Computational Complexity and Level-k Reasoning in Games

by

Zsolt Udvari

Submitted to

Central European University

Department of Economics

In partial fulfillment of the requirements for the degree of Master of Arts

Supervisor: Professor Botond Kőszegi

Budapest, Hungary

2013

Abstract

In this thesis I develop a model which is able to capture the complexity of games and predict the depth of reasoning performed by the players across different games. I use a modified version of a Turing machine and measure the complexity of a level-k strategy with the number of moves the machine has to make to compute the given strategy using the parameters of the game as inputs. This analyzing framework is able to explain some part of the variation in the observed cognitive type distribution found out by experimental papers.

Contents

1 Introduction

The equilibrium approach used in game theory often fails to explain the outcomes of games played by real subjects during an experiment. This well-known fact inspired several scholars to think about non-equilibrium approaches to explain the outcomes of the experiments and to have a better understanding of strategic thinking. This field of research developed several new alternative models. The examples are equilibrium with noise, level-k models (Nagel, 1995; Stahl and Wilson, 1994, 1995) and the closely related cognitive hierarchy model (Camerer, Ho and Chong, 2004), and quantal response equilibrium (McKelvey and Palfrey, 1995). Crawford, Costa-Gomes and Iriberri (2013) provide a survey of both theoretical and empirical papers about these models; based on these findings they suggest that the level-k models should be added to the analyst's tools since they can explain a large fraction of empirical results.

The model of level-k reasoning was first introduced in the 1990's by Nagel (1995) and Stahl and Wilson (1994, 1995). The model assumes a distribution of cognitive types of the players, all type is a level-k type with different k. All of the k types play their best response to the type $k-1$'s actions. The level-0 type is a nonstrategic one that chooses its actions without any reasoning about the strategic situation. The definition of the level-0 strategy is the crucial step in the construction of a level-k model. According to Crawford, Costa-Gomes and Iriberri (2013) there are two methods in the literature for defining the level-0 strategy; one of them is uniform randomness and the other is attraction to salience.

There is an extensive experimental literature which tests the goodness-of-fit of the predictions of a level-k models. Besides the experimental literature, there are some recent theoretical advances related to the level-k reasoning. Strzalecki (2011), Heifetz and Kets (2012), Kets (2012), and Pintér and Udvari (2012) provide type spaces which are able to express finite belief hierarchies (and therefore level-k reasoning).

Despite the large body of empirical papers using level-k models to explain results of experiments, the determinants of the distribution of the cognitive types with level-k depth has not received much attention. In the experimental papers the distribution of the level-k types was assumed to be exogenous and never depended on the parameters of the game played, neither on the subject pool.

To the best of my knowledge, Alaoui and Penta (2013a,b) were the first who investigated this question. Alaoui and Penta (2013b) use an axiomatic approach to provide a theoretical framework in which the reasoning process in a game is an outcome of a cost-benefit analysis. In their model, the costs of reasoning are exogenous and depend on the cognitive abilities of the players, while the benefits are determined by the payoffs of the game played.

They test their theory with an experiment in Alaoui and Penta (2013a). Their subjects played a modified version of the 11-20 game introduced by Arad and Rubinstein (2012) which is the most suitable among game types to a level-k analysis. Their experimental design enabled them to examine how the sophistication (higher analytical skills), beliefs about the opponent's sophistication, incentives (higher rewards) and beliefs about the opponent's incentives influence the depth of reasoning of the players. They found support that the underlying process which determines the depth of reasoning can be modeled as a cost-benefit analysis.

The work of Alaoui and Penta thus successfully applied classical economic concepts to model the initial responses for a given game, however, the further specification of the cost function was outside their paper's focus. They suggested future research for the specification of the cost function which would enable to make predictions about the depth of knowledge *across* different games. This thesis contributes toward this direction of research.

Here I assume that the costs of the reasoning are determined – given the cognitive ability of the player – by the computational complexity of the game played. Joining the research line which started with Rubinstein (1986) I use a kind of automaton (it is named game machine) to capture the complexity of the level-k strategies of different games. The game machine introduced here can capture more of the computational complexity than the finite automaton used by Rubinstein (1986). After the review of the level-k studies about different games, I apply the game machine to evaluate the computational costs of calculating level-k best responses for the game types I investigate. The new model can provide new insights into the strategic thinking process and able to explain some of the variation in the results of experiments.

This thesis is organized as follows. Section 2 summarizes the results of experimental level-k studies. Section 3 describes the models of computational complexity and introduces the game machine model. Section 4 uses the newly developed game machine model to evaluate the games described in Section 2. Section 5 conducts the overall analysis across games, Section 6 concludes and provides recommendations for an empirical study.

2 Overview of experimental results

In this section I overview the experimental findings about the distribution of level-k types of players in different games. The goal is not to summarize all of the game theoretic experiments conducted so far, I only focus to the empirical papers which concentrated on the estimation of the distribution of level-k type players. I do not include papers lacking the description of the level-k distribution.

Different types of games are analyzed in different subsections. In the case of all game types, first I briefly introduce the game in question. The description is followed by the summary of experiment results including the important details of experiment designs. I do not refer to another details of research designs if they are not important for the purposes of this thesis. At the end of each experiment description I present a table which contains the relative frequencies of the identified types reported by the authors.

Subsection 2.1 deals with the Beauty Contest-like outguessing games, Subsection 2.2 with other normal form games with two players. Hide-and-seek games are included in Subsection 2.3, and Subsection 2.4 covers the most recent 11-20 money request game.

2.1 Beauty Contest games

Beauty Contest games are extremely useful tools to analyze the subjects' depth of thinking. The name of this type of game originates with Keynes (1936). Keynes (1936) compared the activity of professional investors to a newspaper competition where the competitors have to choose six faces of hundred photographs, and the winner is the one whose choice is the nearest to the average choice of all competitors. In the conducted experiments the subjects' task is less complicated: they have to pick a number from an interval $[a, b]$, and the winner is whose choice is the nearest to the average times p, where p is a positive number. So the players have to guess the average of the other players' choices. The parameters a, b and p, as well as the number of participants varies across experiments.

This kind of game has many attractive features for a level-k analysis. First, the game has unique symmetric equilibrium for all $p \neq 1$, which can be found by iterated dominance. If $p < 1$, then the symmetric equlibrium guess is a;and it is b if $p > 1$. Knowing the clear prediction of the equlibrium concept, it is easy to compare the predictions of the equilibrium and the level-k approaches.

Second, there is a natural way to define the level-0 strategy. A non-strategic player chooses a number randomly or choose a number which is salient to him – for example his favorite number (Nagel, 1995). Therefore the average of guesses by a population of level-0 players will be the expected value of the numbers in $[a, b]$. With such natural level-0 strategy it is easy to find the level-1 and higher order players' actions. A level-1 player best responding to the level-0 players actions will guess p times the average, and level-n players will guess p^n times the average respectively. Given the large strategy space (another advantage of this game type) and the sharp identification of the level-0 strategy – and the higher level strategies implied by that – these outguessing games are excellent for a level-k analysis.

Outguessing games are extensively discussed by the literature, but most of the papers concentrate only on the winning number in different settings and not the distribution of level-k types. This subsection covers the findings of Nagel (1995), Ho, Camerer and Wiegelt (1998) and Costa-Gomes and Crawford (2006) as they explicitly focused on the type distribution.

2.1.1 Nagel (1995)

The experiment run by Nagel (1995) was the first which examined the level-k distribution in outguessing games. The interval was $[0, 100]$ in each sessions, but three different p was used: $\frac{1}{2}$, $\frac{2}{3}$ and $\frac{4}{3}$. Nagel (1995) allowed some noise during the identification of the level-k guesses, she looked at neighborhoods of the $50p^n$ guesses not only the points $50p^n$. The relative frequency of level-k guesses is presented by Table 1 (the results are from Figure 2 of Nagel (1995)).

Table 1: Results of Nagel (1995)

p	$L0$	$L1$	$L2$	$L3$
1/2	24%	28%	41%	7%
2/3	12%	45%	38%	5%
4/3	4%	59%	20%	18%

A very similar experiment was ran by Ho, Camerer and Wiegelt (1998). In their analysis the authors categorized all the guesses as a level-k guess according to intervals which contain level-k undominated guesses. This method led to too much noise in the data. It is hard to imagine such high frequency (34-49%) of the $L3$ type they found while in other papers the frequency of $L3$ is almost always (well) below 10%. Due to that outlying results I do not include Ho, Camerer and Wiegelt (1998) in my analysis.

2.1.2 Costa-Gomes and Crawford (2006)

Another article I include is Costa-Gomes and Crawford (2006) which analyzed guessing games played by two players. Costa-Gomes and Crawford (2006)'s subjects played a series of 16 similar games without any feedback of the results, which excluded learning effects. The series of games allowed a more precise identification of players' types. In addition, the two-player setting made sure that the players consider themselves and their own actions as influencing factors in the environment.

Another difference from Nagel (1995)'s and Ho, Camerer and Wiegelt (1998)'s design is that the two players had different intervals and different targets (though they were aware of the other player's interval and target as the structure was publicly announced). It was possible that both of the targets are above 1, both of the targets are below 1, or they were mixed. The equilibrium is determined by the product of the two targets and always exists if that product is not equal to 1 (and it never happened in the design). The authors chose this variable structure to help the understanding of the game by the subjects and let them concentrating on the predictions, therefore reduce the noisiness of the guesses.

The experimental design described above enabled to separate the cognitive types of players more reliably. Their estimation included six different types: $L1$, $L2$, $L3$, $D1$, $D2$ and *Equilibrium*. The Dk types are similar to $Lk+1$ types, the difference between them is that the $Lk+1$ simply best responds to Lk, while Dk best responds to a distribution of lower-level types. At the end the authors rejected the Dk types in favor of the $Lk+1$ types. Subjects with $L0$ guesses were not found at all in this experiment.

Costa-Gomes and Crawford (2006) was able to reliable identify guesses of 43 subjects out of 88 that are exactly corresponding to a hypothetical type. The results are presented on Table 2.

Table 2: Reliably identified types (Costa-Gomes and Crawford, 2006)

Type	$L0$	$L1$	$L2$	$L3$
Frequency	0%	56%	34%	10%

2.2 Two-person matrix games

The two-person normal form games were also among the game types that became subjects of the level-k analysis (Stahl and Wilson, 1994, 1995). This subsection covers three articles which analyzed two-person matrix games: Stahl and Wilson (1994), Stahl and Wilson (1995) and Costa-Gomes, Crawford and Broseta (2001). The strategy space of these simple matrix games are much coarser than the strategy space of Beauty Contest games. To ensure the identification of the players' depth of thinking, in all of the experiments presented below the subjects played multiple games in the experiment sessions (similarly as in Costa-Gomes and Crawford (2006)).

Again, there are lots of another papers about normal form game experiments, here I present the papers which explicitly concentrate on the estimation of the k-distributions (this statement is also true for the hide-and-seek games, though I will not mention it again).

2.2.1 Stahl and Wilson (1994)

In the experiment of Stahl and Wilson (1994) the subjects played 10 symmetric 3×3 matrix games, all of them as a row player, and after that the actions were matched.

After considering multiple econometrical estimations, 35 subjects had very high probability to being one type. The authors found the following type distributions (notice that in the case of 3×3 matrix games, $L3$ plays the equilibrium; $L0$ subjects were not identified). Table 3 presents the results.

Table 3: Findings of Stahl and Wilson (1994)

Type	$L0$	$L1$	$L2$	$L3$
Frequency	0%	24%	51%	25%

In an other paper, Stahl and Wilson (1995) conducted an analysis with improved experimental design and econometric analysis. In that setting the authors included a new type, they called it 'worldly', which corresponds to a level-4 player. 38 subjects out of 48 found to be one of these types with at least 0.9 posterior probability, and 17 of them found to be a worldly type. The high number of worldly players is striking. However, this estimation method received criticism in Costa-Gomes, Crawford and Broseta (2001) and Crawford, Costa-Gomes and Iriberri (2013). These authors argue that the over-parametrization of Stahl and Wilson (1995)'s model led to the rejection of $L2$ types in favor of worldly. Costa-Gomes, Crawford and Broseta (2001) got more similar results to Stahl and Wilson (1994) - which does not include the $Worldly$ type - than the results to Stahl and Wilson (1995), despite that the designs of Stahl and Wilson (1994) and Stahl and Wilson (1995) were very similar to each other. I agree to the criticism of Crawford, Costa-Gomes and Iriberri (2013) and due to the irregularly high level of $L4$ types I do not include the results of Stahl and Wilson (1995) into my analysis.

2.2.2 Costa-Gomes, Crawford and Broseta (2001)

The subjects there played 18 matrix games, there were 2×2, 2×3 and 2×4 games that were designed to enable separating the choices of strategic and nonstrategic types as much as possible.

The authors included several types in their estimation, but all of their identified types coincide with one of the Lk types of Stahl and Wilson (1994). According to the econometric analysis of the authors, 58 subjects out of 72 found to be one type with at least 0.9 posterior probability. The results are presented by Table 4.

Table 4: Costa-Gomes, Crawford and Broseta (2001)

Type	$L0$	$L1$	$L2$	$L3$
Frequency	7%	24%	59%	10%

2.3 Hide-and-seek games

The hide-and-seek games first introduced by Rubinstein, Tversky and Heller (1996) are two-person games in normal form with zero-sum payoffs. One of the players is the "hider", her task is to hide a "treasure" in one of the four locations (boxes), while the "seeker" player's task is to find it. The hider wins if the seeker picks a box which is not the one that selected by the hider and vice versa. Therefore the interests of the two players are completely conflicting. (Note that the game is still a simultaneous move game even if the hider acts first since the action is not observed by the seeker).

Another important feature of the game is the non-neutral framing of the boxes. The four boxes are signed by letters A, B, A, A respectively. The standard game theoretical prediction is that all four locations will be chosen with equal probability (mixed strategy equilibrium, notice that the letters do not affect payoffs). However, as Rubinstein, Tversky and Heller (1996) found out, it is not true. The framing makes some locations more salient to the players causing deviations from the equilibrium. Rubinstein, Tversky and Heller (1996) argues that besides the box denoted by B (which is trivially salient) the A boxes at the two ends are also can be salient, so the least salient location is the central A.

Using the salience concept, we can identify the $L0$ choice as the salient location and conduct a level-k analysis. In this subsection I present two level-k analysis of experimental hide-and-seek games, Crawford and Iriberri (2007) and Penczynski (2011).

2.3.1 Crawford and Iriberri (2007)

The paper by Crawford and Iriberri (2007) was the first that proposed a level-k analysis to explain the results of Rubinstein, Tversky and Heller (1996), and their other goal was to explore the correct specification of level-k models for games with non-neutral landscapes. The authors assumed that $L0$ is attracted to the salient locations, and the salient locations are the same for both hiders and seekers (they called it *role-symmetric* $L0$).

In their model the $L0$ player chose the locations A, B, A, A with probabilities $\frac{p}{2}, q, 1-p-q, \frac{p}{2}$ respectively such that $p > \frac{1}{2}$ and $q > \frac{1}{4}$. With this $L0$ specification they get the results on Table 5.[1]

Table 5: Hide-and-seek games in Crawford and Iriberri (2007)

Type	$L0$	$L1$	$L2$	$L3$	$L4$
Frequency	0%	19%	32%	24%	25%

2.3.2 Penczynski (2011)

The main feature of Penczynski (2011)'s paper that he specified the $L0$ beliefs asymmetrically for hiders and seekers. Penczynski (2011) concentrated only on the salience related to B, and level-0 hiders were assumed to B-averse, level-0 seekers were attracted by B.

The played game was the original hide-and-seek game of Rubinstein, Tversky and Heller (1996). The subjects were divided into teams with two members, who were connected by the chat module of the experiment software, and they could send only one message to their team partner with their suggested decisions. After that, they had to state their final decision individually. Given that the messages were observable for the experimenter, this additional information was useful for the examination of reasoning

[1]Crawford and Iriberri (2007) claims that despite the high frequency of $L4$ subjects found, their frequency is not well identified in the distribution.

processes and revealed the salient boxes for the given student. Penczynski (2011) classified the messages according to the corresponding k-level.

The estimated level distribution were given by Table 6.

Table 6: Hide-and-seek games in Penczynski (2011)

Type	Hiders	Seekers	All
$L0$	0.35	0.15	0.27
$L1$	0.4	0.34	0.37
$L2$	0.17	0.4	0.27
$L3$	0.06	0.08	0.07
$L4$	0.02	0.04	0.03

2.4 The 11-20 money request game

The most recently developed game type which was used to level-k analysis is the 11-20 money request game introduced by Arad and Rubinstein (2012). This game type was designed to be the proper framework to study level-k behavior. The original description of the game from (Arad and Rubinstein, 2012, p. 3562):

"You and another player are playing a game in which each player requests an amount of money. The amount must be (an integer) between 11 and 20 shekels. Each player will receive the amount he requests. A player will receive an additional amount of 20 shekels if he asks for exactly one shekel less than the other player.

What amount of money would you request?"

There are several aspects of that game that make it ideal for the level-k study. First of all, there is a natural level-0 specification: choosing the 20 is clearly the most salient to a non-strategic player since it ensures the highest guaranteed amount. Second, given the Lk strategy, the best response of a $Lk+1$ type is straightforward: choosing the number which is smaller than the Lk request by 1. And in addition the type specifications are robust to the type distribution; for example, 19 is best response for a wide range of distributions where 20 is the most likely strategy, and this is true for the higher levels as well.

Putting everything together, it is natural to use the level-k reasoning model to this game. The requests 20, 19, ... correspond to $L0$, $L1$, ... types respectively. The relative frequency of requests are presented by Table 7.

Table 7: Results – 11-20 game of Arad and Rubinstein (2012)

Request	20	19	18	17	16	15	14	13	12	11
Frequency	6%	12%	30%	32%	6%	1%	6%	3%	0%	4%

Looking at the findings of the papers describing experiments with different types of games, we can see significant differences between the reported level-k type distributions (for instance consider the level-3 types were the rarest in the case of Beauty Contest games while they were the most common in the 11-20 game); and we can also recognize a similarity of the type distribution between similar games. In addition, there were no significant differences in terms of payoffs (the reward amounts were fairly similar) or in terms of subject pool (the subjects were undergraduates from economics, business our finance courses without any game theory study in practically every cases). Therefore it seems possible that the different level of the complexity of the different games can explain these differences.

3 Models of computational complexity

Computer scientists use the automata theory to study complexity (see for example Hopcroft and Ullman (1979)). The automata theory has already been applied in Economics and Game Theory to examine the complexity of choice processes or strategies. To the best of my knowledge all of these applications used the finite automaton (FA) model described later. Rubinstein (1986) and Abreu and Rubinstein (1988) used the FA model to analyze strategies in repeated games, while Salant (2011) used the it to study the complexity of procedural choice methods.

In all of these papers, the authors used the total number of needed states in the automaton as measure of complexity. This approach has a drawback, it is explicitly stated in (Abreu and Rubinstein, 1988, p. 1265): "The measure neglects the desire of players to simplify their calculations during the course of play." In other words, the measure does not reflect the computational costs occur during the play of the game. This drawback makes the FA model inappropriate for our purposes, since it does not make any difference in the costs related to different games. In the case of level-k analysis, the levels can be interpreted as states, and a FA which plays a level-k strategy needs exactly $k+1$ states to work regardless of the game type. Performing an analysis across game types is impossible with a model that cannot distinguish between game types.

Here I present a new automaton model which reflects the computation costs that occur during the rounds of iterations and define a complexity measure similar to the time complexity function defined in Hopcroft and Ullman (1979). In this section, first I introduce two models used in Computer Science based on Hopcroft and Ullman (1979): the finite automaton model (Subsection 3.1, this kind of model was used by Rubinstein (1986); Abreu and Rubinstein (1988) and Salant (2011)) and the Turing machine model with the time complexity function (Subsection 3.2). Subsection 3.3 presents the game machine model with the strategy complexity measure.

3.1 The finite automaton model

A finite automaton is a system with discrete inputs and outputs which can be in a finite number of different internal states. The behavior of the system is determined by the history of past inputs, and the states contain the necessary information about past inputs. I provide here the definition of an extension of the baseline FA model, the so-called Moore machine. The Moore machine differs from the baseline FA model in an important notion: the Moore machine has outputs as well, while the FA does not (or only binary outputs). The formal definition presented below are based on the definition in Hopcroft and Ullman (1979).

Definition 3.1. *A Moore machine is given by a 6-tuple* $(Q, \Sigma, \Delta, \delta, \lambda, q_0)$ *where* Q *is the finite set of states,* Σ *is the set of inputs,* Δ *is the set of outputs,* $\delta : Q \times \Sigma \to Q$ *is the transition function,* $\lambda : Q \to \Delta$ *is the output function and* $q_0 \in Q$ *is the initial state.*

I demonstrate the finite automata with the example of repeated game strategies discussed in Rubinstein (1986).

Example 3.2. Consider a baseline 2×2 Prisoner's Dilemma game with actions C(ooperate) and D(eviate) and look at the "Tit-for-Tat" strategy (start with C, then play C if and only if the other player played C in the previous period). In this case

- $Q = \{q_C, q_D\}$;
- $\Sigma = \{C, D\}$, the other player's action in the given period;
- $\Delta = \{C, D\}$, the player's action in the given period;

- $\delta(\cdot, i) = q_i$, where $i \in \{C, D\}$;

- $\lambda(q_i) = i$, where $i \in \{C, D\}$; and

- $q_0 = C$.

The automaton has two states, the transition map is represented by Figure 1. The circles denote the states, the arcs are the transitions, and the letters on the arcs denote the input (other player's action) that triggers the transition. The output function simply assigns C to q_c and D to q_D.

Figure 1: The transition mapping of the Tit-for-Tat automaton by Rubinstein (1986)

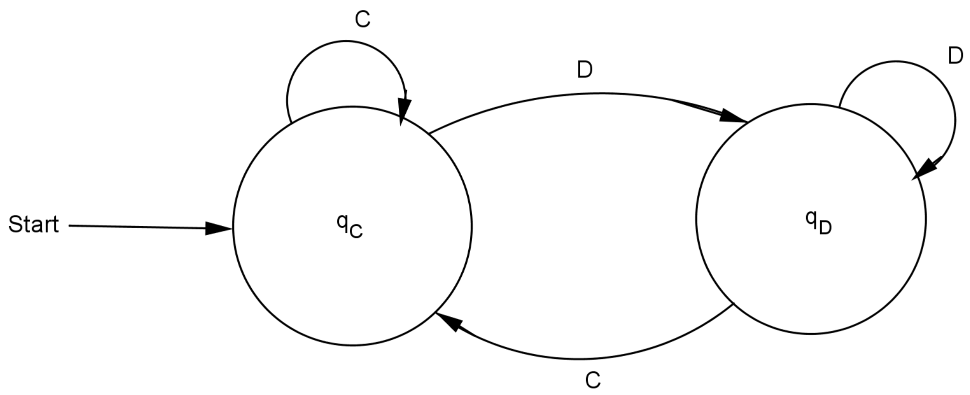

We can see from both of the definition and the example that the FA does not "compute" anything but just changes states according to the inputs, and cannot deduce anything from the information it gets. Therefore while it is a proper tool for analyzing strategies in repeated games and choices from a list, it is not able to express the computational cost of a level-k best reply in a game.[2]

[2]Hopcroft and Ullman (1979) do not define any complexity measure for finite automata

3.2 The Turing machine model

There are many versions of the FA model, but none of them are appropriate for the purposes of this thesis as the possible complexity measures defined on them does not capture the the difference of computational costs across games. However, a modified version of the time complexity function defined on a Turing machine is able to capture the computational complexity in question. Here I present the Turing machine and the time complexity function according to Hopcroft and Ullman (1979), but I do not include formal descriptions or details that are unnecessary for our purposes.

The Turing machine consists of an input tape divided into cells and a tape head that is able to scan the inputs on the tape, one cell at a time. The set of the input symbols is finite. Contrary to the FA, the Turing machine (TM) is able to print out symbols depending on the state of the machine and the inputs. This feature makes possible the computation. Another important difference is that the Turing machine makes *moves* not just changes states as the FA. In one move the Turing machine can change state, print out an output symbol and shift its head left or right one cell.

There are lots of modified versions of the baseline TM model, such as multitape machines. One can construct Turing machines for a given procedure (for example, computing a function). The machine scans the input symbols, changes states and prints out output symbols and halts when the procedure is finished. Hopcroft and Ullman (1979) provide examples of this kind of computing TM. If a reader wants to imagine such a TM, for example, she can think of a TM with binary input and output symbols that has an input tape with m and n zero symbols separated with an one symbol such that $m > n$, and it gives an output of $m - n$ zeros when it halts.

Machines can be classified by computational complexity regarding the amount of time, space or other resources the machine uses for its task. There are several possible measures; here we focus on the time complexity measure. The time complexity measure $T(n)$ is the maximum number of moves the machine makes before halting for every input with n symbols.

The TM is a very general tool which fits for all computational problem. However, the construction of a Turing machine for a given procedure can be very cumbersome. To avoid these complexities, I construct a new model for analyzing the computational complexity of level-k best responses in games which is similar to the Turing machine concept. This model is described in the next subsection.

3.3 The game machine model

The game machine model receives a game as an input and gives a played strategy as an output. The input is a description of a game including the set of players, the set of strategies and the payoff functions. The played strategy is always a strategy played by a level-k type player. The machine (similarly to the FA and the TM) has states as well, the state q_k refers to the level-k strategy, that is, the state determines the output strategy. The strategy played in q_0 is also specified by the machine.

The game machine can also make moves conditional on its input (the game played); the purpose of these moves is to compute the next level best response to the strategy played in the current state. After each move, the machine prints out what is in its 'memory': a strategy, a computed number etc. When the machine prints out the strategy that best responds to the one played in the current state, then its moves to the next state.

At a given state q_k a *move* the machine can perform exactly one step of the following:

- Recognize that the other players also play the same strategy at q_k;

- Perform an algebraic calculation (computing a product or an average, deciding which is the smallest/largest number in a given set etc.); or

- Choose a strategy from the strategy set.

The computational complexity $C(k)$ is the minimum number of required moves to reach state q_k (it is straightforward that this number depends on the game).

It is true that the moves described above may not require the same mental efforts (e. g. it can be easier to compute the product of two numbers than recognizing that the other players think the same), however, since it is very hard to specify the relative cognitive cost of these moves, I maintain the discrete structure of the assumed mental process without distinguishing among the different discrete steps. The model is not perfect, but captures more of the computational complexity than any of the models previously used in game theory.

Another drawback of the model is that it ignores the increasing return to scale of thinking, which is intuitively possible. The model can be extended in a way to taking this increasing return to scale into account; however the same specification problem will occur that we have during the comparison the different moves (if there is increasing return to scale, it is difficult to specify how large is that exactly).

The formal definition of the game machine is given below.

Definition 3.3. *A game machine G is given by a 4-tuple $\langle \Gamma, Q, f, s_i \rangle$, where $\Gamma = (N, S_i, u_i)$ contains the description of the input game: N is the set of players; S_i is the set of strategies for player $i \in N$; u_i is the payoff function of player $i \in N$; Q is the set of states, each of them represents a level-k type; $f : \Gamma \times Q \to Q$ is the transition function using the moves described above; and $s_i : Q \to S_i$ is the output function which chooses the played strategy for all players $i \in N$ in a given state $q \in Q$.*

I illustrate the game machine with an example of the Beauty Contest game with n players, interval $[a, b]$ and parameter $p < 1$.

Example 3.4. We have a game machine $G = \langle \Gamma, Q, f, s_i \rangle$, where

- $\Gamma = (N, S_i, u_i)$; $N = \{1, ..., n\}$; S_i is the numbers in $[a, b]$ for all $i \in N$; u_i is positive if s_i is the closest to $p \cdot \sum_{j \in N} s_j / n$ and zero otherwise,
- $Q = \{q_0, q_1, ...\}$,
- $s_i(q_0) = x$ and $s_i(q_k) = p^k E(x)$ for all $i \in N$ and $k > 0$ where x is a random variable with uniform distribution on $[a, b]$.

The transition mapping is presented in Figure 2. The circles are the states, the arcs represent moves and the squares contain the numbers printed out by the machine after a given move. Note that here the strategies are numbers, but that is not true for all types of games (i. e. the move of choosing the *strategy* $pE(x)$ after computing and printing out the *number* $pE(x)$ is not redundant).

Figure 2: The transition mapping of the game machine playing a Beauty Contest game

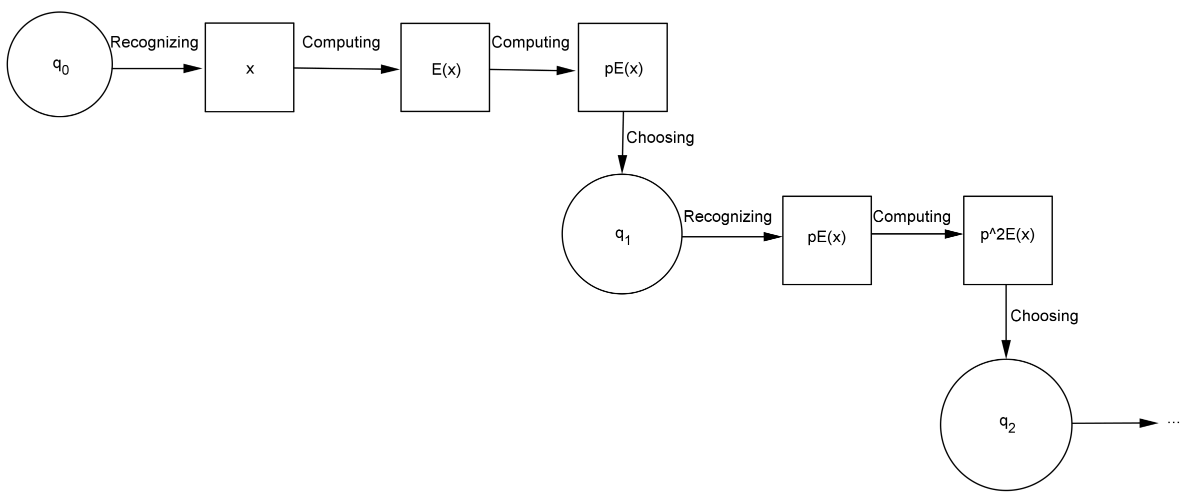

4 Analysis of the computational complexity in games

In this section I use the game machine model introduced in Subsection 3.3 to evaluate the games described in Section 2 in terms of computational complexity. I examine how much moves were required to calculate the different level-k strategies and for all of the experiments I provide a table which contains the $C(k)$ computational costs for each level-k type that was identified by the authors of the given experimental analysis. In this section I report only the complexity values, I discuss in Section 5.

4.1 Beauty Contest games

The Beauty Contest game played in Nagel (1995) and Ho, Camerer and Wiegelt (1998) is exactly the same baseline Beauty Contest game in Example 3.4. Therefore the $C(k)$ values can be computed according to the moves in Figure 2. In the experiment of Costa-Gomes and Crawford (2006) there was an additional step in the reasoning process at the first level: the players had to compute the overall target which is the product of the two individual targets. Table 8 presents the results.

Table 8: $C(k)$ in the Beauty Contest game

	0	**1**	**2**	**3**
Baseline	0	4	7	10
Costa-Gomes and Crawford (2006)	0	5	8	11

4.2 Two-person matrix games

In Stahl and Wilson (1994) the players play 3×3 matrix games. The level-0 type plays a strategy randomly, a level-1 strategy best responds to a random play. Therefore, the level-1 player computes three averages for her three strategies, decides which one is the highest then picks a strategy for which this average is the highest. So the game machine makes five steps to choose the level-1 strategy (one for reasoning about the other player, one for each average and one for choosing a strategy). The level-2 player only has to reason about the other player's choice and pick a strategy best responds to that comparing three numbers with each other (three steps for the machine: reason, compare, choose). That is the same three additional moves to the level-3 strategy.

In Costa-Gomes, Crawford and Broseta (2001) the design was somewhat different. Here the players played 4 2×2, 12 2×3 and 2 2×4 matrix games. While all computations above level 1 are the same as in the 3×3 case, the computation of the level-1 strategy is different with different number of strategies. In the 2×2 case the players only had to compute two averages instead of three, so four moves were needed for the level-1 strategy. Similarly, in the other two cases on average 4.5 and 5 steps were needed to the level-1 strategy in the 2×3 and 2×4 cases respectively. So on average the computation of level-1 needed 4.45 moves from a game machine.

Table 9 represents the $C(k)$ costs for the different levels identified by the authors.

Table 9: $C(k)$ in two-person matrix games

	0	**1**	**2**	**3**
Stahl and Wilson (1994)	0	5	8	11
Costa-Gomes, Crawford and Broseta (2001)	0	4.45	7.45	10.45

4.3 Hide-and-seek games

In the case of hide-and-seek games there are natural best replies for each strategies. The hiders want to pick a different box they think that seekers will choose and the seekers want to pick same as the seekers. In the role symmetric specification of Crawford and Iriberri

(2007), therefore, reaching each of the states need only two moves from the previous state: the first of them is reasoning about the other player's choice, and the second is picking a box according to that.

In the role-asymmetric specification of Penczynski (2011), when hiders and seekers have different level-0 strategies, the players have to think differently about the game than in Crawford and Iriberri (2007). In order to reach the first level, instead of the two moves, the players need three. If a player recognizes that the other player is also level-0, unlike in the other games (due to role asymmetry) an additional move is needed to recognize that the other player's level-0 strategy is different from her. Similarly, at higher states q_k after recognizing that the other player also plays a level-k strategy, an additional move is needed to identify the opponent's level-k strategy and another to find the best response in that. Therefore with role asymmetry three moves are needed for each level. Note that this increase in the computation cost can explain the striking difference between the results of Crawford and Iriberri (2007) and Penczynski (2011). Despite that the players played a similar game, with the assumed role asymmetry Penczynski (2011) found much lower average k. If the players also assumed the role asymmetry, they faced higher computation costs. Table 10 contains the $C(k)$ values.

Table 10: $C(k)$ in hide-and-seek games with and without role symmetry

	0	**1**	**2**	**3**	**4**
Crawford and Iriberri (2007)	0	2	4	6	8
Penczynski (2011)	0	3	6	9	12

4.4 11-20 game

Similarly to the hide-and-seek game in Crawford and Iriberri (2007), the 11-20 game also has a natural best reply for each level-k strategy. Therefore in this game type also only two moves needed to reach a higher state. Table 11 reports the number of game machine moves needed for each level.

Table 11: $C(k)$ the 11-20 game

0	**1**	**2**	**3**	**4**
0	2	4	6	8

Now we have all the data (both observed frequencies of level-k types and the computational costs related to them) to analyze the results across games.

5 Overall analysis across games

While the data I work with is extremely noisy due to the different econometric identifying approaches used by the authors of the papers reviewed in Section 2, I am still able to conduct some simple analyses. In this section I present them.

For instance, we can look at the average k across different games with different cognitive costs. I computed the average k for each experimental result and the average cognitive cost related to a given game (I summed the cognitive costs of the level-k strategies which were played in that experiment and divided that sum with the number of levels identified). A scatter plot with regression line is in Figure 3. The clear negative relationship reveals that the game machine approach can make sense: more complex the game is, on average fewer reasoning steps are performed by the players.

Figure 3: Average k and average cognitive costs across games

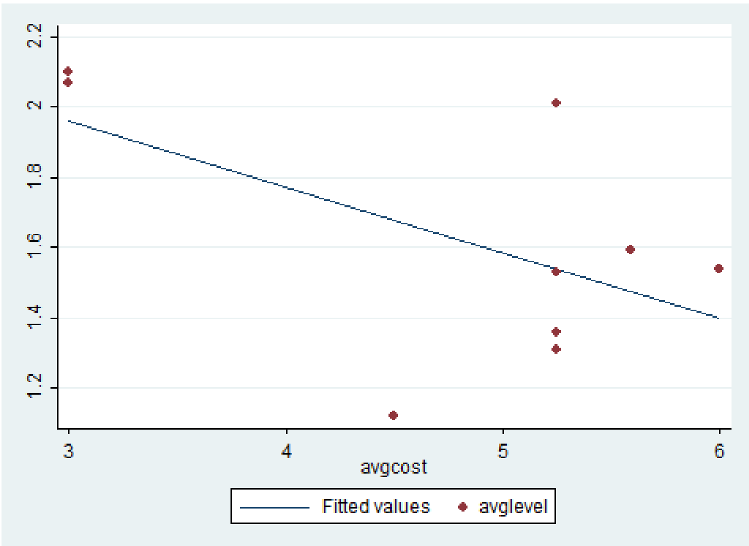

There is another possibility for examining the predictions of the game machine model. Consider the cost-benefit framework of reasoning process according to Alaoui and Penta (2013a). Since the rewards are the same for all player, the benefits are identical for each players. At the cost side there is a cognitive cost of thinking for each player, and that cost is determined by two factors. One of them is the player's cognitive ability, and the other is the complexity of the task the player faces.

Look at a given player i. The player's cognitive ability is described by an increasing and convex cognitive cost function κ_i which is defined on the complexity of the task player i faces. Thus, if the player has to calculate a level-k strategy in a given game, then the complexity of the task she faces is $C(k)$; therefore she faces a cognitive cost of $\kappa_i(C(k))$. When playing a game, the player chooses the strategy for which her marginal cost of thinking equals to the marginal benefit of thinking. More precisely, since the thinking is assumed to be a stepwise procedure and $C(k)$ has discrete values, the player will perform exactly k^* steps where k^* is the highest level when the marginal cost of thinking does not greater than the marginal benefit (the marginal benefit is assumed to be constant). Of course, different players can have different κ_i cognitive costs.

Consider a situation when a group of players play the same game and have the same expected rewards (a game experiment is exactly a situation like that). Since the benefits are the same for each players, the depth of reasoning is determined by the player's cognitive costs. If the thinking is more costly to a player then she will perform fewer steps of thinking, therefore will play a lower level-k strategy. From the other direction, if we have the distribution of level-k strategies from the results of an experiment, we can approximate the distribution of players regarding their cognitive costs.

Take the results of Arad and Rubinstein (2012) which are the most reliable due to the natural application of level-k approach and it has the finest level-cost structure (the differences between the levels are the smallest in this game type) we can approximate the cognitive ability distribution of the subjects.

The benefits are identical to each player, assuming that the marginal benefit is constant and it is equal to x. We know that computing the level-k strategy has exactly $2k$ in the case of the 11-20 game. The players for which $\kappa'(2) > x$ will play the level-0 strategy, the players for which $\kappa'(2) \leq x < \kappa'(4)$ will play the level-1 strategy and so on. If we know from the results that 6% of the players played the level-0 strategy, then it implies that for the 6% of the players do not worth to make two "game machine moves" since the cognitive cost is to high for that players. Using the results of Arad and Rubinstein (2012), the players can be classified by the number of the maximum moves they willing to perform for the rewards. The players can be arranged into intervals of length 2 since 2 is the difference between the complexity of the levels. The frequency of players in these intervals can be seen on the Table 12.

Table 12: Approximated cognitive ability distribution using the results of Arad and Rubinstein (2012)

Maximum moves	[0,2]	(2,4]	(4,6]	(6,8]	(8,10]	(10,∞)
Frequency	6%	12%	30%	32%	6%	14 %

From now we assume that the benefit side and the distribution of players' cognitive abilities identical in all experiments I presented in Section 2 (it is not overly restrictive since the rewards and the educational background of the subject pools were similar). Now the only parameters that can cause variation in the level-k distribution is the difference in the complexity of the games.

Using the distribution from Table 12 and assuming this is the same for the subject pools in the other experiments, we can predict the level-k distribution of the other experiments as well.

Table 13 presents the predicted values for each experiment I analyzed with the observed values in the parentheses. The numbers in the top row refer to Nagel (1995), Costa-Gomes and Crawford (2006), Stahl and Wilson (1994), Costa-Gomes, Crawford and Broseta (2001), Crawford and Iriberri (2007) and Penczynski (2011) respectively. (I assumed uniform distribution within the intervals.)

Table 13: Predicted and observed distributions

Experiment	(1)	(2)	(3)	(4)	(5)	(6)
$L0$	18(13)	33(0)	33(0)	25(7)	6(0)	18(27)
$L1$	46(44)	47(46)	47(24)	46(24)	12(25)	36(37)
$L2$	22(33)	13(28)	13(51)	16(59)	30(43)	36(27)
$L3$	14(10)	7(7)	7(25)	13(18)	32(32)	10(7)

We can see from the table that the predictions are somewhat good for the Nagel (1995), Crawford and Iriberri (2007) and Penczynski (2011) experiments, so for the baseline Beauty Contest and the hide-and-seek games. There are major differences between the actual and predicted frequencies for the Costa-Gomes and Crawford (2006) version of Beauty Contest games and for the normal-form games, especially at the lower levels. Overall, there is a positive correlation between the predicted and actual levels (the correlation coefficient is 0.26).

These two descriptions presented above show that the game machine model can explain some fraction of the variances in the level-k type distribution estimated in previous experiments. Of course, it is not so strong support to the model due to the noisiness I mentioned before. For a precise testing of the model we need to conduct a specific experiment. I give some recommendations for that kind of experiment in the concluding section.

6 Conclusion and recommendations for an experimental analysis

In this thesis I investigated the open question of predicting endogenously the differences in the depth of reasoning performed by players across different games. Using the cost-benefit framework introduced by Alaoui and Penta (2013a,b), I focused on the cost side of the analysis and classified the level-k strategies of the different games by the computational complexity of the given strategy. I used a modified version of a Turing machine to evaluate the computational complexity of the played strategies. The model was able to explain a part of the variation observed in the level-k distribution by authors of experimental papers.

This model is the first step made in this research area and it is somewhat rudimentary. One can improve it by incorporating the increasing return to scale in thinking and the distinction between the cognitive cost of the different type of moves of the machine. However, as I mentioned I do not see a straightforward way to these extensions.

Although I tested the model using the data from previous experiments, a proper test would require conducting its own experiment. The subjects must have identical educational background and in all games the expected reward must be the same to ensure that the differences in the level-k distributions would come from the difference in the games. I recommend dividing the subjects into (at least) eight different sessions. Two sessions would play each of the game types (Beauty Contest game, matrix games, hide-and-seek game and the 11-20 game) with different amount of time available to choose strategies. The reasoning is assumed to be stepwise, therefore with more time available we can expect higher number of performed reasoning steps. The outcomes of an experiment like this can reliably test the game machine model and can shed light to its strong and weak points.

References

Abreu, Dilip, and Ariel Rubinstein. 1988. "The Structure of Nash Equlibrium in Repeated Games with Finite Automata." *Econometrica*, 56: 1259–1281.

Alaoui, Larbi, and Antonio Penta. 2013*a*. "Depth of Reasoning and Incentives." http://cess.nyu.edu/wp-content/uploads/2013/04/alaoui_penta_levelk-2.pdf.

Alaoui, Larbi, and Antonio Penta. 2013*b*. "Strategic Thinking and Value of Reasoning: Theory and Applications to 'Five Little Treasures' of Game Theory." http://www.econ.wisc.edu/~apenta/value%20of%20reasoning_ap.pdf.

Arad, Ayala, and Ariel Rubinstein. 2012. "The 11-20 Money Request Game: A Level-*k* Study." *American Economic Review*, 102: 3561–3573.

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

Costa-Gomes, Miguel, and Vincent P Crawford. 2006. "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study." *American Economic Review*, 96: 1737–1768.

Costa-Gomes, Miguel, Vincent P Crawford, and Bruno Broseta. 2001. "Cognition and Behavior in Normal-Form Games: An Experimental Study." *Econometrica*, 69: 1193–1235.

Crawford, Vincent P, and Nagore Iriberri. 2007. "Fatal Attraction: Salience, Naiveté, and Sophistication in Experimental 'Hide-and-Seek' Games." *American Economic Review*, 97: 1731–1750.

Crawford, Vincent P, Miguel A Costa-Gomes, and Nagore Iriberri. 2013. "Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications." *Journal of Economic Literature*, 51. Forthcoming.

Heifetz, Aviad, and Willemien Kets. 2012. "All Types Naive and Canny." http://www.kellogg.northwestern.edu/research/math/papers/1550.pdf.

Hopcroft, John E, and Jeffrey D. Ullman. 1979. *Introduction to Automata Theory, Languages and Computation.* Addison-Wesley.

Ho, Teck-Hua, Colin Camerer, and Keith Wiegelt. 1998. "Iterated Dominance and Iterated Best Response in Experimental 'p-Beauty Contests'." *American Economic Review*, 88: 947–969.

Kets, Willemien. 2012. "Bounded reasoning and higher-order uncertainty." http://www.econstor.eu/bitstream/10419/59686/1/715698427.pdf.

Keynes, John Maynard. 1936. *The General Theory of Employment, Interest and Money.* Harcourt Brace and Co.

McKelvey, Richard D, and Thomas R. Palfrey. 1995. "Quantal Response Equilibria for Normal-Form Games." *Games and Economic Behavior*, 10: 6–38.

Nagel, Rosemarie. 1995. "Unraveling in Guessing Games: An Experimental Study." *American Economic Review*, 85: 1313–1326.

Penczynski, Stefan P. 2011. "Strategic Thinking: The Influence of the Game." http://www.penczynski.de/attach/Focality.pdf.

Pintér, Miklós, and Zsolt Udvari. 2012. "Generalized type spaces." http://unipub.lib.uni-corvinus.hu/623/1/PinterM_UdvariZs_2012.pdf.

Rubinstein, Ariel. 1986. "Finite Automata Play the Repeated Prisoner's Dilemma." *Journal of Economic Thoery*, 39: 83–96.

Rubinstein, Ariel, Amos Tversky, and Dana Heller. 1996. "Naive Strategies in Competitive Games." In *Understanding Strategic Interaction - Essays in Honor of Reinhard Selten*. 394–402. Springer-Verlag. ed. Wulf Albers, Werner Güth, Peter Hammerstein, Bemmy Moldovanu and Eric van Damme.

Salant, Yuval. 2011. "Procedural Analysis of Choice Rules with Applications to Bounded Rationality." *Aconomic Economic Review*, 101: 724–748.

Stahl, Dale O, and Paul R. Wilson. 1994. "Experimental Evidence on Players' Models of Other Players." *Journal of Economic Behavior and Organization*, 25: 309–327.

Stahl, Dale O, and Paul R. Wilson. 1995. "On Players' Model of Other players: Theory and Experimental Evidence." *Games and Economic Behavior*, 10: 218–254.

Strzalecki, Tomasz. 2011. "Depth of Reasoning and Higher-Order Beliefs." http://scholar.harvard.edu/files/tomasz/files/strzalecki-depth.pdf.