Department of Economics
Working Paper No. 211

Cognitive Hierarchies in the Minimizer Game

Ulrich Berger
Hannelore De Silva
Gerlinde Fellner-Röhling

January 2016

WU
WIRTSCHAFTS
UNIVERSITÄT
WIEN VIENNA
UNIVERSITY OF
ECONOMICS
AND BUSINESS

EFMD
EQUIS
ACCREDITED

Cognitive Hierarchies in the Minimizer Game

Ulrich Berger[*], Hannelore De Silva[†], Gerlinde Fellner-Röhling[‡]

June 12, 2015

Abstract

Experimental tests of choice predictions in one-shot games show only little support for Nash equilibrium (NE). Poisson Cognitive Hierarchy (PCH) and level-k (LK) are behavioral models of the thinking-steps variety where subjects differ in the number of levels of iterated reasoning they perform. Camerer et al. (2004) claim that substituting the Poisson parameter $\tau = 1.5$ yields a parameter-free PCH model (pfPCH) which predicts experimental data considerably better than NE. We design a new multi-person game, the *Minimizer Game*, as a testbed to compare initial choice predictions of NE, pfPCH and LK. Data obtained from two large-scale online experiments strongly reject NE and LK, but are well in line with the point prediction of pfPCH.

JEL classification: C72; C90; D01; D83

Keywords: behavioral game theory; experimental games; Poisson cognitive hierarchy; level-k model; minimizer game

---

[*]Institute for Analytical Economics, WU Vienna University of Economics and Business (ulrich.berger@wu.ac.at)

[†]Institute for Finance, Banking and Insurance, WU Vienna University of Economics and Business (hannelore.desilva@wu.ac.at)

[‡]Institute of Economics, Ulm University (gerlinde.fellner@uni-ulm.de)

1 Introduction

1.1 Nash equilibrium in one-shot games

Nash equilibrium (NE) is the central solution concept in noncooperative game theory, but it is well known that NE imposes extremely demanding assumptions on players' rationality as well as on the consistency of their beliefs (Aumann and Brandenburger, 1995). Experimental tests of game theoretic strategy choice predictions in one-shot games, where players are inexperienced and have no possibility of learning, have consistently shown only little support for NE except in rather specific games (Camerer, 2003). Game theorists have therefore worked out a variety of alternative explanations and models for prediction of choices in one-shot games. One such approach is to keep the rationality assumption and weaken the mutual consistency requirements. Any concept following this approach must address the question of how players form their beliefs about other players' actions. This has led to the formulation of so-called thinking steps models.

1.2 Level-k models

The basic assumption of the thinking steps approach is that players differ in the number of steps of iterated reasoning they apply when deliberating which strategy they should pick in a strategic choice problem. Nagel (1995) used a simple thinking steps model to explain the results of her experiments about number choices in the p-beauty-contest game (or p-guessing game), where players choose a number from $[0, 100]$ and whoever comes closest to p times the average of the chosen numbers wins a fixed prize. Her model explains the "spikes" around choices of 33 and 22, which are often observed in experimental data for the typical parameter $p = 2/3$. Similar thinking steps models have been proposed by Stahl and Wilson (1994, 1995) and by Ho et al. (1998).

The most prominent model of the thinking steps variety is the *level-k model* (LK-model) introduced by Costa-Gomes et al. (2001). It proposes that most players can be classified as level-k (lk) types, which anchor their beliefs in an $l0$ type who does not think strategically at all but just chooses from a uniform random distribution on the set of pure strategies.[1] lk then simply best responds (possibly with noise) to $l(k-1)$. These lk types are complemented by types dk, who best respond to a uniform distribution of beliefs on strategies surviving k rounds of iterated dominance, respectively. Finally, some players might be *equilibrium* types, choosing an equilibrium strategy, or *sophisticated*, best responding to an accurate distribution of beliefs on other types. LK models have been applied in econometric analyses of various experimental data by Costa-Gomes and Crawford (2006), Crawford and Iriberri (2007a, 2007b), and Costa-Gomes et al. (2009), among others. These studies spawned a large body of literature (see the recent review by Crawford et al., 2013). By and large, the common view that emerged from this literature is that $l1$ and $l2$ types are predominant in subject populations, complemented by smaller fractions of $l3$ and possibly $l4$, $d1$ and equilibrium types. $l0$ as well as $l5$ or higher, $d2$ or higher, and sophisticated types, however, are virtually absent from the population.

1.3 The Poisson Cognitive Hierarchy model

A closely related thinking steps model is the *Poisson Cognitive Hierarchy model* (PCH-model) of Camerer et al. (2004). This model uses only a single parameter, τ. It is based on the view that players differ in their level k of iterated thinking, and that k is distributed in the population of players following

---

[1] A vast majority of applications uses this specification of $l0$ behavior. Alternatively $l0$ has also been suggested to choose the most salient strategy in games with non-neutral frames (Crawford and Iriberri, 2007b), but this approach is not without problems itself (Heap et al., 2104). Burchardi and Penczynski (2014) use an innovative experimental design to identify $l0$ reasoning in beauty-contest games. For a recent systematic approach to $l0$ behavior see Wright and Leyton-Brown (2014a).

a Poisson distribution with mean (and therefore variance) τ. Moreover, while a level-k (Lk) player thinks that all other players do less steps of reasoning than he himself, he is aware of the presence of all levels of reasoning from 0 to $k-1$ in the population. The frequency he believes these lower levels to occur are the true (Poisson) frequencies, truncated at $k-1$ and normalized so as to add up to 1.

The PCH-model has been shown to predict reasonably well in a variety of games, among them p-beauty-contest games with $p < 1$, market entry games, 3×3 bimatrix games (Camerer et al., 2004), coordination games (Costa-Gomes et al., 2009), and the action commitment game (Carvalho and Santos-Pinto, 2014). While, as expected, the best-fitting value of τ is game- and population-specific, Camerer et al. (2004) report that a value of $\tau = 1.5$, corresponding to a population dominated by $L1$ and $L2$ types, is able to explain experimental data considerably better than Nash equilibrium across a variety of experimental games.

While the PCH-model is simple to apply and has proven useful in a number of games, it also seems to fail in some specific classes of games. For example, it is well known that in Prisoners Dilemma games and in Public Good games initial cooperation levels are substantial (see Camerer, 2003). Such behavior cannot be explained by a PCH model, since there all types but $L0$ optimize and hence never choose dominated strategies. For the same reason, PCH cannot account for the puzzling majority choices of dominated strategies in the two-person beauty contest of Grosskopf and Nagel (2008). As Camerer et al. (2004) report, the PCH-model also predicts almost random choice in p-beauty-contest games with $p > 1$. Another problem arises in games with large strategy spaces. Camerer et al. (2002) mention that in such games PCH predicts only a small fraction of the strategies actually chosen. An example for this is found in Gneezy (2005), where prior to grouping the data PCH cannot account for bid choices in a first-price auction with 100 pure strategies.

4

Camerer et al. (2002) devote section 4.1 of their working paper to investigate what went wrong in games 2, 6, and 8 of the 12 games of Stahl and Wilson (1995). These are symmetric 3×3-games where the best-fitting PCH-model is $\tau = 0$, predicting purely random choice. But actually roughly half of the subjects picked their Nash equilibrium strategy in these three games. Camerer et al. (2002) speculate that the experimental procedure of Stahl and Wilson may have catalyzed a large fraction of Nash play. However, a more parsimonious explanation derives from the observation that in these three games the Nash equilibrium strategy also happened to be the unique maximin choice. Maximin choices actually predict the majority choices in ten out of the 12 Stahl-Wilson games. In games where strategic thinking is cognitively demanding, the nonstrategic and risk-averse maximin choice is an easy option and might often be a more salient anchor than uniform randomization for level-0 types (see also Van Huyck et al., 1991).

1.4 Predicting choice with the parameter-free Poisson Cognitive Hierarchy model

As noted by Wright and Leyton-Brown (2014b), the bulk of the literature on thinking steps models is concerned more with explaining than with predicting behavior. Typically, type distributions and other parameters are estimated from experimental training data while direct prediction performance comparisons are rare. In this paper we focus on prediction in a very strict sense. What we aim at is the prediction of initial choices without any prior parameter estimation. For this we need a parameter-free model which can be directly pitted against Nash equilibrium.

To our knowledge, within the thinking steps variety the only parameter-free model to be found in the literature is the PCH model which results from substituting $\tau = 1.5$ for the Poisson parameter, suggested by Camerer et al. (2004). In their abstract they state that *[a]n average of 1.5 steps fits data from many games* (p. 861); they note that *values of τ between 1 and 2 explain*

5

empirical results for nearly 100 games, suggesting that assuming a τ value of 1.5 could give reliable predictions for many other games as well (p. 863) and that the data *suggest that the Poisson-CH model with $\tau = 1.5$ can be used to reliably predict behaviors in new games*. In their conclusion, Camerer et al. (2004) stress that *[t]he value $\tau = 1.5$ is a good omnibus guess which makes the Poisson-CH theory parameter-free and is very likely to predict as accurately as Nash equilibrium, or more accurately, in one-shot games* (p. 890).

We call the PCH model with $\tau = 1.5$ the *parameter-free PCH-model* (pfPCH-model). The pfPCH-model states that the population wide frequencies of levels $L0$, $L1$, $L2$, and $L3$ are given by 22.3%, 33.5%, 25.1%, and 12.6%, respectively, with only 6.6% accruing from levels 4 and higher. The value $\tau = 1.5$ for the mean (and the variance) of the number of thinking-steps is based on experimental data from various games scrutinized by Camerer at al. (2004). Does this value also predict reasonably out of sample, i.e. in games beyond the classes of games it was derived from? To evaluate the predictive performance of the pfPCH-model we use it to predict the distribution of initial choices in a new game, the *minimizer game* described below, which we motivate and construct specifically for this purpose.

We do not only compare the pfPCH-prediction to the Nash prediction, but also to the LK-model's prediction. The LK-model is not parameter-free and we are not aware of any suggestions for a "good omnibus guess" of LK type frequencies from the literature. We therefore take a generous approach and allow for *all* type distributions of the LK-model. We find that in our minimizer game the pfPCH prediction easily outperforms both Nash equilibrium and all LK-model specifications. Indeed, the pfPCH prediction is not even statistically significantly different from the distribution observed in the experiment.

2 The Minimizer Game

How well does the pfPCH-model predict players' initial response to a strategic choice problem? We planned to answer this question experimentally and started by asking what kind of game would be appropriate for an experimental test of this question. In our opinion, three issues had to be considered:

- Experimental outcomes in Prisoners Dilemma games, Public Goods games, Dictator games, Ultimatum games and the like are strongly influenced by the presence of altruistic motives, fairness considerations, or other social preferences. These preferences "contaminate" the experimental results, since they may override the strategic incentives created by the monetary payoffs.[2] We should therefore avoid games where choices are sensitive to the presence of social preferences. This basically rules out almost all two-player games. It seems therefore wise to look for a multi-person game where other-regarding preferences are unlikely to influence choices.

- Since the premise of the thinking steps models is that given their beliefs, players optimize, we should choose a *simple* game for our test. If due to computational complexity subjects get the arithmetics wrong when optimizing, choices will be biased even if the PCH-model accurately describes belief formation and the distribution of thinking steps. It is to a large extent a matter of taste what kind of game to deem simple. However, bearing in mind Grosskopf and Nagel's (2008) stunning results for two-person beauty contest games, where even a majority of professionals failed to realize that 0 is a dominant strategy, we would strongly opt for avoidance of the need of any arithmetics having to be done by subjects trying to optimize. Moreover, simplicity seems to

---

[2]Wright and Leyton-Brown (2014a) find that the feature of *fairness* of an action is especially prone to influence level-0 behavior.

require a rather small set of possible choices, since with a larger number of choices both formation of beliefs and optimization given beliefs become computationally demanding.

- A third aspect to be aware of are attitudes towards risk. Strategic uncertainty may lead subjects with different risk attitudes to different choices, which could then falsely be attributed to different levels of iterated reasoning. As discussed above, nonstrategic and strongly risk-averse subjects might tend to choose their maximin strategy in the face of strategic uncertainty, again biasing the distribution of thinking steps. Ideally we would therefore construct an experimental game where maximin does not restrict the set of available choices.

2.1 Definition of the minimizer game

Considering these three points we specifically designed a game for our experiments which has to the best of our knowledge not been studied before. For reasons which will become clear in a moment, we call this game the *minimizer game* (MG). The formal definition of an MG is the following.

Let $P = \{1, \ldots, I\}$ be a finite set of players, let N be a finite subset of $\mathbb{N}$, and let $n = |N|$ be the size of N. Players' strategy sets are $S_i = N$ for all i, so each player chooses a number from the set N. For a pure strategy profile $s = (s_i)_{i \in I}$ and for $k \in N$ let $c_k(s) = |\{i \in P : s_i = k\}|$ count the number of players choosing k in the profile s. Let $M(s) = argmin_{k \in N}\{c_k(s) : c_k(s) \geq 1\}$ be the set containing the numbers which have been chosen least often among those chosen at all in profile s. Let $m(s) = \min M(s)$ be the smallest of these numbers. The payoff function is identical for all players and is given by

$u_i(s) = m(s)\, c_{m(s)}^{-1}(s)$ if $s_i = m(s)$ and $u_i(s) = 0$ else.

Despite its technically sounding definition, it is easy to explain the MG in an extremely simple and intuitive way. The rules of the game state that each subject may choose its desired payoff from a given set of possible integer

8

payoffs. The 'winning amount' is the amount chosen least often in total - the *minimizer*. Among all players who chose the minimizer, one player is randomly drawn to receive this amount, while all others receive zero.[3] Ties are broken by declaring the smallest of the least often chosen amounts to be the minimizer.

While this is neither necessary nor specified in the definition, when it comes to experiments we implicitly think of the MG as being played with a small number of possible choices and a large number of players. The reasons for this are explained in the section on equilibrium analysis of the MG.

2.2   Advantages of the MG

The MG appears markedly dissimilar from strategic choice situations which are frequently encountered in everyday life. It may thus be considered artificial, but we consider this an explicit advantage. The reason is that we focus on choices in truly one-shot games, and for this we have to make it unlikely that subjects can transfer experiences from related games they "played" in the past. We think that this is the case for the MG. While the MG has many features of a congestion game,[4] its peculiar rules make it very unlike the "typical" congestion games people unconsciously play, like e.g. choosing the fastest road to their office in the morning.

Let us now reconsider our list of three points.

---

[3]An even simpler, deterministic variant of the MG lets *all* players having chosen the minimizer receive this amount, i.e. $u_i(s) = m(s)$ if $s_i = m(s)$ and $u_i(s) = 0$ else. While this variant was the first we came up with, it is technically almost infeasible in our large-population online experimental approach, which is why we proceeded to work with the stochastic variant of the MG described here.

[4]The MG should not be confused with the superficially similar *minority game*, which builds upon the *El Farol bar problem* (Arthur, 1994) and has been studied intensively in the statistical physics literature. It is related, but not identical to the *LUPI Lottery game* (Östling et al., 2011) either. Note that in the LUPI game, contrary to the MG, the prize for the winner is independent of the winning number. Moreover, if there is no uniquely chosen number, then all payoffs are zero in the LUPI game. Thus, the LUPI game is interesting if there are many more available choices than players, while we study the MG in the exactly opposite case.

- Social preferences do not interfere in the MG, at least if the number of players I is large. Since the influence of one's own choice on the overall minimizer is typically negligible, there is no obvious way social preferences, even if present, could influence choice behavior.

- When it comes to optimizing, the MG is extremely simple. Given a belief about the individual distribution of number choices of other players, maximization of expected payoff is straightforward if I is large: Choose the number which has the lowest frequency. As opposed to beauty contest games, bimatrix games, or first-price auction games, this does not require subjects to perform any arithmetical operations.

- In the MG, any choice might result in a zero payoff in the worst case. Therefore maximin has no bite in this game.

Taken together, these advantages indicate that the MG (with large I and small n) is better suited than known experimental games to test the pfPCH-model in a "purified" context.[5] Note, however, that we do not bias our test in favor of pfPCH. In principle, eliminating the three confounders could work for or against the pfPCH-prediction.

2.3 Nash equilibria of the MG

Consider a MG with $n \geq 2$ and let s be a pure-strategy profile with $c_k(s) \geq 2$ for all k and $|M(s)| \geq 2$. By the tie-breaking rule, the smallest element of $M(s)$ is the minimizer and only the players having chosen this minimizer get a nonzero expected payoff. But no player has an incentive to deviate

---

[5]A potential exception is Arad and Rubinstein (2012), who proposed the 11-20 Money Request Game as a simple game which elicits LK-reasoning and used somewhat similar arguments as we do here. (At the time we designed the MG their paper had not yet been circulated.) However, they do not try to predict choice behavior. Moreover, as they explain, the 11-20 game is not well suited to distinguish between LK- and PCH-models. According to our criteria it would also not qualify as "pure", since it has a unique maximin choice.

unilaterally, because switching to the minimizer causes this amount to lose its minimizer status and leads to a zero payoff. Hence s is a Nash equilibrium. The MG has a variety of such (and other) asymmetric equilibria. However, all these asymmetric equilibria, requiring explicit coordination in a symmetric one-shot setting, are implausible solutions. Classical game theoretic choice prediction would instead point to a symmetric Nash equilibrium of the MG. Such an equilibrium always exists, since the game is symmetric and finite.

The MG can not be solved analytically and the simplifying approach of modeling the MG as a Poisson game (Myerson, 1998, 2000) does not work.[6] However, for our purposes we do not need to explicitly calculate an equilibrium. Indeed, every symmetric equilibrium of the MG approaches the uniform distribution on the set N of available choices as the number I of players grows to infinity.[7] Since we have a large number of participants in our experiments, we can therefore safely approximate the symmetric equilibrium by the uniform distribution. As an example, if $N = \{100, 150, 200\}$, as in our basic experimental treatment, then the equilibrium distribution is $E \approx (0.320, 0.329, 0.351)$ for $I = 300$ players. Whenever the subject pool is sufficiently large, the distribution of choice frequencies as predicted by the Nash hypothesis will be close to uniform.

2.4 PCH predictions for the MG

What choice frequencies (p_1, p_2, p_3) does the PCH-model predict in the MG with $N = \{1, 2, 3\}$ in a large population? The answer depends on the exact

---

[6]The latter route proved successful for LUPI games (Östling et al., 2011) and lowest unique bid auctions (Pigolotti et al., 2012), but the uniqueness of winning choices in these games, which the MG lacks, is crucial for the Poisson games approach to generate a tractable solution.

[7]Let e be in the ω-limit of the sets of symmetric equilibria for $I \to \infty$. Assume e is not uniform. Then there are integers $\hat{k}$ and $\check{k}$ such that $e_{\hat{k}} > e_{\check{k}}$ are the maximal and minimal frequencies, respectively. But then by the law of large numbers we can make it arbitrarily more likely for $\check{k}$ to be the minimizer than $\hat{k}$ by choosing the number I of players large enough. This contradicts equality of expected payoffs in mixed equilibria.

values of τ and I, but for large I the prediction is that $p_1 < p_2 < p_3$ holds uniformly for $0 < \tau < \bar{\tau}$, where $\bar{\tau} \approx 1.8$. To see this, note that the nonstrategic $L0$ types choose each number with probability $1/3$ by assumption. Hence $L1$ types believe that choices are distributed uniformly and maximize their payoff by picking the highest number, 3. $L2$ types believe that all others are $L0$ and $L1$, hence in their opinion 3 will be chosen most often and 1 and 2 have equal chances of turning out as minimizers. $L2$ types therefore choose 2, if I is large. For the same reason $L3$ types opt for 1 as the minimizer. The choices of $L4$ and higher-level types are less straightforward, as they depend on the order of the frequencies of $L1$, $L2$, and $L3$. However, numerical computation shows that if $0 < \tau < \bar{\tau}$, then the lowest number, 1, has the lowest frequency according to the beliefs of $L3$ and all higher types, which therefore also pick 1, if I is not too small.

For $\tau > \bar{\tau}$ the PCH-prediction is near and for $\tau \to \infty$ converges to the equilibrium E which equals the uniform distribution for an infinite number of players. In the range $\tau \in [0, \bar{\tau}]$, the PCH-predictions for increasing τ describe a loop as depicted in Figure 1.[8] The loop starts at E, dives into the triangular section of the simplex where $p_1 < p_2 < p_3$, takes a turn at around $\tau = 0.8$ and heads back towards E until it intersects itself at $\tau = \bar{\tau}$. The pfPCH point prediction is $(p_1, p_2, p_3) \approx (0.266, 0.325, 0.409)$.

2.5 Level-k predictions for the MG

Level-k models have several parameters, viz. the frequencies of the various types of players. These frequencies have to be estimated from the data or transferred from estimations in similar games, so unlike Nash equilibrium or the pfPCH-model, an unconstrained level-k model does not give a point prediction for the MG. For predictions we therefore take the generous approach

---

[8]Technically, this loop is not smooth near E if I is finite, since for every given I there exists a threshold value of τ below which the $L2$ type picks 3 instead of 2. But for large I this threshold value and the resulting discontinuity in the loop are so small that the latter is not visible in Figure 1.

to allow for all point predictions possibly arising from the "typical" estimates outlined in section 1.2 above: The population predominantly consists of $l1$ and $l2$, complemented by smaller fractions of $l3$, $l4$, $d1$, and equilibrium types. We call the level-k model with these characteristics the *standard level-k model* or SLK-model.

In the MG, the SLK-model's types $l0$, $l1$, and $l2$ are behaviorally indistinguishable from the PCH-model's corresponding types $L0$, $L1$, and $L2$. These types choose uniformly, 3, and 2, respectively. Note, however, that contrary to the PCH-model higher types in the SLK-model, by best responding to level-$(k-1)$, never choose number 1 but switch between choosing 3 and 2 only. The dominance type $d1$ behaves like $l1$ and chooses 3, since there are no dominated strategies. Finally, the equilibrium type of the SLK-model behaves like the $l0$ type, since the symmetric equilibrium has an approximately uniform distribution of number choices. An important constraint in the SLK-model is that, as $l1$ and $l2$ are predominant, any other type's frequency is restricted to be at most 1/3. For the MG this means that at least 2/3 of the population chooses 3 or 2 (types $l1$ to $l4$ and $d1$) and at most 1/3 chooses uniformly (equilibrium type). This translates into the "prediction set" $\{(p_1, p_2, p_3) : p_1 \leq \min(p_2, p_3, 1/9)\}$ for the SLK-model. This prediction set is depicted by the shaded area in Figure 1. Note that at most about 11% of all choices fall on the low amount 1, since this amount is only picked by 1/3 of the equilibrium type, whose frequency is itself constrained to be less than 1/3.

3 Experiments

3.1 Basic setup and hypotheses

Since the predicted choice frequencies of the pfPCH-model are not too far from the Nash equilibrium E, rejecting the Nash hypothesis with an adequate statistical power requires a large number of experimental subjects. We there-
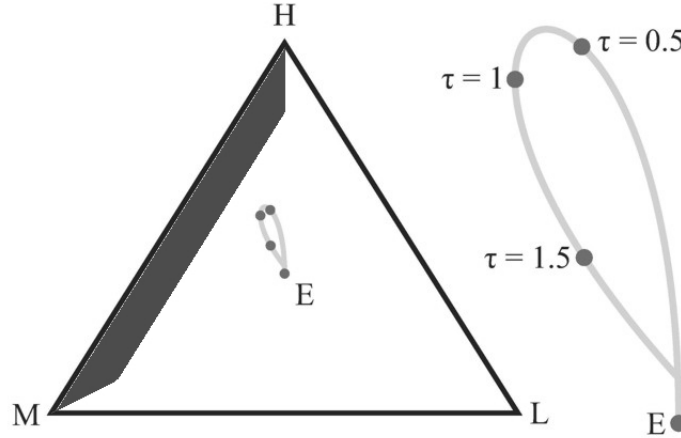
Figure 1: Nash (E), SLK (shaded area), and PCH predictions (loop) in the simplex

fore conducted two large-scale online experiments. A total of 1360 subjects were recruited from first-year undergraduates at the Vienna University of Economics and Business. These subjects were unlikely to have been exposed to game theory, as this is only taught later on.

To provide appropriate, yet feasible incentives, we decided to use substantial amounts of money as prizes. The basic setup of the game required participants in the experiment to choose between the three amounts: €100, €150, and €200. All experiments presented in the following use variations of this basic game.

Based on the Nash equilibrium of this game for a sufficiently large subject pool and on the prediction set of the SLK-model, we can state the following two null-hypotheses:

(H0a) Choice frequencies are uniformly distributed ($p_1 = p_2 = p_3$).

(H0b) Choice frequencies belong to the SLK-model's prediction set $\{(p_1, p_2, p_3) : p_1 \leq \min(p_2, p_3, 1/9)\}$.

Opposed to these two benchmarks, the predictions derived from the PCH-

model gives rise to three increasingly sharp hypotheses:

(H1) The choice frequencies p are ordered $p_1 < p_2 < p_3$.

(H2) The choice frequencies p can be derived from a PCH-model for some Poisson parameter τ.

(H3) The choice frequencies p can be derived from the pfPCH model, i.e. from the PCH-model with Poisson parameter $\tau = 1.5$.

For hypotheses testing and reporting of significant results we apply a 1% level of significance. Summarizing the results, we found a statistically significant deviation both from Nash equilibrium and from SLK predictions, while we were not able to reject any of the three hypotheses (H1)-(H3). The best-fitting Poisson parameter value was $\tau = 1.37$.

3.2 Experiment 1

The first experiment provides an initial analysis of the one-shot minimizer game. It gathers first evidence on the choice distribution in this game, tests the null-hypothesis that the distribution is consistent with a Nash equilibrium and, additionally, examines the predictions of the PCH-model and the SLK predictions.

When implementing a game for the first time, appropriate incentives are a main concern. Experimental studies frequently investigate whether behavior is sensitive to changes in incentives. This question becomes even more central in the domain of cognitive models, where incentives might crucially affect subjects' cognitive effort. To address this issue appropriately, our first experiment varies the stakes of the game in a between-subjects design. We denote the basic setup of Experiment 1 described above as the *low-stakes treatment* of Experiment 1. In addition, we introduced a *high-stakes treatment* where the amounts to choose from were quadrupled to €400, €600 and €800. If choice distributions are equal across both treatments, the basic findings of our experiment are not the result of a specific level of incentives.

3.2.1 Procedure

To conduct the online experiment we chose to use the free software LimeSurvey. A cohort of 3824 first-year undergraduate students from the Vienna University of Economics and Business were invited by e-mail. The e-mail asked students to participate in an online experiment that could earn them a substantial amount of money while requiring only a few minutes of their time. To reduce the transaction costs of participating as much as possible, the invitation e-mail contained a link to the university web page hosting the experiment. Each invited student was randomly assigned to one of the two treatments and received a unique seven-digit identification number to ensure that he or she could participate only once. Additionally, the e-mail invitation told subjects how many fellow students had been invited to the experiment as well, which helped them to assess the strategic situation and create homogeneous expectations about the possible pool of participants. 1905 students were invited to the low stakes treatment, of which 312 (154 females and 158 males) actually took part. For the high stakes treatment, 1919 students were invited of which 305 (164 females and 141 males) actually participated.[9] The experiment was open for participation for four days.

On the experiment website subjects were instructed that they will have to choose one of three options that correspond to three different payoff levels.[10] They were told that the 'winning number' is the one that is chosen least often among all participants in the treatment. After their choice among the three amounts of money, subjects had to state their gender and age before submitting their choice and completing the experiment. The winning number was announced in an e-mail after the experiment was closed and all participants were invited for the public draw of the winner.

---

[9]After clicking the link to the web page in the invitation e-mail virtually no-one dropped out of the experiment.

[10]For instructions and screen shots, see the Appendix.

3.2.2 Results

As the predictions for the distribution of choices across the three options are independent of the stake size, we denote the three amounts – irrespective of the treatment – Small (corresponding to €100 or €400), Medium (€150 or €600) and Large (€200 or €800). Table 1 presents absolute (n) and relative (f) choice frequencies for the three amounts in each treatment.

Table 1: Experiment 1: Summary of results

Treatment	Low stakes		High stakes		Pooled	
Payoff	n	f	n	f	n	f
Small	64	0.21	69	0.23	133	0.22
Medium	108	0.35	103	0.34	211	0.34
Large	140	0.45	133	0.44	273	0.44
Total	312	1	305	1	617	1

Table 1 suggests that the distribution of choices across the three possible amounts is nearly identical across treatments. Indeed, a χ^2-test cannot reject the null hypothesis of equality of distributions across the low and high stakes treatment ($p = .8160$); columns 6 and 7 therefore show absolute and relative frequencies of choices when pooling the two treatments.

The choice frequencies of the three payoffs in both treatments are, however, significantly different from a uniform distribution (χ^2-Tests, both $p < .0001$) and thus from the Nash equilibrium. The minimizer is always the small amount, while the most frequently chosen number is the large amount. The difference in choice frequencies between the three amounts, and therefore the ranking of payoffs with respect to the choice frequencies ($p_1 < p_2 < p_3$) is highly significant in all pairwise comparisons (one sample test of proportions, $p < .0001$ for $p_1 = p_2$ and $p = .0004$ for $p_2 = p_3$ for both, low and high stakes). These decreasing choice frequencies from the large to the small

amount support an explanation in line with the PCH-model.[11] Since the low amount is chosen much more frequently than allowed by the SLK-model, the choice frequencies are also significantly different from those predicted by the SLK-model ($p < .0001$ for both treatments).

Evidence of this first experimental test of the minimizer game seems to be well in line with the PCH-model and rejects the classical game theoretic predictions of the Nash equilibrium as well as the predictions arising from the SLK-model.

3.3 Experiment 2

Experiment 2 basically replicates the low stakes treatment of Experiment 1, but here the three amounts to choose from were permuted to avoid possible framing by the increasing order in which the amounts were presented in Experiment 1.[12] As in the previous experiment, a (fresh) cohort of first-year undergraduate students were approached by e-mail invitations. 3680 e-mails were sent, resulting in a total number of 743 participants (404 females and 339 males).[13] All other procedural details were the same as in the previous experiment.

3.3.1 Results

Table 2 provides an overview of the choice distribution across the three payoffs. The distribution is significantly different from a uniform distribution

---

[11]The background characteristics of participants allows to look into potential gender differences with respect to thinking steps. Tests on the equality of distributions across female and male participants, as well as across stakes for both females and males, do not reveal any significant differences with respect to gender.

[12]Actually, Experiment 2 consisted of six rounds of the minimizer game. The intention of this was to enable the analysis of potential learning effects, the results of which will be published elsewhere. Here we focus on one-shot games and therefore consider exclusively data from the first round of Experiment 2.

[13]The slightly higher response rate to the invitations might be due to increased advertising efforts for experiments at the university.

(χ^2-Test, $p < .0001$) as well as from all SLK predictions ($p < .0001$). The ranking of choice frequencies ($p_1 < p_2 < p_3$, significant at $p < .0001$ for $p_1 = p_2$ and $p = 0.0054$ for $p_2 = p_3$) replicates the findings of the previous experiment and encourages an interpretation of the data in line with the PCH-model.

Table 2: Experiment 2: Summary of results

	Low stakes	
Payoff	n	f
Small	183	0.25
Medium	260	0.35
Large	300	0.40
# of subjects	743	1

3.4 Aggregate data

Data aggregation is facilitated by the fact that the basic game structure remains constant across both experiments. Indeed, the choice distributions in Experiment 1 and in Experiment 2 are not significantly different (χ^2-Test, $p = .2684$), and both exhibit the expected ranking of $p_1 < p_2 < p_3$. This result suggests that the underlying Poisson distribution of thinking steps and its parameter τ are indeed fairly constant across initial plays of the minimizer game. We thus pooled the data of Experiments 1 and 2 and present aggregate data the last two columns of Table 3.

Table 4 presents the pooled choice frequencies vis-a-vis the Nash prediction, the closest of all SLK predictions[14], and the parameter-free PCH prediction. Like for the single experiments, the overall choice distribution

---

[14]This is the maximum-likelihood point-prediction within the prediction set of the SLK-model.

Table 3: Aggregate choice frequencies

| | Experiment 1 | | Experiment 2 | | Pooled | |
Payoff	n	f	n	f	n	f
Small	133	0.22	183	0.25	316	0.23
Medium	211	0.34	260	0.35	471	0.35
Large	273	0.44	300	0.40	573	0.42
Total	617	1	743	1	1360	1

is significantly different from uniform ($p < .0001$), rejecting the Nash prediction (hypothesis H0a), and the closest from the set of SLK predictions (hypothesis H0b, $p < .0001$). Moreover, the observed ranking $p_1 < p_2 < p_3$ is highly significant in all pairwise comparisons ($p < .0001$ for each one). Hence we cannot reject hypothesis H1.

Table 4: Model predictions vs. choice frequencies

	Nash	SLK	PCH	actual choice
Small	0.33	0.11	0.27	0.23
Medium	0.33	0.40	0.33	0.35
Large	0.33	0.49	0.41	0.42

Figure 2(a) shows this aggregate choice distribution in relation to theoretically expected ones under the PCH-model. The grey shaded area is a 99% pointwise confidence band along the loop and represents choice distributions in the simplex that are consistent with a PCH prediction but neither the Nash nor the SLK prediction. The aggregate distribution, marked by "pooled", lies within this area. This illustrates that we cannot reject hypothesis H2 and our aggregate findings are well in line with the PCH prediction.

τ-values ranging from $\tau = 1.20$ to $\tau = 1.52$ predict choice frequencies which are statistically indistinguishable on a 1% margin from the sample

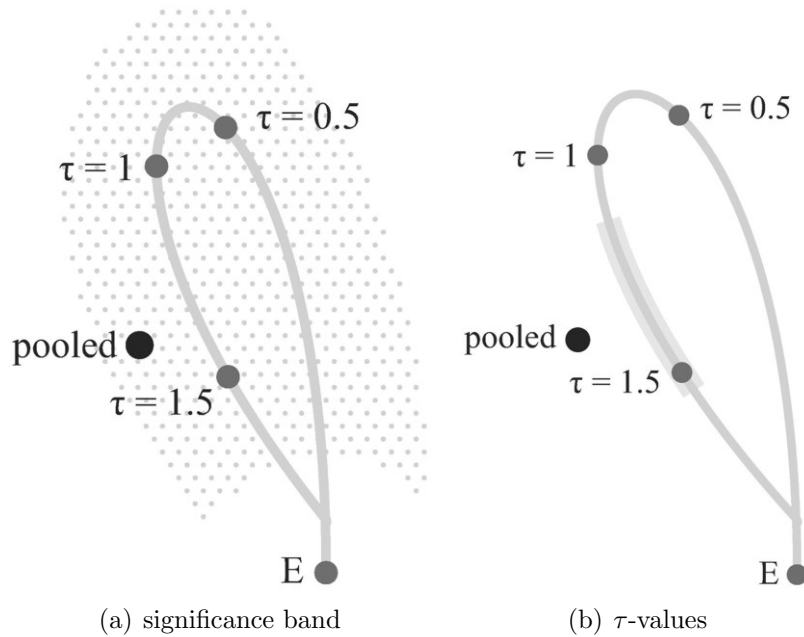(a) significance band   (b) τ-values

Figure 2: (a) Aggregate choice distribution and significance band along the loop. (b) PCH-predictions that are not significantly different from the pooled data.

distribution based on our pooled data. Figure 2(b) illustrates this interval of τ-values as the corresponding loop segment of choice frequencies. According to a maximum likelihood estimator, the best fitting value for τ under the assumption that the PCH-model applies is 1.37. Remarkably enough, we cannot even reject the point prediction from the pfPCH-model, i.e. hypothesis H3 ($p = 0.0187$). On the contrary, our results corroborate the suggestion of Camerer et al. (2004) that the PCH-model with a parameter of 1.5 could be used as a parameter-free model of choice prediction.

4 Conclusion

We discussed the question why the PCH-model seems to predict well in some games and fails in others. Based on conjectures about possible biases due

to social preferences, complexity-induced infeasibility of maximizing behavior, and maximin-principle interference, we constructed a multi-player game, the *minimizer game*, that avoids these obstacles. We then formulated three increasingly sharp hypotheses from the PCH approach, where the last one corresponds to a context- and parameter-free prediction. We tested these hypotheses in two large-scale Internet experiments. Despite considerable statistical power, none of the three hypotheses could be rejected. We thus confirmed the PCH-prediction in its strongest, parameter-free form. Stake size did not appear to influence the distribution of thinking steps. The competing predictions derived from Nash equilibrium and the SLK-model are clearly rejected by the data from our experiments.

We do not mean to imply that $\tau = 1.5$ is the best overall choice for the PCH-model; we just used this suggestion to avoid any appearance of *post hoc* model fitting. But we believe that our results show that when a game is "pure" and "simple", i.e. stripped of all complications introduced by social preferences, algebraic complexities and risk issues, then the thinking steps approach to predicting behavior is useful. Furthermore, the pfPCH-model seems to predict remarkably well, while no specification of the standard level-k model is able to do so, since it ignores subjects' taking into account that the behavior of others might be heterogenous. We therefore propose the pfPCH model as a useful context- and parameter-free alternative to Nash equilibrium in predicting initial choices in simple games.

References

[1] Arad, A., Rubinstein, A. (2012). The 11-20 money request game: A level-k reasoning study." American Economic Review 102 (7), 3561-3573.

[2] Arthur, W. B. (1994). Inductive reasoning and bounded rationality. American Economic Review 84 (2), 406-411.

[3] Aumann, R.J., Brandenburger, A. (1995). Epistemic conditions for Nash equilibrium. Econometrica 64, 1161-1180.

[4] Binmore, K. (2007). Does Game Theory Work? The Bargaining Challenge. MIT Press, Cambridge, MA.

[5] Burchardi, K.B., Penczynski, S.P. (2014). Out of your mind: Eliciting individual reasoning in one shot games. Games and Economic Behavior 84, 39-57.

[6] Camerer, C.F. (2003). Behavioral Game Theory: Experiments in Strategic Interaction. Princeton University Press, Princeton, NJ.

[7] Camerer, C.F., Ho, T.-H., Chong, J.K. (2002). A cognitive hierarchy theory of one-shot games: Some preliminary results. Working paper.

[8] Camerer, C.F., Ho, T.-H., Chong, J.K. (2004). A cognitive hierarchy model of games. Quarterly Journal of Economics 119 (3), 861-898.

[9] Carvalho, D., Santos-Pinto, L. (2014). A cognitive hierarchy model of behavior in the action commitment game. International Journal of Game Theory 43 (3), 551-577.

[10] Costa-Gomes, M.A., Crawford, V.P. (2006). Cognition and behavior in two-person guessing games: An experimental study. American Economic Review 96, 1737-1768.

[11] Costa-Gomes, M.A., Crawford, V.P., Broseta, B. (2001). Cognition and behavior in normal-form games: An experimental study. Econometrica 69, 1193-1235.

[12] Costa-Gomes, M.A., Crawford, V.P., Iriberri, N. (2009). Comparing models of strategic thinking in Van Huyck, Battalio, and Beil's coordination games. Journal of the European Economic Association 7, 365-376.

[13] Crawford, V.P., Iriberri, N. (2007a). Level-k auctions: Can a non-equilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions? Econometrica 75, 1721-1770.

[14] Crawford, V.P., Iriberri, N. (2007b). Fatal attraction: Salience, naivete, and sophistication in experimental hide-and-seek games. American Economic Review 97, 1731-1750.

[15] Gneezy, U. (2005). Step-level reasoning and bidding in auctions. Management Science 51 (11), 1633-1642.

[16] Grosskopf, B., Nagel, R. (2008). The two-person beauty contest. Games and Economic Behavior 62 (1), 93-99.

[17] Ho, T.-H., Camerer, C.F., Weigelt, K. (1998). Iterated dominance and iterated best response in experimental p-beauty contests. American Economic Review 88, 947-969.

[18] Maynard Smith, J. (1982). Evolution and the Theory of Games. Cambridge University Press, Cambridge, UK.

[19] Myerson, R.B. (1998). Population uncertainty and Poisson games. International Journal of Game Theory 27(3), 375-392.

[20] Myerson, R.B. (2000). Large Poisson games. Journal of Economic Theory 94(1), 7-45.

[21] Nagel, R. (1995). Unraveling in guessing games: An experimental study. American Economic Review 85 (5), 1313-1326.

[22] Nash, J.F. (1950). Non-Cooperative Games. PhD thesis, Princeton University, Princeton, NJ.

[23] Nash, J.F. (1951). Non-cooperative Games. The Annals of Mathematics 54 (2), 286-295.

[24] Östling, R., Wang, J.T., Chou, E., Camerer, C.F. (2011). Testing game theory in the field: Swedish LUPI lottery games. American Economic Journal: Microeconomics 3 (3), 1-33.

[25] Pigolotti, S., Bernhardsson, S., Juul, J., Galster, G., Vivo P. (2012) Equilibrium strategy and population-size effects in lowest unique bid auctions. Physical Review Letters 108 (8), 088701.

[26] Stahl, D.O., Wilson P.W. (1994). Experimental evidence on players' models of other players. Journal of Economic Behavior and Organization 25 (3), 309-327.

[27] Stahl, D.O., Wilson P.W. (1995). On players' models of other players: Theory and experimental evidence. Games and Economic Behavior 10 (1), 218-254.

[28] Van Huyck, J.B., Battalio, R.C., Beil, R.O. (1991). Strategic uncertainty, equilibrium selection, and coordination failure in average opinion games. Quarterly Journal of Economics 106 (3), 885-910.

[29] Wright, J.R., Leyton-Brown, K. (2014a). Level-0 meta-models for predicting human behavior in games. In Proceedings of the fifteenth ACM conference on Economics and Computation. ACM, New York, 857-874.

[30] Wright, J.R., Leyton-Brown, K. (2014b). Predicting human behavior in unrepeated, simultaneous-move games. Preprint, arXiv:1306.0918v2.

A   Appendix: Instructions

These instructions have been translated from German. Original instructions are available from the authors upon request.

A.1   Experiment 1: Instructions and choice screen

Screen 1: Instructions

Thank you for participating in this online experiment!

Instructions

On the following page, you will find three amounts of money to choose from. Please select one of the three amounts!

The amount that is selected **least often by all participants** is the **winning amount**.

Of all participants who selected this winning amount, one will be randomly drawn as the winner. This participant will be paid out the winning amount.

Note:
To keep the chances of being drawn the winner, you have to select the amount that you believe will be selected least often in total.

Proceed >>

Screen 2: Choice in low (high) stakes treatment

Please choose one of the three amounts:

☐ 100 € (400 €)
☐ 150 € (600 €)
☐ 200 € (800 €)

Help: *Here are the rules again: The amount that is chosen **least often** by all participants is the winning amount. Of all participants who have selected this winning amount, one will be randomly drawn as the winner. She/he receives this winning amount as payoff.*

Proceed >>

A.2 Experiment 2: Instructions and choice screen

Screen 1: Instructions

Thank you for participating in this online experiment!

Instructions

1. This online experiment consists of 6 rounds. In each of the following 6 rounds, you must select **one of 3 amounts** of money.
2. When all participants have completed the experiment, one of the 6 rounds is randomly drawn. This round is called **decision round**.
3. The amount of money that is in the decision round **selected least often** by all participants is the **winning amount**.
4. Of all participants who **have chosen the winning amount** in the decision round, one participant is randomly drawn. She/he is notified by mail and receives the winning amount in cash.

   Note: Each round can be the decision round. To keep the chances of being drawn the winner, you have to select the amount that you believe will be selected least often in this round.

Proceed >>

Screen 2: Choice in round 1

Round 1:
Please choose one of the three amounts:

☐ 100 €   ☐ 200 €   ☐ 150 €

*Help: This is round 1. If this round is drawn as the decision round, the amount that is chosen **least often** by all participants in this round is the winning amount. Of all participants who have chosen the winning amount in this round, one will be randomly drawn as the winner. She/he receives the winning amount in cash.*

<div style="text-align:center">Proceed >></div>

A.3   Both experiments: final screens

Screen: Gender

Please state your gender:

☐ Female
☐ Male

<div style="text-align:center">Proceed >></div>

Screen: Submission

To finally submit your choice(s), please click on the submit button.

<div style="text-align:center">Submit >></div>

Screen: Confirmation

Thank you for participating!

To confirm that we received your decision, you will obtain an e-mail shortly. In case you are the lucky winner, you will be notified by e-mail as well.