used fully over time at a rate that is determined by the speed of learning and by the horizon of the agent's optimisation. Introducing shocks to the true talent allowed to nurture the uncertainty and therefore inhibits the depletion of this potential.

In a model with additional heterogeneity like Kőszegi and Li (2008) or mine, the fact that there is a new parameter that impacts the effort level does not change the notion of potential or the fact that at some t high enough, it must be that $\alpha_t = 0$. Here I showed that the effort level as well decreases steadily over time. However, since this new parameter influences the current effort level and is influenced by past and future effort levels, we observe a more complex path of effort exertion. The interplay between backward attribution and forward attribution becomes a new dimension in the dynamic characteristic of the model.

3.3.4 Bootstrapping

In order to fill figure 3.2 with life, we can get the first intuition about the dynamics by considering the special case of certainty with respect to the agent's talent. For this, consider the equation that illustrates the two attribution effects

$$\frac{\partial w_{t+1}}{\partial y_t} = \alpha_t = \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_\varepsilon^2} - \underbrace{\frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_\varepsilon^2} \alpha_t \frac{\partial E(k|y_t)}{\partial y_t}}_{\text{backward attribution}} + \underbrace{\alpha_{t+1} \frac{\partial E(k|y_t)}{\partial y_t}}_{\text{forward attribution}} \quad . \tag{3.6}$$

If the type is fully known, then $\frac{\partial E(k|y_t)}{\partial y_t}$ is zero, and the only motivation for exerting effort is to get a high talent estimation from the part of the principal. This refers back to the original Holmström-model, as reflected on the left edge of the graph.

On the other hand, if the talent is fully known and $\sigma_\theta^2 = 0$, then the two first expressions are zero, and only the forward attribution might give an incentive to exert effort if the type is unknown to the principal. In this case, the implicit equation in 3.6 results in the relationship between α_t and α_{t+1} of

$$\alpha_t^2 + \frac{\sigma_\varepsilon^2}{\sigma_k^2} = \alpha_{t+1}$$
$$\alpha_t = \sqrt{\alpha_{t+1} - \frac{\sigma_\varepsilon^2}{\sigma_k^2}}$$

Proposition 11 *If $\sigma_\theta^2 = 0$ and $4\sigma_\varepsilon^2 < \sigma_k^2$ there exists a range*

$$\alpha_t \in \left[\frac{1}{2}\left(1 - \sqrt{1 - \frac{4\sigma_\varepsilon^2}{\sigma_k^2}}\right), \frac{1}{2}\left(1 + \sqrt{1 - \frac{4\sigma_\varepsilon^2}{\sigma_k^2}}\right)\right]$$

for which $\alpha_t > \alpha_{t+1}$, i. e. the current effort level is higher than the next period's effort level due to the forward attribution.

This is what Kőszegi and Li (2008) refer to as "bootstrapping" of the incentives, since the mere signalling of drive in the next period allows the current periods effort to be even higher. It therefore propagates backward in time. The interval, however, depends on the information available on the type and the output error. The condition $4\sigma_\varepsilon^2 < \sigma_k^2$ implies that the output observation has to be precise enough to allow for a meaningful update on the type. As the information increases over time, the variance σ_k^2 decreases, reducing the channel for the bootstrapping and the difference between α_t and α_{t+1}. If the initiating implicit incentive is too far in the future, the channel might vanish completely, resulting in insignificant effort levels.

To finish this consideration the question regarding the initiating incentive has to be posed. If the talent is fully revealed ($\sigma_\theta^2 = 0$) the only possible incentive is the forward attribution that signals a high responsiveness to incentives. By the finiteness of the horizon and the knowledge that at some point $\alpha = 0$ irrespective of the type, the forward attribution cannot sustain a positive level of effort by itself. Hence, it is interesting to see that there is a way to have incentives propagate backwards in time through this forward attribution, but it is as well clear that they themselves cannot represent a starting point for such a propagation.

However, in a richer setting with information asymmetries between principals it might be conceivable that at a point in time there will be a "true" implicit incentive from talent uncertainty due to, say, a job change, which then allows for the backward propagation. A single stochastic shock in form of a change in the job description or in the environment might also be a realistic situation that could be at the root of a forward propagation. Of course, one could incorporate explicit incentives and would find that a forward attribution takes place.

In this model, there is no such starting point in the future, leaving $\alpha = 0$ when $\sigma_\theta^2 = 0$ as depicted at the bottom edge of the graph in figure 3.2.

3.3.5 Attribution dynamics

So far, I discussed the two regions along the axes of figure 3.2. In the following the most important part of the graph, the updating process with twofold uncertainty, will be analysed and the effort dynamics be put into relation with the model under certainty.

The path of effort exertion, which I suggest to be the virtual third dimension in the graph, is changed depending on whether the backward attribution is higher or lower than the forward attribution. For example, in the two-period model of Proposition 1 on page 89, the always negative backward attribution and the – due to the end of the game – non-existing forward attribution resulted in a strictly lower effort level for all types of agents. This illustrates that both the intertemporal effort allocation as well as the total effort exertion can ultimately be influenced by the timing of the updating of information.

Comparing the path of effort exertion in the model presented here and in the original model by Holmström (1999)[61], equation 3.6 shows that the sum of forward and backward attribution determines whether the effort in period t is higher than in the model without heterogeneity. Let $\hat{\alpha}$ denote effort levels under the same conditions in the original career concerns model.

Proposition 12 *The level of effort is higher in the model with heterogeneity* ($\alpha_t > \hat{\alpha}_t$) *if forward attribution is higher than backward attribution* ($\alpha_{t+1} \geq \alpha_t \frac{\sigma_\theta^2}{\sigma_\theta^2+\sigma_\varepsilon^2}$) *and vice versa.*

Proof. Follows immediately from equation 3.6. ■

Since $\alpha_{t+1} \leq \alpha_t$, the forward attribution diminishes continuously and goes to 0 over time. This fact leads to the conclusion that in the beginning of the career it is more likely to see an increased effort level due to signalling in another dimension than talent. This is in line with the intuition that Kőszegi and Li (2008) develop in their

[61] I compare to the equivalent with a one-period horizon, where $\alpha_t = \frac{\sigma_\theta^2}{\sigma_\theta^2+\sigma_\varepsilon^2}$.

three-period model, where the first period exhibits a higher forward attribution than backward attribution while the second period naturally exhibits a forward attribution of 0.

Proposition 13 *While the change over time in σ_θ^2 is independent of the effort level, the updating process on the type reflected by σ_k^2 is accelerated for higher effort levels.*

Proof. The variance of talent $\bar{\sigma}_\theta^2 = \frac{\sigma_\theta^2\sigma_\varepsilon^2}{\sigma_\theta^2+\sigma_\varepsilon^2}$ is independent of α_t. The next period variance $\bar{\sigma}_k^2$ decreases with α_t,

$$\frac{\partial\bar{\sigma}_k^2}{\partial\alpha_t} = -\frac{2\alpha_t(\sigma_k^2)^2(\sigma_\theta^2+\sigma_\varepsilon)}{\sigma_\theta^2+\alpha_t^2\sigma_k^2+\sigma_\varepsilon^2} < 0.$$

■

Whether this last fact leads immediately to a lower effort exertion in the period after the forward attribution has been higher and has led to a higher effort level depends on the resulting changes of α_t. It cannot be ruled out yet, that the effort level increases when more type information is available.

Proposition 14 *An increase in information on talent cannot lead to an increase in effort.*

Proof. Using the Implicit Function Theorem on $F(\alpha_t, \sigma_\theta^2) = \frac{\sigma_\theta^2+\alpha_t\alpha_{t+1}\sigma_k^2}{\sigma_\theta^2+\alpha_t^2\sigma_k^2+\sigma_\varepsilon^2} - \alpha_t$ (see equation 3.8) results in $\frac{\partial F}{\partial\alpha_t} < 0$. Furthermore, $\frac{\partial F}{\partial\sigma_\theta^2} > 0$ if $\frac{\sigma_\varepsilon^2}{\sigma_k^2} > \alpha_t(\alpha_{t+1}-\alpha_t)$, which is always the case since we proved that $\alpha_t > \alpha_{t+1}$. Therefore, $\frac{\partial\alpha_t}{\partial\sigma_\theta^2} = \frac{\frac{\partial F}{\partial\sigma_\theta^2}}{\frac{\partial F}{\partial\alpha_t}} > 0$. ■

Proposition 15 *An increase in information on type leads to an increase in effort exertion only if backward attribution is higher than forward attribution.*

Proof. Using the Implicit Function Theorem again we get $\frac{\partial\alpha_t}{\partial\sigma_k^2} = \frac{\frac{\partial F}{\partial\sigma_k^2}}{\frac{\partial F}{\partial\alpha_t}} > 0$ if $\alpha_{t+1} > \alpha_t\frac{\sigma_\theta^2}{\sigma_\theta^2+\sigma_\varepsilon^2}$, the condition from proposition 12. ■

The intuition behind this result is that $\frac{\partial E(k|y_t)}{\partial y_t}$ is higher for higher σ_k^2, impacting the effort level positively only if the coefficient is positive. An increase in information affects the effort level positively if the backward attribution is higher than the forward attribution, so that the additional information on talent keeps the principal from attributing past performance to a too high extent to a high effort level.

3.3.6 Dynamics

In the following I will use the insights gathered so far in order to present a relatively complete picture of the dynamics of the effort level in a multi-period setting.

A starting point is the fact that the speed of updating for the talent is independent of the effort level. Therefore, the normal Holmström result, just with a one-period horizon, will represent the basis of the present model. Deviations from this will result from the interplay of forward and backward attribution. Those two effects are depending on both past and future effort levels since the updating of the type is depending on the effort level and the forward attribution is a function of future effort and hence future uncertainty. This future uncertainty is itself again impacted by the current effort level and so on.

Whether the effort level is above or beneath the standard result depends on the magnitude of the two attributions. By proposition 15 the dynamics in comparison with the standard model also depend on the same conditions. While proposition 14 says that the information gain from one period to another always leads to a lower effort exertion for talent signalling reasons, this effort reduction is accelerated if the effort level decreases as well with the gain in information on type. Therefore, for higher than standard effort levels there is a tendency to relatively faster effort decrease, while for lower than standard effort levels the reduction in effort due to information gain is mitigated due to the gain in information on type and its effort increasing effect.

In addition to this, proposition 13 allows to say that the mitigation from type updating is the higher the higher the effort levels are since the amount of information gained from one period to the other is increasing in the absolute level of effort. With the knowledge that in the second last period the effort level in the model with heterogeneity has to be lower than the standard model, the following picture gives a first idea of the implied dynamics.

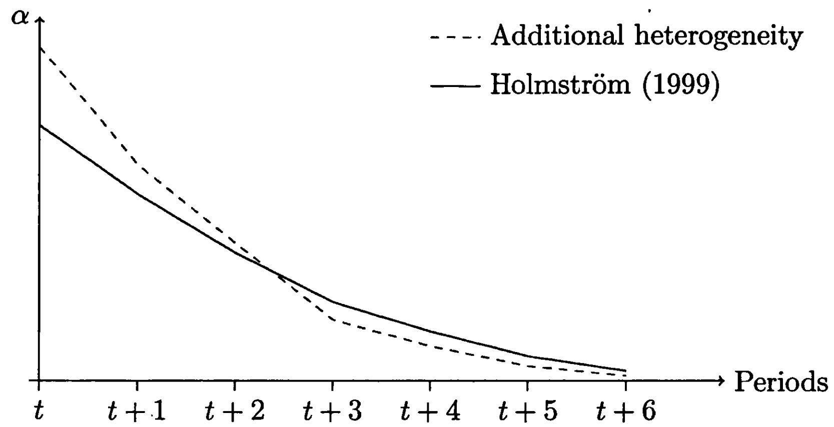

Figure 3.3: The effort dynamics with and without type uncertainty.

Qualitatively, this fits the results found by Kőszegi and Li (2008) for the three-period model which could be represented graphically as follows:

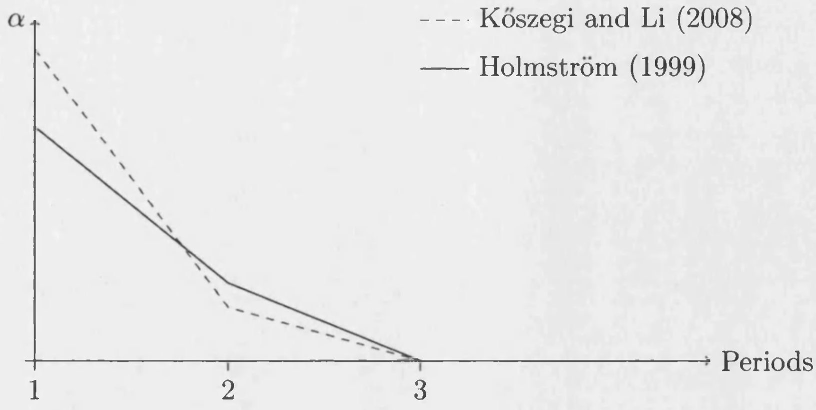

Figure 3.4: The effort dynamics in a three-period model.

However, it is not clear that initially the forward attribution needs to be larger than the backward attribution. Note that the condition for this is strictest when the uncertainty about the talent is highest ($\frac{\sigma_\theta^2}{\sigma_\theta^2+\sigma_\varepsilon^2}$), i. e. in the very beginning. Therefore, it is also conceivable to obtain the following effort dynamics.

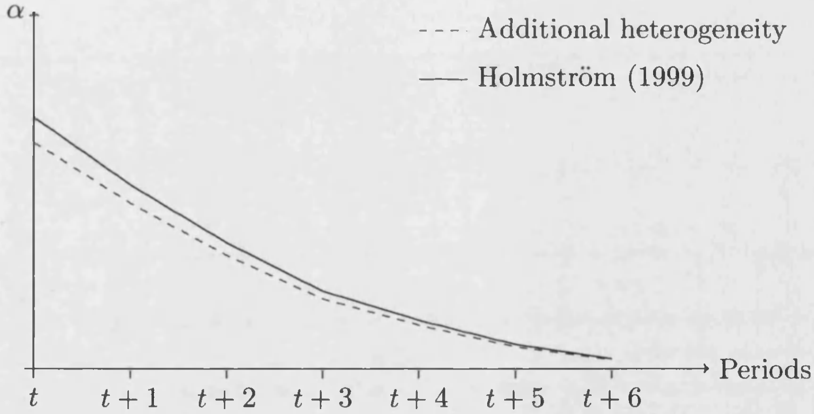

Figure 3.5: Possible effort dynamics with and without type uncertainty.

Obviously, the scope for a higher forward attribution depends on the planning horizon of the agent. In my model, for computational purposes I restricted myself to a one-period planning horizon which gives rise to rather restrictive conditions. With a longer horizon, like the two-period horizon in Kőszegi and Li (2008), the forward attribution obtains a higher weight, making the process depicted in figure 3.3 more likely. At the same time, the horizon consideration points towards the relevance of not only the higher career concerns as considered so far, but also towards the necessity of planning ahead long enough.

Finally, the condition on the relative size of the two effects dictates that there is no oscillation around the standard effort level. Without being able to prove it, in my opinion it is also true that the effort pathways cross one time at most.

3.3.7 Interpretation and extensions

The results of my model show a congruence with the intuition that was developed by Kőszegi and Li (2008). While the analysis of the two-period model indicates the relevance of an incentive muting due to additional heterogeneity, in the longer horizon model both an increased effort in the beginning and constantly lower effort exertion can occur depending on the parameters.

This indicates that it is not necessarily true that the beginning of a career is also connected with a higher effort due to type signalling. We can observe that the signalling of the type can be a hinderance for effort exertion in the situation where agents are afraid that the past talent evaluation is in jeopardy when they display too large career concerns.

The possibility of an increase in effort thanks to a noisy signal of effort hinges on the fact that the forward attribution is higher than the backward attribution. Letting the principal receive a noisy signal on effort implies that the relevance of this observed period's effort is increased and the updating accelerated. However, this measure will exhibit an adverse effect if it is done in situations of higher backward attribution. Here, effort would be muted even more and the increased precision in the inferred talent and type leads to a reduction of effort for all remaining periods.

Of course, this statement needs to be qualified and revisited when it becomes feasible to examine a longer period game with a planning horizon of the agent that is larger than just the one period I assumed here. This qualifier seems to point towards the question as to whether the horizon in the planning can also be considered to be part of a certain heterogeneity, since career concerns are obviously all about the planning horizon and future payoffs.

3.4 Conclusion

In this chapter I presented an analysis of additional heterogeneity in the career concerns model. Starting out with the general consideration of heterogeneity in the framework of a two-period career concern model the effect of backward attribution was introduced and found to be a constant feature of any incentive structure under additional heterogeneity. While this effect's direction could also be determined to be effort reducing, the analysis of the infinite-period model is not as straightforward.

In the multi-period model the analysis hinges on the interplay between backward attribution and forward attribution, a second effect that results from the incentive to signal high career concerns. Since the interplay is relative intricate, the analysis is confined to point out the main qualitative features of the effort dynamics. In particular, it cannot be ruled out that the additional heterogeneity does mute incentives throughout the duration of the game. However, there is also the possibility of a dynamic pathway that is similar to the one found in Kőszegi and Li (2008). Here incentives are increased in the beginning of the career and muted in the end. This indicates that measures like additional noisy signals on effort work only under certain circumstances in the desired direction.

A main take-away point should be the strong possibility of a non-monotonicity of effort levels in uncertainty. While the original model of Holmström (1999) highlights the purpose of uncertainty for achieving an efficient effort level, the present analysis shows that uncertainty about other dimensions of relevant type characteristics of the agent might work against this, reducing effort levels under increased uncertainty.

C Appendix

C.1 Proof of proposition with more than 2 types

Proof. (Proposition 1 with N types) Analogously to equation 3.2 the objective function with N types becomes

$$\int_y k \left[\sum_{k=k^0}^{k^N} Pr(k|y) \left(\int_\theta \theta \frac{f(\theta, y|a^{k*})}{\hat{f}(y|a^{k*})} \, d\theta \right) \right] \hat{f}(y|a) \, dy - c(a). \tag{3.9}$$

By the MLRP arguments we can state the relationships between talent inferences as

$$E(\theta|y, a^{0*}) \geq E(\theta|y, a^{1*}) \geq \ldots \geq E(\theta|y, a^{N*}).$$

Since $\sum_{k=k^0}^{k^N} Pr(k|y) = 1$ it follows that $\sum_{k=k^0}^{k^N} \frac{\partial}{\partial y} Pr(k|y) = 0$ and $\frac{\partial}{\partial y} Pr(k = k^0|y) = -\sum_{k=k^1}^{k^N} \frac{\partial}{\partial y} Pr(k|y)$. We analyse the change in the integrand that comes from a change in output, which is a direct consequence of a marginal increase in effort.

The expression in square brackets is

$$\sum_{k=k^0}^{k^N} Pr(k|y) \cdot E(\theta|y, a^{k*})$$

and its derivative with respect to y becomes

$$\sum_{k=k^0}^{k^N} \frac{\partial}{\partial y} Pr(k|y) \cdot E(\theta|y, a^{k*}) + \sum_{k=k^0}^{k^N} Pr(k|y) \cdot \frac{\partial}{\partial y} E(\theta|y, a^{k*}).$$

Using the mentioned results this can be rewritten as

$$\frac{\partial}{\partial y} Pr(k = k^0|y) \cdot \left[E(\theta|y, a^{0*}) - \frac{\sum_{k=k^1}^{k^N} \frac{\partial}{\partial y} Pr(k|y) \cdot E(\theta|y, a^{k*})}{\sum_{k=k^1}^{k^N} \frac{\partial}{\partial y} Pr(k|y)} \right] + \frac{\partial}{\partial y} E(\theta, \hat{a}^*).$$

The very last expression is the one known from the standard model, where the higher output level positively impacts the expectation of the talent. The expressions before result from the association of the output to different types of agents.

The sign of this additional expressions can be found by observing that the first expression $\frac{\partial}{\partial y} Pr(k = k^0|y) < 0$ as a higher output decreases the probability that the agent is the one with lowest career concerns.

Because the fraction inside the brackets is necessarily greater than $E(\theta|y, a^{1*}) > E(\theta|y, a^{0*})$, the expression in the brackets is negative. Therefore, as in the case of two types, it follows that the marginal incentives are muted for all types as well. ■

C.2 Proof of Proposition 3

To calculate the different required conditional expectations, we need to consider a multi-variate normal distribution that consists of the 'talent' and 'type' variables that are updated over time and the observables y_t and y_{t+1} that are used for the updating.

$$E\left[\begin{pmatrix} \theta \\ k \\ \theta+\alpha_t\cdot k+\varepsilon_t \\ \theta+\alpha_{t+1}\cdot k+\varepsilon_{t+1} \end{pmatrix}\right] = \begin{pmatrix} 0 \\ \mu_k \\ \alpha_t\cdot\mu_k \\ \alpha_{t+1}\cdot\mu_k \end{pmatrix}$$

$$Var\left[\begin{pmatrix} \theta \\ k \\ \theta+\alpha_t\cdot k+\varepsilon_t \\ \theta+\alpha_{t+1}\cdot k+\varepsilon_{t+1} \end{pmatrix}\right] = \begin{pmatrix} \sigma_\theta^2 & 0 & \sigma_\theta^2 & \sigma_\theta^2 \\ 0 & \sigma_k^2 & \alpha_t\sigma_k^2 & \alpha_{t+1}\sigma_k^2 \\ \sigma_\theta^2 & \alpha_t\sigma_k^2 & \sigma_\theta^2+\alpha_t^2\sigma_k^2+\sigma_\varepsilon^2 & \sigma_\theta^2+\alpha_t\alpha_{t+1}\sigma_k^2 \\ \sigma_\theta^2 & \alpha_{t+1}\sigma_k^2 & \sigma_\theta^2+\alpha_t\alpha_{t+1}\sigma_k^2 & \sigma_\theta^2+\alpha_{t+1}^2\sigma_k^2+\sigma_\varepsilon^2 \end{pmatrix}$$

The conditional of a multivariate normal distribution of the form

$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \sim N\left(\begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix}, \begin{pmatrix} \Sigma_{11} & \Sigma_{21}' \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix}\right),$$

is given by the following expression

$$x_1|x_2 \sim N\left(\mu_1+\Sigma_{21}'\Sigma_{22}^{-1}(x_2-\mu_2), \Sigma_{11}-\Sigma_{21}'\Sigma_{22}^{-1}\Sigma_{21}\right).$$

The following four expressions for expected value and variance obtain.

$$\begin{aligned} E(k|y_t) &= \mu_k+\alpha_t\sigma_k^2(\sigma_\theta^2+\alpha_t^2\sigma_k^2+\sigma_\varepsilon^2)^{-1}(y_t-\alpha_t\mu_k) \\ &= \mu_k+\frac{\alpha_t\sigma_k^2}{\sigma_\theta^2+\alpha_t^2\sigma_k^2+\sigma_\varepsilon^2}(y_t-\alpha_t\mu_k) \end{aligned}$$

$$Var(k|y_t) = \bar{\sigma}_k^2 = \sigma_k^2-\frac{(\alpha_t\sigma_k^2)^2}{\sigma_\theta^2+\alpha_t^2\sigma_k^2+\sigma_\varepsilon^2} \leq \sigma_k^2 \tag{3.10}$$

$$= \frac{\sigma_k^2(\sigma_\theta^2+\sigma_\varepsilon^2)}{\sigma_\theta^2+\alpha_t^2\sigma_k^2+\sigma_\varepsilon^2} \tag{3.11}$$

$$\begin{aligned}
E(\theta|k,y_t) &= 0+\begin{pmatrix} 0 & \sigma_\theta^2 \end{pmatrix}\begin{pmatrix} \sigma_k^2 & \alpha_t\sigma_k^2 \\ \alpha_t\sigma_k^2 & \sigma_\theta^2+\alpha_t^2\sigma_k^2+\sigma_\varepsilon^2 \end{pmatrix}\begin{pmatrix} k-E(k|y_t) \\ y_t-\alpha_t E(k|y_t) \end{pmatrix} \\
&= \begin{pmatrix} \sigma_\theta^2\cdot I_{21} & \sigma_\theta^2 I_{22} \end{pmatrix}\begin{pmatrix} k-E(k|y_t) \\ y_t-\alpha_t E(k|y_t) \end{pmatrix} \\
&= \frac{\sigma_\theta^2}{\sigma_\theta^2+\sigma_\varepsilon^2}\left(y_t-\alpha_t E(k|y_t)\right)
\end{aligned}$$

Note that $k-E(k|y_t)=0$ since at that point in time, there is no perceivable difference between the two since k is not observed as opposed to y_t.

$$Var(\theta|k,y_t)=\bar{\sigma}_\theta^2=\frac{\sigma_\theta^2\sigma_\varepsilon^2}{\sigma_\theta^2+\sigma_\varepsilon^2}<\sigma_\theta^2 \tag{3.12}$$

With these expressions the wage in the next period is

$$\begin{aligned}
w_{t+1} &= \frac{\sigma_\theta^2}{\sigma_\theta^2+\sigma_\varepsilon^2}\left(y_t-\alpha_t E(k|y_t)\right)+\alpha_{t+1}E(k|y_t) \\
&= \frac{\sigma_\theta^2}{\sigma_\theta^2+\sigma_\varepsilon^2}y_t+\left(\alpha_{t+1}-\frac{\sigma_\theta^2}{\sigma_\theta^2+\sigma_\varepsilon^2}\alpha_t\right)E(k|y_t).
\end{aligned}$$

Deriving with respect to y_t proves the proposition.

A similar derivation can be done for the subsequent period $t+1$, in order to get an idea as to what the next effort parameter α_{t+1} looks like.

$$\begin{aligned}
w_{t+2} &= E(y_{t+1}|y_t,y_{t+1})=E(\theta|y_t,y_{t+1})+\alpha_{t+2}E(k|y_t,y_{t+1}) \\
\alpha_{t+1} &= \frac{\partial w_{t+2}}{\partial y_{t+1}} \\
&= \frac{\sigma_\theta^2\left(\alpha_t(\alpha_t-\alpha_{t+1})\sigma_k^2+\sigma_\varepsilon^2\right)}{(\sigma_\theta^2+\alpha_t^2\sigma_k^2+\sigma_\varepsilon^2)(\sigma_\theta^2+\alpha_{t+1}^2\sigma_k^2+\sigma_\varepsilon^2)-(\sigma_\theta^2+\alpha_t\alpha_{t+1}\sigma_k^2)^2} \\
&\quad -\frac{\alpha_{t+1}\sigma_\theta^2\left(\alpha_t(\alpha_t-\alpha_{t+1})\sigma_k^2+\sigma_\varepsilon^2\right)\left[\alpha_{t+1}\sigma_k^2(\sigma_\theta^2+\alpha_t^2\sigma_k^2+\sigma_\varepsilon^2)-\alpha_t\sigma_k^2(\sigma_\theta^2+\alpha_t\alpha_{t+1}\sigma_k^2)\right]}{\left((\sigma_\theta^2+\alpha_t^2\sigma_k^2+\sigma_\varepsilon^2)(\sigma_\theta^2+\alpha_{t+1}^2\sigma_k^2+\sigma_\varepsilon^2)-(\sigma_\theta^2+\alpha_t\alpha_{t+1}\sigma_k^2)^2\right)^2} \\
&\quad +\alpha_{t+2}\frac{\alpha_{t+1}\sigma_k^2(\sigma_\theta^2+\alpha_t^2\sigma_k^2+\sigma_\varepsilon^2)-\alpha_t\sigma_k(\sigma_\theta^2+\alpha_t\alpha_{t+1}\sigma_k^2)}{(\sigma_\theta^2+\alpha_t^2\sigma_k^2+\sigma_\varepsilon^2)(\sigma_\theta^2+\alpha_{t+1}^2\sigma_k^2+\sigma_\varepsilon^2)-(\sigma_\theta^2+\alpha_t\alpha_{t+1}\sigma_k^2)^2}
\end{aligned}$$

References

Attali, Yigal and Maya Bar-Hillel, "Guess Where: The Position of Correct Answers in Multiple-Choice Test Items as a Psychometric Variable," *Journal of Educational Measurement*, Sommer 2003, *40* (2), 109–128.

Bacharach, Michael, "Variable Universe Games," in Ken Binmore, Alan Kirman, and Piero Tani, eds., *Frontiers of Game Theory*, The MIT Press, 1993, pp. 255–275.

__ **and Dale O. Stahl**, "Variable-Frame Level-N Theory," Working Paper, Oxford University November 1997.

__ **and** __ , "Variable-Frame Level-n Theory," *Games and Economic Behavior*, August 2000, *32* (2), 220–246.

Bosch-Domènech, Antoni, José G. Montalvo, Rosemarie Nagel, and Albert Satorra, "One, Two, (Three), Infinity, . . . : Newspaper and Lab Beauty-Contest Experiments," *American Economic Review*, December 2002, *92* (5), 1687–1701.

__ , __ , __ , **and** __ , "Finite Mixture Analysis of Beauty-Contest Data from Multiple Samples," Economics Working Papers 737, Department of Economics and Business, Universitat Pompeu Fabra January 2004.

Brosig, Jeannette, "Identifying cooperative behavior: some experimental results in a prisoner's dilemma game," *Journal of Economic Behavior and Organization*, March 2002, *47* (3), 275–290.

Butler, David J. and Graham C. Loomes, "Imprecision as an account of the preference reversal phenomenon," *American Economic Review*, March 2007, *97* (1), 277–297.

Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong, "A Cognitive Hierarchy Model of Games," *The Quarterly Journal of Economics*, August 2004, *119* (3), 861–898.

Christenfeld, Nicholas, "Choices from identical options," *Psychological Science*, January 1995, *6* (1), 50–55.

Coate, Stephen and Stephen Morris, "On the Form of Transfers to Special Interests," *Journal of Political Economy*, December 1995, *103* (6), 1210–1235.

Cooper, David J. and John H. Kagel, "Are Two Heads Better than One? Team versus Individual Play in Signaling Games," *American Economic Review*, June 2005, *95* (3), 477–509.

Costa-Gomes, Miguel A. and Georg Weizsäcker, "Stated Beliefs and Play in Normal-Form Games," *Review of Economic Studies*, July 2008, *75* (3), 729–762.

__ **and Vincent P. Crawford**, "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study," *American Economic Review*, December 2006, *96* (5), 1737–1768.

__ **and** __ , "Studying Cognition via Information Search in Two-Person

Guessing Game Experiments," Working Paper, University of York 2009.

__ , __ , **and Bruno Broseta**, "Cognition and Behavior in Normal-Form Games: An Experimental Study," *Econometrica*, September 2001, *69* (5), 1193–1235.

Crawford, Vincent P., "Lying for strategic advantage: Rational and boundedly rational misrepresentation of intentions," *American Economic Review*, March 2003, *93* (1), 133–149.

__ **and Nagore Iriberri**, "Fatal Attraction: Focality, Naïveté and Sophistication in Experimental 'Hide and Seek' Games," *American Economic Review*, December 2007, *97* (5), 1731–1750.

__ **and** __ , "Level-k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?," *Econometrica*, November 2007, *75* (6), 1721–1770.

__ , **Uri Gneezy, and Yuval Rottenstreich**, "The Power of Focal Points Is Limited: Even Minute Payoff Asymmetry May Yields Large Coordination Failures," *American Economic Review*, July 2008, *98* (4), 1443–1458.

Dewatripont, Mathias, Ian Jewitt, and Jean Tirole, "The Economics of Career Concerns, Part I: Comparing Information Structures," *Review of Economic Studies*, January 1999, *66* (1), 183–198.

__ , __ , **and** __ , "The Economics of Career Concerns, Part II: Applications to Missions and Accountability of Government Agencies," *Review of Economic Studies*, January 1999, *66* (1), 199–217.

Ericsson, K.A., "Towards a procedure for eliciting verbal expression of non-verbal experience without reactivity: interpreting the verbal overshadowing effect within the theoretical framework for protocol analysis," *Applied Cognitive Psychology*, December 2002, *16* (8), 981–987.

Eyster, Erik and Matthew Rabin, "Naïve Herding," Working Paper, University of California Berkeley May 2008.

Fama, Eugene F., "Agency Problems and the Theory of the Firm," *Journal of Political Economy*, April 1980, *88* (2), 288–307.

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

Geng, Hong and Heike Hennig-Schmidt, "In Your Shoes - A Video Experimental Study on Communication and Quasi-communication in Ultimatum Bargaining with Groups," Working Paper, University of Bonn July 2007.

Halpern, Benjamin R. and George M. Bergman, "A Fixed-Point Theorem for Inward and Outward Maps," *Transactions of the American Mathematical Society*, February 1968, *130* (2), 353–358.

Hennig-Schmidt, Heike and Hong Geng, "Decision Making in Tacit Bargaining. A Video Experimental Investigation on the Power-to-take Game," Working Paper, University of Bonn October 2005.

Holmström, Bengt, "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies*, January 1999, *66* (1), 169–182.

Kőszegi, Botond and Wei Li, "Drive and Talent," *Journal of the European Economic Association*, March 2008, *6* (1), 210–236.

Kocher, Martin G. and Matthias Sutter, "The Decision Maker Matters: Individual versus Group Behaviour in Experimental 'Beauty-Contest' Games," *The Economic Journal*, January 2005, *115* (500), 200–223.

Kremer, Michael, "The O-Ring Theory of Economic Development," *The Quarterly Journal of Economics*, August 1993, *108* (3), 551–75.

Levy, Gilat, "Decision Making in Committees: Transparency, Reputation, and Voting Rules," *American Economic Review*, March 2007, *97* (1), 150–168.

Lewis, David K., *Convention: A Philosophical Study*, Harvard University Press, 1969.

Maskin, Eric and Jean Tirole, "The Politician and the Judge: Accountability in Government," *American Economic Review*, September 2004, *94* (4), 1034–1054.

Mehta, Judith, Chris Starmer, and Robert Sugden, "The Nature of Salience: An Experimental Investigation of Pure Coordination Games," *American Economic Review*, June 1994, *84* (3), 658–673.

Milgrom, Paul R., "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, Autumn 1981, *12* (2), 380–391.

Morris, Stephen, "Political Correctness," *Journal of Political Economy*, April 2001, *109* (2), 231–265.

Nagel, Rosemarie, "Reasoning and Learning in Guessing Games and Ultimatum Games With Incomplete Information: An Experimental Investigation." PhD dissertation, Universität Bonn 1994.

__ , "Unraveling in Guessing Games: An Experimental Study," *American Economic Review*, December 1995, *85* (5), 1313–1326.

__ , "A Survey on Experimental Beauty Contest Games: Bounded Rationality and Learning," in David V. Budescu, Ido Erev, and Rami Zwick, eds., *Games and Human Behavior*, Lawrence Erlbaum Associates, 1999, pp. 105–142.

O'Neill, Barry, "Nonmetric test of the minimax theory of two-person zerosum games," *Proceedings of the National Academy of Sciences of the United States of America*, 1987, *84* (7), 2106–2109.

Ottaviani, Marco and Peter N. Sorensen, "Professional advice," *Journal of Economic Theory*, January 2006, *126* (1), 120–142.

Penczynski, Stefan P., "Career Concerns under Multiplicity of Tasks and Principals," December 2006. Mimeo, London School of Economics.

Prat, Andrea, "The Wrong Kind of Transparency," *American Economic Review*, June 2005, *95* (3), 862–877.

Prendergast, Canice and Lars Stole, "Impetuous Youngsters and Jaded Old-timers: Acquiring a Reputation for Learning," *Journal of Political Economy*, December 1996, *106* (4), 1105–1134.

Puri, Manju and David T. Robinson, "Optimism and economic choice," *Journal of Financial Economics*, October 2007, *86* (1), 71–99.

Rubinstein, Ariel, "Experience from a Course in Game Theory: Pre- and Postclass Problem Sets as a Didactic Device," *Games and Economic Behavior*, July 1999, *28* (1), 155–170.

__ **, Amos Tversky, and Dana Heller**, "Naive Strategies in Zero-Sum Games," in Wulf Albers, Werner Güth, Peter Hammerstein, Benny Molduvano, and Eric van Damme, eds., *Understanding Strategic Interaction–Essays in Honor of Reinhard Selten*, Springer Verlag, 1996, pp. 394–402.

__ **and** __ , "Naive Strategies in Zero-Sum Games," Working Paper 17-93, The Sackler Institute of Economic Studies 1993.

Rydval, Ondrej, Andreas Ortmann, and Michal Ostatnicky, "Three Very Simple Games and What It Takes to Solve Them," Jena Economic Research Papers 2007-092, Friedrich-Schiller-University Jena, Max-Planck-Institute of Economics, Thueringer Universitaets- und Landesbibliothek November 2007.

Salmon, Timothy C., "An Evaluation of Econometric Models of Adaptive Learning," *Econometrica*, November 2001, *69* (6), 1597–1628.

__ , "Evidence for Learning to Learn Behavior in Normal Form Games," *Theory and Decision*, June 2004, *56* (4), 367–404.

Scharfstein, David S. and Jeremy C. Stein, "Herd Behavior and Investment," *American Economic Review*, June 1990, *80* (3), 465–479.

Schelling, Thomas C., *The Strategy of Conflict*, Harvard University Press, 1960.

Selten, Reinhard, "Features of experimentally observed bounded rationality," *European Economic Review*, May 1998, *42* (3), 413–436.

Stahl, Dale O. and Paul W. Wilson, "On Players' Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior*, July 1995, *10* (1), 218–254.

Tirole, Jean, "The Internal Organization of Government," *Oxford Economic Papers*, January 1994, *46* (1), 1–29. New Series.

von Neumann, John, "A Certain Zero-Sum Two-Person Game Equivalent to the Optimal Assignment Problem," in Harold W. Kuhn and Albert W. Tucker, eds., *Contributions to the Theory of Games, II*, Vol. 28 of Annals of Mathematics Studies, Princeton University Press, 1953, pp. 5–12.

__ **and Oskar Morgenstern**, *Theory of Games and Economic Behaviour*, John Wiley and Sons, 1944.

Wilson, James Q., *Bureaucracy: What Government Agencies Do and Why They Do It*, BasicBooks, 1989.