



On the persistence of strategic sophistication

Sotiris Georganas^a, Paul J. Healy^{b,*}, Roberto A. Weber^c

^a Dept. of Economics, City University London, Northampton Square, London, EC1V 0HB, UK

^b Dept. of Economics, The Ohio State University, 1945 North High street, Columbus, OH 43210, USA

^c Dept. of Economics, University of Zurich, Blumlisalpstrasse 10, 8006 Zurich, Switzerland

Received 2 April 2014; final version received 10 July 2015; accepted 15 July 2015

Available online 20 July 2015

Abstract

We examine whether the “Level- k ” model of strategic behavior generates reliable cross-game predictions at the individual level. We find no correlation in subjects’ estimated levels of reasoning across two families of games. Furthermore, estimating a higher level for Ann than Bob in one family of games does not predict their ranking in the other. Direct tests of strategic reasoning generally do not predict estimated levels. Within families of games, we find that levels are fairly consistent within one family, but not the other. Our results suggest that the use of Level- k reasoning varies by game, making prediction difficult.

© 2015 Elsevier Inc. All rights reserved.

JEL classification: C72; C91; D03

Keywords: Level- k ; Cognitive hierarchy; Behavioral game theory

1. Introduction

Following a considerable literature demonstrating deviations from Nash equilibrium play (see, for example, [Camerer, 2003](#)), behavioral research has sought to model the processes determining individual play and aggregate behavior in experimental games. One widely-used approach for modeling behavioral deviations from Nash equilibrium in one-shot games involves

* Corresponding author.

E-mail addresses: sotiris.georganas.1@city.ac.uk (S. Georganas), healy.52@osu.edu (P.J. Healy), roberto.weber@econ.uzh.ch (R.A. Weber).

the use of heterogeneous types, based on varying levels of strategic sophistication (Nagel, 1993; Stahl and Wilson, 1994; Costa-Gomes et al., 2001; Camerer et al., 2004).¹ In this framework—often referred to as *Level- k* or *Cognitive Hierarchy*—players' strategic sophistication is represented by the number of iterations of best response they perform in selecting an action.

In the simplest version of these models, Level-0 types randomize uniformly over all actions and, for all $k > 0$, the Level- k type plays a best response to the actions of Level- $(k - 1)$. Thus, the model suggests that a subject's level is a measure of her strategic sophistication—or, more precisely, her belief about her opponents' strategic sophistication. The application of such models to data from one-shot play in experiments has yielded several instances in which the model accurately describes the aggregate distributions of action choices. We provide a review of this literature in the next section.

The value of the Level- k framework as a *post hoc* descriptive model of the aggregate distribution of actions in laboratory games has been widely documented. There is also evidence that the overall distribution of levels may possess some stability across games (e.g., Camerer et al., 2004), meaning that one might be able to predict the distribution of actions in a novel game based on the distributions in other games.

However, an open question remains regarding whether Level- k types correspond to some meaningful individual characteristic that one might label as “strategic sophistication.” That is, does a particular individual's estimated level correspond to a persistent trait that can be used to predict play across games? If levels are indicative of strategic sophistication, and if strategic sophistication is an invariant characteristic of a person, then there should exist reliable cross-game patterns in players' observed levels. Estimated levels in one game could then be used to predict players' behavior in novel games. Moreover, estimates of a player's level could be improved by using direct psychometric measures that correlate with strategic sophistication.

On the other hand, if players' levels appear to be randomly determined from game to game, then one of two negative conclusions must be reached: Either iterative best response is not an accurate description of players' reasoning, or the model is accurate but players' levels vary from game to game in a manner that is difficult to predict. In either case, knowledge of a player's level in one game provides neither information about their play in another game, nor a useful measure of that person's strategic sophistication in general.²

In this paper we test for persistence of individuals' strategic sophistication across games. We begin by identifying several plausible, testable restrictions on cross-game behavior in the Level- k framework. For example, the most stringent testable restriction is that players' levels are constant across all games. A weaker restriction requires only that players' relative levels be invariant, so that a ranking of players based on their levels remains constant across games, even if their absolute levels do not.

We then conduct a laboratory experiment in which subjects play several games drawn from two distinct families of games. The first family of games consists of four novel matrix games

¹ An alternative approach involves modeling deviations from Nash equilibrium as noise (or unobservable utility shocks) in players' best response. For an example, see the Quantal Response Equilibrium model proposed by McKelvey and Palfrey (1995). Rogers et al. (2009) bridge the Quantal Response approach with the Level- k approach studied here. Other directions in behavioral game theory include the study of dynamics following initial play (see Crawford, 1995; Erev and Roth, 1998; Camerer and Ho, 1998, for example).

² We do not suggest that levels must be constant across games for the model to have predictive power. Camerer et al. (2004) and Chong et al. (2005), for example, suggest that levels *will* change in certain situations. Predictive power simply requires that situational changes be predictable.

developed for this study, which we refer to as “undercutting games.” The second family is a set of two-person guessing games studied by [Costa-Gomes and Crawford \(2006\)](#) (henceforth CGC06).³

Within each family of games, we identify an individual’s level in a Level- k framework, following a standard approach for classifying individual behavior based on the observation of play in several games. We then test whether these observed levels satisfy any of the cross-game restrictions we have identified. To complement this analysis, we also attempt to identify individual levels separately for each game, and use these classifications to conduct cross-game comparisons within each family of games.

We also consider two additional ways in which strategic sophistication might be detectable. First, we elicit several direct measures of strategic intelligence using brief quizzes that have been found to identify strategic reasoning ability or general intelligence. We explore the relationship between such measures and subjects’ levels identified from their behavior. Second, we have subjects play each game against three different opponents: a subject randomly selected from the population in the session, the opponent who scored highest on the strategic intelligence measures discussed above, and the opponent who scored lowest. Thus, we are able to detect whether sophisticated types vary their behavior based on the expected sophistication of their opponent.

The degree of persistence in strategic sophistication that emerges from our data is mixed. The key results are summarized as follows:

- (1) The *aggregate* distribution of levels is similar to that found in previous studies for both families of games.
- (2) Individual levels show little persistence between the two families of games. Moreover, the relative ordering of players is also unstable between the two game families.
- (3) Looking within families of games, the aggregate distribution of levels is remarkably stable across undercutting games, but quite unstable across two-person guessing games. Individual levels and relative ordering are moderately persistent within the family of undercutting games, but have no persistence within the family of guessing games.
- (4) The quizzes generally fail to predict players’ levels in either family of games, though Level-1 play is correlated with a test for autism in undercutting games.
- (5) Some players adjust strategies against stronger opponents, but neither quiz scores nor levels predict which subjects make this adjustment.

Our interpretation of these results is that the congruence between Level- k models and subjects’ actual decision processes depends on the context. Players confronted with a novel game may have many alternative processes for determining what strategy to select, and different environments trigger the use of different decision processes.⁴ Level- k reasoning might be one process that is triggered in some contexts (undercutting games), but not in others (guessing games). Of course, if Level- k reasoning is employed in some games, it is critical for the theory’s predictive ability to be able to identify what factors trigger its use.

³ A two-person guessing game is different from the two-person beauty contest studied by [Grosskopf and Nagel \(2008\)](#); the latter has a (weakly) dominant strategy while the former does not.

⁴ This interpretation is similar to the idea of a “toolbox” of various decision making approaches or heuristics, which are employed varyingly depending on the context ([Gigerenzer, 2001](#)).

Additional insight into robustness comes from comparing our guessing game data with that of CGC06.⁵ In their data the Level- k model receives stronger support—especially when considering subjects’ “lookup” behavior—though we also find substantial cross-game instability in their data as well. We believe this difference in model accuracy stems from two differences in protocols. First, their instructions are far more detailed (spanning 31 computer screens) and include four practice rounds with feedback on aggregate choices. Ours consist of a 5-page handout with no practice rounds. Second, they require that subjects pass an understanding test in which they must calculate their best response to an opponent’s choice, and their opponent’s best response to their own choice. Subjects who fail this test are dismissed. We have no such understanding test, and include all subjects. It is possible that the more extensive instructions, the practice rounds with feedback, and the best-response understanding test all trigger the use of Level- k reasoning in a greater fraction of the subjects. Alternatively, it is possible that our design fails to eliminate confusion sufficiently, leading to noisier data. Regardless, the sensitivity of results to the protocol suggests that the Level- k model’s applicability may be limited in this regard.

2. Review of relevant literature

The notion of heterogeneous strategic sophistication operating through limited iterations of best response dates back at least to the “beauty contest” discussion of Keynes (1936). Motivated by this, Nagel (1993, 1995), Ho et al. (1998), and others study behavior in laboratory p -beauty contest games, in which all players submit a number in $[0, 100]$ and the closest guess to p times the average wins a prize. The observed distributions of guesses show clear spikes consistent with Level-1 and Level-2 play. This finding is robust to the structure of the game (Duffy and Nagel, 1997) and varied populations (Bosch-Domènech et al., 2002).⁶

Stahl and Wilson (1994) study Level- k behavior in ten 3×3 matrix games.⁷ They find that roughly 25 percent of players are Level-1, 50 percent are Level-2, and 25 percent are Nash equilibrium players. Level-0 play is virtually non-existent. Stahl and Wilson (1995) examine play in twelve normal-form games played without feedback, adding Worldly and Rational Expectations types. In both studies, many subjects fit strongly into one type, with posterior probabilities of their maximum likelihood type exceeding 0.90. Stahl and Wilson (1995) also provide a test of individual cross-game stability: They select a subset of nine games, estimate individuals’ types from these games, calculate the predicted choice probabilities for the remaining three games for each type, and then estimate the posterior probability that a subject has a particular type. They classify as “stable” those subjects for whom the posterior probability of having the same type is at least 15 percent. Using this relatively low threshold, they find that 35 of 48 subjects are stable.

⁵ Our data are comparable to the CGC06 “Baseline” and “Open-Box” treatments. See the online appendix for this and other comparisons with CGC06.

⁶ When the game is made into a global game in which players are also rewarded for guessing an unknown state (Morris and Shin, 2002), however, the Level- k model fits poorly when the state-guessing incentive is emphasized (Shapiro et al., 2009).

⁷ In their model Level-0 players are assumed to randomly choose strategies, Level-1 players best respond to Level-0, and Level-2 players best respond to a Level-1 strategy with noise added. This works similarly to best responding to a mixture of Level-0 and Level-1.

In contrast, we estimate a player's type independently in two sets of games, and directly compare whether the two estimated types are identical.⁸

Costa-Gomes et al. (2001) fit a Level- k model, with 9 possible types, to behavior in 18 matrix games. In their experiment, payoffs in the games are initially hidden to subjects, so that estimation of a player's level based on strategy choice can be augmented by analyzing which pieces of information subjects choose to view before making a decision. The model fits well, and they generally find higher levels in simpler games.

Camerer et al. (2004) introduce the *Cognitive Hierarchy* variation of the Level- k model, in which players best respond to the distribution of levels truncated below their own level. Thus, a Level- k player believes all other players are Level-0 through Level- $(k - 1)$ and his belief about the relative frequencies of those levels is accurate. Using a Poisson distribution of levels reduces the model to a single parameter τ (after defining the Level-0 distribution) that describes the mean level in the population. They estimate this distribution for a wide range of games. In p -beauty contests, for example, they estimate higher mean levels in more educated populations, in simpler games, and when subjects are asked their beliefs about opponents' play. They also show that the model suffers relatively little loss in likelihood scores when restricting τ to be constant across games, indicating a fair amount of cross-game stability in the aggregate distribution of levels; however, they do not explore individual-level cross game stability.

In Costa-Gomes and Crawford (2006) (CGC06), players participate in 16 two-person guessing games in which a player and her opponent are each assigned an interval $[a_i, b_i]$ and a 'target,' $p_i \in \{0.5, 0.7, 1.3, 1.5\}$. Players' payoffs decrease in the distance between their own guess and p_i times their opponent's guess. As in the earlier paper (Costa-Gomes et al., 2001), lookup behavior is used to strengthen type estimation. Again the results support the Level- k model: A reasonably large percentage of players play exactly the strategy predicted by one of the Level- k types. Six of the ten games we study in this paper are two-person guessing games; we compare our findings to CGC06 in our analysis. Chen et al. (2009) study similar two-person games on a two-dimensional grid. They use eye-tracking technology to augment the type estimation based on behavior alone. They find distributions of types that are somewhat more uniform than in past studies. When subjects' data are randomly re-sampled to generate new bootstrapped samples, however, only 8 of 17 subjects receive the same classification in at least 95% of the bootstrapped samples as they did in the original sample. This suggests that roughly half of the subjects are not strongly consistent with any one level across these games.

Arad and Rubinstein (2012a) introduce the 11–20 money request game, which is similar to our undercutting game in that it is a simple game designed to trigger Level- k behavior while allowing a clean separation of levels. Although they are not explicitly testing for cross-game stability, they do find that subjects behave differently across variations of the game that do not change the equilibrium or Level- k predictions. De Sousa et al. (2012) identify non-strategic players by observing play in 10 beauty contest games, and find that these subjects are more likely to play non-strategically in the 11–20 money request game as well. Arad and Rubinstein (2012b) develop a model of multi-dimensional iterative reasoning, focused on 'features' of strategies, and apply it to behavior in Colonel Blotto games. They find that subjects who apply more iterative reasoning in an 11–20 money request game also seem to exhibit more multi-dimensional iterative reasoning in the Colonel Blotto game.

⁸ Burchardi and Penczynski (2014) and Penczynski (2011) find that players' estimated levels are altered after communicating with others. Although this represents one notion of type instability, it is unlikely that it stems from true randomness in players' types.

Batzilis et al. (2013) fit a Level- k model to a very large number of rock-paper-scissors ('Roshambull') games among Facebook users. They find that aggregate play frequencies vary slightly from the equilibrium prediction, and that most players' strategies are not consistently aligned with any single level.

The Level- k model has also been applied to extensive-form games. Kawagoe and Takizawa (2012) study centipede games and compare 12 different specifications of the Level- k /Cognitive Hierarchy model against two specifications of the Agent Quantal Response Equilibrium (AQRE) model (McKelvey and Palfrey, 1998). They find that an AQRE specification fits best for their constant-sum centipede game, while a Cognitive Hierarchy model with a uniform Level-0 strategy fits best their increasing-sum centipede game. Ho and Su (2013) develop a dynamic version of the Level- k model for centipede games and show that it fits both first-round play and the pattern of earlier taking as the game is repeated. The model also matches patterns of behavior in a dynamic bargaining experiment.

Relatively few authors test whether estimated levels correlate with personal traits such as intelligence. Camerer et al. (2004) find higher average levels in subject pools with greater academic training, such as Caltech undergraduates and game theorists. Burnham et al. (2009) show that individuals' choices in a p -beauty contest game correlate with scores on a 20-minute test of cognitive ability. Gill and Prowse (2013) also find a correlation between cognitive ability and levels in a p -beauty contest, and show that higher-ability players are more likely to converge toward equilibrium over time and earn higher payments. Chong et al. (2005) find that cognitive effort matters along with intelligence. They let their subjects play 22 mixed-equilibrium matrix games in a fixed order and report a positive correlation between thinking time and levels. Furthermore, average levels are higher in games 12–22 than in games 1–11, indicating a learning-by-doing increase in sophistication over time.⁹ Rubinstein (2014) finds correlation in players' reaction times across games, suggesting that some systematically engage in more contemplation than others. But he reports that he could not find interesting cross-game correlations in strategy choices, and that the level of contemplativeness is not very predictive of strategy choices. Similarly, De Sousa et al. (2012) do not find a correlation between the Elo ranking (a measure of a chess player's quality) and the likelihood of playing strategically.

The Level- k model has also been applied successfully to a variety of other games, including "hide-and-seek" games (Crawford and Iriberry, 2007b), incomplete-information betting games (Brocas et al., 2014), betting games and matrix games (Rogers et al., 2009), coordinated attack games (Kneeland, 2012), sender-receiver games augmented with eye-tracking data (Wang et al., 2009), and cheap-talk games (Kawagoe and Takizawa, 2009). In the field, Level- k has been shown to fit behavior in Swedish lowest-unique-positive-integer lottery games (Ostling et al., 2011) and to explain the fact that movies that were not released to critics before their public opening earn higher revenues (Brown et al., 2012). A functional MRI study even suggests differences in brain activity between subjects who exhibit varying degrees of "strategic sophistication" (Bhatt and Camerer, 2005). Finally, a few recent papers apply the Level- k concept to study departures from Nash equilibrium play in auctions, finding that the Level- k approach often, though not always, yields a significantly better fit than the Nash equilibrium (Crawford and Iriberry, 2007a; Georganas, 2011). However, Ivanov et al. (2010) show that models with misguided beliefs (such

⁹ In a personal communication, Camerer reported that a regression of individuals' average second-half level on their first-half level yields an R^2 value of 0.37, indicating reasonable predictive power in these games despite the learning-by-doing effect. Our experiment reduces the incidence of learning-by-doing effects by allowing subjects to revise any of their past decisions after making choices in all ten games.

as Level- k) cannot explain the winner's curse in common value auctions, because subjects who play against their own past actions still exhibit substantial overbidding.

For a more comprehensive survey of studies on the Level- k model, see Crawford et al. (2013).

3. A formulation of level- k models

The usual applications of the Level- k model generally treat it as an *ex post* descriptive model. As such, prior analyses typically omit cross-game or cross-individual testable restrictions, or test only how the aggregate distribution of types varies across games or populations (e.g. Camerer et al., 2004). In this section we introduce a formal framework in which such testable restrictions can be defined clearly. Our experiment then examines several possible cross-game testable restrictions to see which have empirical merit.

Specifically, we build a simple type-space model for two-player games where an agent's type describes her *capacity* for iterated best-response reasoning and her realized *level* of iterated best-response reasoning. Under Harsanyi's (1967) interpretation, types would also describe beliefs about opponents' types, second-order beliefs about opponents' beliefs, and all higher-order beliefs. Following the Level- k literature, however, we make the simplifying assumption that a player's level is a sufficient statistic for her entire hierarchy of beliefs, and that all players believe all others to have strictly lower levels than themselves.¹⁰

In our experiment, subjects play several two-person games. Let $\gamma = (\{i, j\}, S, u)$ represent a typical two-person game with players i and j , strategy sets $S = S_i \times S_j$, and payoffs $u_i : S \rightarrow \mathbb{R}$ and $u_j : S \rightarrow \mathbb{R}$. The set of all such two-player games is Γ . When players use mixed strategies $\sigma_i \in \Delta(S_i)$ we abuse notation slightly and let $u_i(\sigma_i, \sigma_j)$ and $u_j(\sigma_i, \sigma_j)$ represent their expected payoffs. In some cases players receive signals about the type of their opponent; we denote the signal i receives by $\tau_i \in T$, and let $\tau^0 \in T$ represent the uninformative "null" signal.

Player i 's type is given by $\theta_i = (c_i, k_i)$ where $c_i : \Gamma \rightarrow \mathbb{N}_0 := \{0, 1, 2, \dots\}$ identifies i 's capacity for each game $\gamma \in \Gamma$, and $k_i : \Gamma \times T \rightarrow \mathbb{N}_0$ identifies i 's level for each game $\gamma \in \Gamma$ and signal $\tau_i \in T$. The capacity bounds the level, so $k_i(\gamma, \tau_i) \leq c_i(\gamma)$ for all i, γ , and τ_i .¹¹ Let Θ be the space of all possible types. Note that c_i does not vary in τ_i since the capacity represents a player's underlying ability to "solve" a particular game, regardless of the type of her opponent. The realized level k_i may vary in τ_i , however, because the realized level stems directly from i 's belief about her opponent's strategy.

Beliefs are fixed by the model. Each player i 's pre-defined first-order beliefs are given by a mapping $\nu : \mathbb{N}_0 \rightarrow \Delta(\mathbb{N}_0)$ such that $\nu(k_i)(\{0, 1, \dots, k_i - 1\}) = 1$ for all $k_i \in \mathbb{N}_0$.¹² For example,

¹⁰ For example, Costa-Gomes et al. (2001), Costa-Gomes and Crawford (2006), Crawford and Iriberry (2007a), and Crawford and Iriberry (2007b) assume that all players with a level of $k > 0$ believe all other players' level to be $k - 1$ with probability one. Camerer et al. (2004), on the other hand, assume that all players with a level of $k > 0$ believe the realized levels of opponents to follow a truncated Poisson distribution over $\{0, 1, \dots, k - 1\}$. Whatever the assumption on first-order beliefs, all higher-order beliefs are then assumed to be consistent with this assumption (i believes j believes his opponent's levels follow this distribution, *et cetera*). Strzalecki (2009) builds a similar—though more general—type-space model that encompasses all Level- k models. It does not explicitly allow for levels to vary by game or for agents to update their beliefs upon observing signals, though both features could easily be incorporated.

¹¹ Technically, the inclusion of capacities is extraneous. A player's type could simply be defined as $k_i : \Gamma \times T \rightarrow \mathbb{N}_0$ and then a capacity would then be derived by setting $c_i(\gamma) = \sup_T k_i(\gamma, \tau_i)$ for each γ . We include capacities in the model to emphasize that agents' upper bounds on k_i may vary in γ .

¹² The simple interpretation of this assumption is that each player believes they are more sophisticated than all of their opponents. An alternative interpretation is that players are aware that they may be less sophisticated than some of their

in [Camerer et al. \(2004\)](#), $\lambda > 0$ is a free parameter and $v(k)(l) = (\lambda^l/l!)/\sum_{\kappa=0}^{k-1}(\lambda^\kappa/\kappa!)$ if $l < k$ and $v(k)(l) = 0$ otherwise. The function v is common knowledge and therefore is not included in the description of θ_i . Thus, the k_i component of a player’s type completely identifies her beliefs since v is a function only of k_i ; this is a common implicit assumption in the literature.

Behavior in a Level- k model is defined inductively. The Level-0 strategy for each player i in γ is given exogenously as $\sigma_i^0 \in \Delta(S_i)$. If $k_i(\gamma, \tau_i) = 0$ then player i plays σ_i^0 . For each level $k > 0$ the Level- k strategy $\sigma_i^k \in \Delta(S_i)$ for player i with $k_i(\gamma, \tau_i) = k$ is a best response to beliefs $v(k)$, given that each level $\kappa < k$ of player j plays σ_j^κ .¹³ Formally, for each $k > 0$, the strategy σ_i^k is such that for all $s'_i \in S_i$,

$$\sum_{\kappa=0}^{k-1} u_i(\sigma_i^k, \sigma_j^\kappa) v(k)(\kappa) \geq \sum_{\kappa=0}^{k-1} u_i(s'_i, \sigma_j^\kappa) v(k)(\kappa).$$

Finally, we define a Nash type, denoted by $k = N$, whose beliefs are $v(N)(N) = 1$. The profile σ_i^N is then the best response to the other player’s Nash-type strategy σ_j^N .¹⁴

When σ_i^k is degenerate (the Level- k strategy is a unique pure strategy) we let s_i^k be the strategy such that $\sigma_i^k(s_i^k) = 1$.

To see how this construction operates, fix a game γ and signal τ_i . If player i ’s type in this situation is $(c_i, k_i) = (0, 0)$ then she plays σ_i^0 . If i ’s capacity is one then her type is either $(1, 0)$ or $(1, 1)$. In the former case she plays σ_i^0 ; in the latter case her beliefs are $v(1)$, which has $v(1)(0) = 1$, and so she plays σ_i^1 . If i ’s type is $(2, 2)$ then she has beliefs $v(2)$, which puts pre-defined probabilities on her opponent being Level-0 and Level-1. In this case she plays σ_i^2 . For any (c_i, k_i) player i ’s beliefs are $v(k_i)$ and her best response to those beliefs is $\sigma_i^{k_i}$. Note that beliefs depend only on k_i , so player types $(4, 2)$, $(3, 2)$, and $(2, 2)$ all have the same hierarchy of beliefs, for example.

Once σ_i^0 and v are defined, the only testable prediction of this model is that in each game and for each signal all players must select a strategy from the set $\{\sigma_i^0, \sigma_i^1, \sigma_i^2, \dots\} \cup \{\sigma_i^N\}$.¹⁵ In many applications, the researcher assumes that each level k plays σ_i^k with noise (usually with a logistic distribution) and then assigns each subject to the level that maximizes the likelihood of their data across all games played.

As specified, a player’s level $k_i(\gamma, \tau_i)$ can be any arbitrary function of γ and τ_i . If no structure is imposed on the k_i function then the model is incapable of cross-game or cross-signal predictions; knowing that player i plays Level-2 in one game doesn’t provide information about i ’s level in another game. Our goal is to consider a set of reasonable cross-game or cross-signal testable restrictions on k_i and explore which (if any) receive empirical support. Understanding which restrictions on k_i apply will then lead to an understanding of the out-of-sample predictions

opponents, but they have no model of how more sophisticated players choose strategies. More sophisticated players are then treated as though they are Level-0 players. This second interpretation does suggest that $v(k_i)(0)$ should be positive for all k_i , which is inconsistent with the commonly-used assumption that $v(k_i)(k_i - 1) = 1$ for all k_i .

¹³ If there are multiple pure-strategy best responses then σ_i^k can be any distribution over those best responses, and that distribution is assumed to be known by all higher levels.

¹⁴ As is standard, we assume $v(k)(N) = 0$ for all $k \neq N$. If multiple Nash equilibria exist then multiple Nash types could be defined, but all of our games have a unique Nash equilibrium.

¹⁵ If σ^0 is not restricted then there are no testable predictions; letting σ^0 equal the empirical distribution of strategies provides a perfect fit.

that can be made through this model. If no restrictions on k_i can be found then no out-of-sample predictions can be made for an individual.

Examples of possible restrictions on k_i that one can test using experiments are:

- (1) **Constant:** $k_i(\gamma, \tau_i) = k_i(\gamma', \tau'_i)$ for all $i, \gamma, \gamma', \tau_i$, and τ'_i .
- (2) **Constant Across Games:** $k_i(\gamma, \tau_i) = k_i(\gamma', \tau_i)$ for all i, γ, γ' , and τ_i .
- (3) **Constant Ordering:** If $k_i(\gamma, \tau) \geq k_j(\gamma, \tau)$ for some γ and τ then $k_i(\gamma', \tau') \geq k_j(\gamma', \tau')$ for all γ' and τ' .
- (4) **Responsiveness to Signals:** For every γ and i there is some τ and τ' such that $k_i(\gamma, \tau) > k_i(\gamma, \tau')$.
- (5) **Consistent Ordering of Games:** For any τ , if $k_i(\gamma, \tau) \geq k_i(\gamma', \tau)$ for some i, γ and γ' , then $k_j(\gamma, \tau) \geq k_j(\gamma', \tau)$ for all j .

The first restriction represents a very strict interpretation of the Level- k model in which each person's level never varies, regardless of the difficulty of the game or the information received. The second restriction weakens the first by allowing players' beliefs to respond to information, but to otherwise keep levels constant across games.

Instead of forcing absolute levels to be constant, the third restriction requires only that players' relative levels be fixed. Thus, if Anne plays a (weakly) higher level than Bob in one game when they have identical information, then Anne should play a (weakly) higher level than Bob in all games where they have identical information. Certainly this would be violated with differing degrees of game-specific experience; recall, however, that the Level- k model applies only to the first-time play of novel games.¹⁶

The fourth restriction requires that there exist a pair of signals in each game over which a player's level will differ. Thus, a minimal amount of responsiveness to information, for at least some players, is assumed.

The last restriction listed implies that the observed levels can be used to order the games in Γ . If, at some fixed signal, all players play a lower level in γ' than in γ then it can be inferred that γ' is a more difficult or complex game. This enables future out-of-sample predictions, since a player who subsequently plays a given level in γ can be expected to play a lower level in γ' .

It is certainly easy to imagine plausible functions k_i that violate each of these restrictions, or that violate any other restriction we may consider. But each restriction that is violated means the loss of a testable implication for the model. If the most empirically accurate version of the Level- k model requires k_i functions that satisfy no cross-game or cross-signal restrictions, then the model cannot be used to make out-of-sample predictions about individual behavior. Thus, the predictive power of the model hinges on the presence of some identifiable restrictions.

4. The games

We study two families of games: a novel family of games that are useful for identifying player types—which we call undercutting games (UG)—and the two-person guessing games (2PGG) studied by [Costa-Gomes and Crawford \(2006\)](#).

¹⁶ Cross-game learning may still generate violations of this restriction; a chess master may play to a higher level than a professional soccer player in checkers, but to a lower level in an asymmetric matching pennies game. For this reason the boundaries of applicability of the Level- k model are sometimes ambiguous.

	1	2	3	4	5	6	7
1	1 1	10 -10	0 0	0 0	0 0	0 0	-11 0
2	-10 10	0 0	10 -10	0 0	0 0	0 0	0 0
3	0 0	-10 10	0 0	10 -10	0 0	0 0	0 0
4	0 0	0 0	-10 10	0 0	10 -10	10 -10	10 -10
5	0 0	0 0	0 0	-10 10	0 0	0 0	0 0
6	0 0	0 0	0 0	-10 10	0 0	0 0	0 0
7	0 -11	0 0	0 0	-10 10	0 0	0 0	-11 -11

Fig. 1. Undercutting game 1 (UG1).

	1	2	3	4	5	6	7	8	9
1	1 1	10 -10	0 0	0 0	0 0	0 0	0 0	0 0	-11 0
2	-10 10	0 0	10 -10	0 0	0 0	0 0	0 0	0 0	0 0
3	0 0	-10 10	0 0	10 -10	0 0	0 0	0 0	0 0	0 0
4	0 0	0 0	-10 10	0 0	10 -10	10 -10	10 -10	10 -10	10 -10
5	0 0	0 0	0 0	-10 10	0 0	0 0	0 0	0 0	0 0
6	0 0	0 0	0 0	-10 10	0 0	0 0	0 0	0 0	0 0
7	0 0	0 0	0 0	-10 10	0 0	0 0	0 0	0 0	0 0
8	0 0	0 0	0 0	-10 10	0 0	0 0	0 0	0 0	0 0
9	0 -11	0 0	0 0	-10 10	0 0	0 0	0 0	0 0	-11 -11

Fig. 2. Undercutting game 2 (UG2).

4.1. Undercutting games

An undercutting game is a symmetric, two-player game parameterized by two positive integers m and n with $m < n$. Each player $i \in \{1, 2\}$ picks a positive integer, $s_i \in \{1, 2, \dots, m, \dots, n\}$. Player i wins \$10 from player j if either $s_i = m < s_j$ or $s_i + 1 = s_j \leq m$. Thus, if player i expects her opponent to choose $s_j > m$, then her best response is to choose $s_i = m$; otherwise her best response is to “undercut” her opponent by choosing $s_i = s_j - 1$. If no player undercuts the other then one of the following situations apply: If both choose $s_i = 1$ (the unique Nash equilibrium choice) then both earn a payoff of one. If both choose n then both lose \$11. If i chooses one and j chooses n then i loses \$11 and j earns nothing. In all other cases both players earn zero. The cases where a player loses \$11 are designed to rule out any mixed-strategy Nash equilibria.

The payoff matrices of the undercutting games used in this experiment are shown in Figs. 1–4. Consider UG1, shown in Fig. 1. A levels-of-reasoning model that assumes uniformly random play by Level-0 types will predict that Level-1 types play $s^1 = 4$ as it maximizes the sum of row payoffs, Level-2 types play $s^2 = 3$, Level-3 types play $s^3 = 2$, and all higher levels play the

	1	2	3	4	5	6	7	8	9
1	1	10	0	0	0	0	0	0	-11
2	1	-10	0	0	0	0	0	0	0
3	-10	0	10	0	0	0	0	0	0
4	10	0	-10	0	0	0	0	0	0
5	0	-10	0	10	0	0	0	0	0
6	0	10	0	-10	0	0	0	0	0
7	0	0	-10	0	10	0	0	0	0
8	0	0	10	0	-10	0	0	0	0
9	0	0	0	-10	10	-10	0	0	0
10	0	0	0	0	0	10	10	10	10
11	0	0	0	0	-10	10	-10	-10	-10
12	0	0	0	0	0	-10	0	0	0
13	0	0	0	0	0	10	0	0	0
14	0	0	0	0	0	-10	10	10	10
15	0	0	0	0	0	10	-10	-10	-10
16	0	0	0	0	0	0	10	0	0
17	0	0	0	0	0	0	0	0	0
18	0	0	0	0	0	-10	0	0	0
19	0	0	0	0	0	10	0	0	0
20	0	0	0	0	0	-10	0	0	-11
21	-11	0	0	0	0	0	0	0	-11

Fig. 3. Undercutting game 3 (UG3).

	1	2	3	4	5	6	7
1	1	10	0	0	0	0	-11
2	1	-10	0	0	0	0	0
3	-10	0	10	0	0	0	0
4	10	0	-10	0	0	0	0
5	0	-10	0	10	0	0	0
6	0	10	0	-10	0	0	0
7	0	0	-10	0	30	10	10
8	0	0	10	0	-30	-10	-10
9	0	0	0	-30	0	0	0
10	0	0	0	30	0	0	0
11	0	0	0	-10	0	0	0
12	0	0	0	10	0	0	0
13	0	0	0	-10	0	0	-11
14	-11	0	0	10	0	0	-11

Fig. 4. Undercutting game 4 (UG4).

equilibrium strategy of $s^N = 1$. This enables a unique identification of a player’s level (up to Level-4) from a single observation of their strategy.

The game in Fig. 4, UG4, departs from UG2 only in that three dominated actions have been ‘compressed’ into one (which is now itself also dominated by another dominated action). Since dominated actions are never predicted for types above Level-0, this modification should have little impact on the distribution of types.

This family of games was designed explicitly for testing the Level- k model. Its undercutting structure is intended to focus players’ attention on the strategies of their opponents, encouraging Level- k -type thinking. The strategy space is relatively small, unlike p -beauty contest games, but the only strategy that confounds multiple levels (other than the Level-0 type, which may randomize over many strategies) is the Nash equilibrium strategy since all levels greater than m are predicted to play this action. There are no other Nash equilibria in pure or mixed strategies. Moreover, variations in the assumed Level-0 strategy have no impact on the ordering of players’ inferred levels. For example, if i plays 3 and j plays 2 then we infer that $k_j = k_i + 1$, regardless of the Level-0 specification.

4.2. Two-person guessing games

Two-person guessing games are asymmetric, two-player games parameterized by a lower bound $a_i \geq 0$, upper bound $b_i > a_i$, and target $p_i > 0$ for each player. Strategies are given by $s_i \in [a_i, b_i]$ and player i is paid according to how far her choice is from p_i times s_j , denoted by $e_i = |s_i - p_i s_j|$.

Each player i 's payment is a quasiconcave function of e_i that is maximized at zero. Specifically, players receive $15 - (11/200)e_i$ dollars if $e_i \leq 200$, $5 - (1/200)e_i$ dollars if $e_i \in (200, 1000]$, and zero if $e_i \geq 1000$. The unique best response is to set $e_i = 0$ by choosing $s_i = p_i s_j$. If $p_i s_j$ lands outside of i 's strategy space then the nearest endpoint of the strategy space is the best response. In a levels-of-reasoning model, Level-0 may be assumed to randomize uniformly over $[a_i, b_i]$ or to play the midpoint of $[a_i, b_i]$ with certainty. In either case Level-1 types will play $s_i^1 = p_i(a_j + b_j)/2$; if this is not attainable then the Level-1 player will select the nearest endpoint of her interval. A Level-2 type will play $s_i^2 = p_i s_j^1$ (or the nearest endpoint), and so on. This iterative reasoning converges to a Nash equilibrium with one player playing on the boundary of her interval and the other best-responding to that boundary strategy (see [Costa-Gomes and Crawford, 2006](#)).

5. Experimental design

In total, 116 undergraduate students from Ohio State University participated as subjects in these experiments. After reading through the experiment instructions, each subject completed five tasks, intended to measure general cognitive ability and strategic reasoning:

- (1) an IQ test,
- (2) the Eye Gaze test for adult autism,
- (3) the Wechsler digit span working memory test,
- (4) the Cognitive Reflection Test (CRT), and
- (5) the one-player Takeover game.

Each of these quizzes represents a previously-used measure of general intelligence or strategic sophistication. The IQ test consists of ten questions taken from the Mensa society's "workout" exam.¹⁷ Similar tests of cognitive ability have been shown to correlate with higher levels of reasoning in p -beauty contest games ([Burnham et al., 2009](#)).

The Eye Gaze test ([Baron-Cohen et al., 2001](#)) asks subjects to identify the emotions being expressed by a pair of eyes in a photograph. See [Fig. 5](#) for sample problems. Poor performance on this task is diagnostic of high-functioning adult autism or Asperger's Syndrome ([Baron-Cohen et al., 2001](#)) and strong performance is correlated with the ability to determine whether or not price movements in a market are affected by a trader with inside information ([Bruguier et al., 2010](#)).

The Wechsler Digit Span memory test tests subjects' abilities to recall strings of digits of increasing length. It is one component of the Wechsler Adult Intelligence Scale ([Wechsler, 1939](#)) to assess overall intelligence. [Devetag and Warglien \(2003\)](#) had 67 subjects take this short-term memory test and then play three games against a computerized opponent that always selected

¹⁷ See <http://www.mensa.org/workout2.php>.



jealous panicked arrogant hateful



aghast fantasizing impatient alarmed

Fig. 5. Sample questions from the Eye Gaze test.

the equilibrium strategy. The three games all required iterated reasoning to solve the equilibrium best response. They found a positive and significant correlation between subjects' memory test score and the frequency with which they selected the best response.¹⁸

The CRT contains three questions for which the initial intuitive response is often wrong. Performance on the test is correlated with measured time preferences, risk taking in gains, risk aversion in losses, and other IQ measures (Frederick, 2005). This measure also correlates with a tendency to play default strategies in public goods games (Altmann and Falk, 2009).

Finally, the one-player Takeover game is a single-player adverse selection problem in which the subject is asked to make an offer to buy a company knowing that the seller will only sell if the company's value is less than the offer. Given the parameters of the problem, all positive offers are unprofitable in expectation, yet many subjects fall victim to the "winner's curse" by submitting positive offers (Samuelson and Bazerman, 1985), even after receiving feedback and gaining experience (Ball et al., 1991).

We normalized each of the quiz scores to a scale of ten possible points. For scoring purposes during the experiment, we combined the CRT and Takeover game into one four-question, ten-point quiz, with answers coded using a binary (correct or incorrect) classification. For the Takeover game, subjects received a positive score if and only if their bid was exactly zero—the unique profit-maximizing bid.¹⁹ The sum of the four quiz scores was calculated for each player. Players were given no feedback about any player's absolute or relative performance on

¹⁸ Camerer et al. (2004) use this observation as a plausible justification for their assumption that the relative frequencies of two consecutive levels k and $k - 1$ ($f(k)/f(k - 1)$) is declining in k , which then motivates their restriction to Poisson distributions of levels.

¹⁹ In the data analysis below we disaggregate the CRT and Takeover game quizzes and treat them separately. The rationale for combining them in the experiment was to prevent a single question (the Takeover quiz) from having an excessively disproportionate weight. Also, in our analysis we use a score for the Takeover game that is linearly decreasing

Table 1
The ten games used in the experiment.

Game ID	Game type	Player's limits & target	Opponent's limits & target
UG1	Undercutting game	See Fig. 1	
UG2	Undercutting game	See Fig. 2	
UG3	Undercutting game	See Fig. 3	
UG4	Undercutting game	See Fig. 4	
GG5	Guessing game	([215, 815], 1.4)	([0, 650], 0.9)
GG6	Guessing game	([100, 500], 0.7)	([300, 900], 1.3)
GG7	Guessing game	([100, 500], 0.5)	([100, 900], 1.3)
GG8	Guessing game	([0, 650], 0.9)	([215, 815], 1.4)
GG9	Guessing game	([300, 900], 1.3)	([100, 500], 0.7)
GG10	Guessing game	([100, 900], 1.3)	([100, 500], 0.5)

the quizzes until the end of the experiment, at which point they learned only their own total quiz score.

After completing the quizzes, the subjects played ten games against varying opponents. The first four games are undercutting games and the last six are guessing games.²⁰ The parameters of each game are given in Table 1 and Figs. 1–4. The final three guessing games are identical to the first three, with the players' roles reversed. As in Costa-Gomes and Crawford (2006), this allows players to play both roles and also allows subjects' decisions in GG5, for example, to be matched with another subject's player-1 decision in GG8 to determine payoffs.

In each game subjects were asked to choose a strategy against a random opponent, against the opponent (other than themselves) with the highest total score on all of the quizzes, and against the opponent (other than themselves) with the lowest score on the quizzes. All choices were made without feedback. After making these three choices in all ten games, players learned that they could "loop back" through the games to revise their choices if desired. This could be done up to four times, for a total of five iterations through the ten games, all without feedback.²¹

Once subjects finished all five iterations—or declined the opportunity to loop back—their play was recorded, four of their choices were randomly selected (two from the undercutting games and two from the guessing games), and they were matched with another player and paid for their decisions. Subjects earning less than \$6 (the standard show-up fee) were paid \$6 for their time. Subjects earned an average of \$24.85 overall.

6. Data analysis procedures

Each subject played ten games, each against three different opponents, for a total of thirty game-play observations per subject. We employed three signals ($T = \{\tau^{\text{LO}}, \tau^0, \tau^{\text{HI}}\}$) indicating, respectively, whether the opponent had the lowest quiz score, was randomly selected, or had the highest quiz score. Following CGC06 (and others), we focus on the case where $v(k)(k-1) = 1$ for all $k > 0$ and σ_i^0 is uniform over S_i . We chose games so that the estimated levels (or, at

in a subject's bid. Specifically, a subject who submitted a bid of b_i was scored as earning $10(1 - b_i / \max_j b_j)$ points in our analysis.

²⁰ As we discuss later, in Result 4, we conducted a second iteration of the experiment in which we reversed the order of the two families of games. The results are very similar, indicating that the order of games did not affect the results.

²¹ Most subjects do not use the "loop back" option. In the undercutting games 15.5% submit a final choice different from their initial choice. In the guessing games this drops to 7.2%.

least, players’ relative rankings of levels) are fairly robust to these assumptions. Furthermore, the guessing-game parameters were chosen from among the CGC06 parameters to maximize the distance between any two levels’ predicted strategy choices; this helps to minimize the error in subjects’ level estimates.

For each subject i , signal τ , and set of games $G \subseteq \Gamma$ we can estimate a level, $k_i(G, \tau)$, using a simple maximum-likelihood approach that follows closely CGC06. Specifically, for each player i and level k we define a likelihood function $L(s_{i\gamma\tau}|k, \lambda_i, \varepsilon_i)$ for each level k based on the assumption that player i plays the Level- k strategy $s_{i\gamma}^k$ with probability ε_i , and otherwise maximizes a random expected utility with beliefs $\nu(k)$, extreme-value-distributed noise, and sensitivity parameter λ_i . We allow both parameters (ε_i and λ_i) to differ across players.

Formally, let $\bar{s}_{i\gamma\tau}$ be the value of the observed $s_{i\gamma\tau}$ rounded to the nearest integer, and let $I_i(s_{i\gamma\tau}, k)$ be an indicator function that equals one if $\bar{s}_{i\gamma\tau} = s_{i\gamma}^k$, where $s_{i\gamma}^k$ denotes the Level- k strategy for player i in game γ .²² $I_i(s_{i\gamma\tau}, k)$ equals zero otherwise. Thus, $I_i(s_{i\gamma\tau}, k) = 1$ indicates that i played exactly the Level- k strategy, allowing for rounding. The likelihood function for $k \neq 0$ is then given by

$$L(s_{i\gamma\tau}|k, \lambda_i, \varepsilon_i) = \varepsilon_i I_i(s_{i\gamma\tau}, k) + (1 - \varepsilon_i)(1 - I_i(s_{i\gamma\tau}, k)) \times \left(\frac{\exp\left(\lambda_i \sum_{\kappa} u_i(s_{i\gamma\tau}, \sigma_j^{\kappa}) \nu(k)(\kappa)\right)}{\int_{S_i} \exp\left(\lambda_i \sum_{\kappa} u_i(z_i, \sigma_j^{\kappa}) \nu(k)(\kappa)\right) dz_i} \right).$$

For $k = 0$ we set $L(s_{i\gamma\tau}|0, \lambda_i, \varepsilon_i)$ equal to $\sigma_{i\gamma}^0$, which is assumed to be the uniform distribution over S_i .

For any set of games $G \subseteq \Gamma$, denote i ’s strategies given signal τ by $s_{iG\tau} = (s_{i\gamma,\tau})_{\gamma \in G}$. For each level $k \in \mathbb{N}_0 \cup \{N\}$, the maximum likelihood of observing $s_{iG\tau}$ is given by

$$L^*(s_{iG\tau}|k) = \max_{\lambda_i > 0, \varepsilon_i \in [0,1]} \prod_{\gamma \in G} L(s_{i\gamma\tau}|k, \lambda_i, \varepsilon_i).$$

In practice, we search over a non-uniform grid of 122 possible values for λ_i and a uniform grid of 19 possible values for ε_i for each player i . The maximum-likelihood level for player i is then given by

$$k_i(G, \tau) = \arg \max_{k \in \mathbb{N}_0 \cup \{N\}} L^*(s_{iG\tau}|k).$$

Our games and our model of noisy play are such that the maximum-likelihood level is generically unique. Given that levels greater than three are very rarely observed in past data, we only calculate likelihood values for $k \in \{0, 1, 2, 3, N\}$.

We consider two types of analyses. First, we estimate for each subject one level for all undercutting games ($G = \{1, \dots, 4\}$) and another level for all guessing games ($G = \{5, \dots, 10\}$). This enables us to compare stability of levels across families of games. This pooling of several games per estimate also matches the standard procedure for estimating levels in the literature.²³ Second,

²² All strategies are integers in the undercutting games, in which case $\bar{s}_{i\gamma\tau} = s_{i\gamma\tau}$.

²³ As a robustness check, we apply our procedure to CGC06’s data, pooling all games to generate a single estimated level per subject (as in their paper), and find exact subject-by-subject agreement between our estimated levels and theirs.

Table 2

Frequency of levels in each game, and when pooling each family of games.

Game	L0	L1	L2	L3	Nash
UG1	7.76%	32.76%	19.83%	10.34%	29.31%
UG2	7.76%	32.76%	22.41%	7.76%	29.31%
UG3	5.17%	27.59%	18.10%	5.17%	43.97%
UG4	6.03%	31.03%	29.31%	5.17%	28.45%
UGs pooled	4.31%	28.45%	26.72%	5.17%	35.34%
GG5	6.03%	70.69%	9.48%	12.07%	1.72%
GG6	0.86%	65.52%	17.24%	11.21%	5.17%
GG7	43.10%	37.07%	13.79%	1.72%	4.31%
GG8	6.90%	39.66%	24.14%	21.55%	7.76%
GG9	5.17%	42.24%	23.28%	4.31%	25.00%
GG10	9.48%	38.79%	24.14%	19.83%	7.76%
GGs pooled	1.72%	50.00%	10.34%	10.34%	27.59%

we estimate for each subject a level in *every* game ($G = \{\gamma\}$). This enables us to compare stability of levels within each family of games.²⁴ In the online appendix, we also explore intermediate cases where two or three games per estimate are used.

7. Results

Result 1: Aggregate distributions of levels

The distributions of levels, both for each game family and for each individual game, are shown in Table 2. The aggregate game family distributions represent fairly typical distributions of estimated levels: Level-0 is observed fairly infrequently, Level-1 is the modal type, and Level-2 and Level-3 are observed less frequently. The distribution for guessing games is similar to the distribution found by CGC06. We do find that Nash play in our undercutting games is noticeably higher than what is found in many other games.

Within the family of undercutting games, the distribution of types is generally stable across games. In all four games, there is a high proportion of L1, L2 and Nash behavior, and relatively little behavior corresponding to L0 and L3.²⁵

Within the guessing games, however, distributions of levels vary substantially from one game to the next. For example, the fraction of Level-0 play jumps from 0.86% in guessing game 6 (GG6) to 43.10% in GG7. The fraction of Level-1 play nearly doubles from GG7 to GG5. Nash

²⁴ In this case $k_i(G, \tau)$ represents an assignment rule rather than an econometric estimate since only one observation is used for each “estimate” and no standard errors can be calculated. For the case of $|G| = 1$, we alternatively estimated levels in the guessing games by eliminating ε and $I_i(s_{i\gamma\tau}, k)$ and setting $\lambda = 1.33$ (the average estimated value of λ in CGC06 using only subjects’ guesses). We then assigned a k to each observation using maximum likelihood as described above. Under this new procedure, 85.5% of observations receive the same level assignment as in our original procedure. Roughly half of the observations whose level changes became Level-0 observations, implying their likelihood value simply falls below the uniform distribution likelihood. None of the key results of the paper change under these alternative estimates.

²⁵ Level-0 is necessarily under-counted here, since a proportion of all observed actions should be coming from Level-0 players. Although this cannot be corrected at an individual level, the aggregate frequency can be adjusted. The result simply shifts mass uniformly from the higher levels down to L0.

Table 3

Markov transitions from the pooled undercutting games to the pooled guessing games.

From ↓ to →	L0	L1	L2	L3	Nash
L0	0.0%	60.0%	0.0%	20.0%	20.0%
L1	6.1%	42.4%	6.1%	9.1%	36.4%
L2	0.0%	51.6%	16.1%	9.7%	22.6%
L3	0.0%	33.3%	33.3%	0.0%	33.3%
Nash	0.0%	56.1%	7.3%	12.2%	24.4%
Overall	1.7%	50.0%	10.3%	10.3%	27.6%

play ranges from 1.72% in GG5 to 25% in GG9. This suggests that either the Level- k model lacks descriptive power in these games, or else players' levels shift substantially between games.

We also find that 14.22% of all observations in the guessing games correspond exactly to one of the four (non-zero) levels' predictions, after rounding. This is clearly greater than the 0.7% frequency which would occur if actions were random with a uniform distribution. The most frequently-observed exact hit is the Level-1 action, in which players best respond to the midpoint of their opponent's interval, which accounts for roughly one-half of all the exact hits.²⁶

In an online appendix, we compare a graphical illustration of the likelihood functions for each level with a histogram of actions. This analysis shows that there does not appear to be a substantial and regular concordance between actions and the predicted behavior of different types across games. That is, the spikes in the likelihood functions do not consistently coincide with spikes in the data for any type across the different games.²⁷

Result 2: Persistence of absolute levels

To examine the hypothesis that levels are constant across games ($k_i(\gamma, \tau^0) = k_i(\gamma', \tau^0)$ for all γ and γ'), we generate a Markov transition matrix of levels between the two families of games. Table 3 reports the frequency with which a subject moves from each level in the pooled undercutting games to each level in the pooled guessing games. From the table, it is apparent that most of the transitions are into Level-1 and Nash types in the guessing game, and that these transitions do not show great correlation with a subject's type in the undercutting games. The

²⁶ By contrast, 48.9% of the observations in CGC06's data exactly correspond to one of the four levels' predictions. Cross-game variation in the distribution of levels remains high, however. See the online appendix for details. Again, we conjecture that these differences are due to differences in experimental instructions and their use of a best-response understanding test.

²⁷ This exercise also reveals purely mechanical reasons why subjects are classified more frequently as the Level-1 and Nash types. First, Level-1 beliefs are disperse, which means its likelihood function is quite flat. At the same λ_i , all higher levels have 'spike-shaped' likelihoods that exceed the Level-1 likelihood only in a small neighborhood around the predicted action. With uniformly-distributed random data, for example, Level-1 would be estimated to be the modal type for this reason. Second, the Nash type's predicted play is often at a boundary, so that logistic-response trembles can only occur in one direction. This truncation doubles the likelihood function on the interior of the strategy space, giving it a relative advantage over types with an interior prediction. We do find that estimated values of λ_i differ significantly across levels, but this appears to happen because those with noisier decisions are more likely to be classified as Level-1 due to its flatter likelihood function. In fact, the same correlation between λ_i and k is found when estimated on randomly-generated strategy data. This suggests that these two parameters do not capture independent traits—a Level-2 subject who becomes noisier is likely to be re-classified as Level-1—and that the correlation between them should not be interpreted as an insightful result.

Table 4

Markov transition between single-game levels within the four undercutting games.

From ↓ to →	L0	L1	L2	L3	Nash
L0	43.0%	22.6%	7.5%	9.7%	17.2%
L1	4.9%	59.7%	14.6%	4.4%	16.4%
L2	2.2%	20.2%	57.1%	9.0%	11.5%
L3	9.1%	19.2%	28.3%	18.2%	25.3%
Nash	3.5%	15.6%	7.9%	5.5%	67.5%
Overall	6.7%	31.0%	22.4%	7.1%	32.8%

Table 5

Markov transition between single-game levels within the six guessing games.

From ↓ to →	L0	L1	L2	L3	Nash
L0	8.7%	48.2%	18.1%	12.3%	12.8%
L1	11.7%	53.1%	16.8%	11.2%	7.1%
L2	11.5%	44.2%	27.4%	10.0%	6.9%
L3	12.4%	46.6%	15.9%	13.2%	12.0%
Nash	17.7%	40.3%	15.0%	16.3%	10.7%
Overall	11.9%	49.0%	18.7%	11.8%	8.6%

distributions in separate rows of Table 3 are generally similar to the overall distribution in the final row, which would occur if types were independent across families of games.

As a measure of the stability of levels across games, consider the prediction accuracy of the Level- k model assuming k_i is constant. This is simply the probability that a player plays the same level in two different games. We refer to this probability as the *constant-level prediction accuracy*, or *CLPA*. Mathematically, the CLPA equals the main diagonal of the Markov matrix weighted by the overall probability of each level. If types are constant then the main diagonal entries are all one, as is the CLPA. If types are randomly drawn then each row of the Markov matrix equals the overall distribution, and so the CLPA is simply the sum of squared overall probabilities in any row. In Table 3 the overall frequencies of the levels would imply a 29.4% CLPA under the null hypothesis of independent, randomly-drawn levels. The actual CLPA is 27.3%, suggesting a slight *negative* correlation in types across games.

To test whether levels are uncorrelated, we generate 10,000 samples of 116 randomly-drawn levels, with each sample drawn independently using the overall distribution from Table 3. For each sample we calculate the CLPA, generating an approximate distribution of CLPA values under the null hypothesis. A comparison of the actual CLPA with this distribution fails to reject the null that levels are randomly drawn across game families (p -value 0.68).²⁸

Tables 4 and 5 also show these transition matrices for the single-game levels in the undercutting and guessing games, respectively. Clearly, players' levels are more stable in the undercutting games than in the guessing games. In the undercutting games, over half of all Level-1, Level-2

²⁸ The cross-game (or cross-family) correlations can be also be tested statistically for any pair of games by calculating the Cramér correlation coefficient for categorical data (see Siegel and Castellan, 1988, p. 225) and comparing it against the null hypothesis of independently-drawn levels, which would give an expected Cramér correlation of zero. When comparing between the two families of games using pooled-game estimates (Table 3), the null hypothesis of independently-drawn types again cannot be rejected, with a Cramér correlation of only 0.177 and a p -value of 0.562.

and Nash types keep the same type across games. In the undercutting games, the overall frequencies of the levels (given in the last row of Table 4) would imply a CLPA of 26.3% if types were randomly drawn. In fact we observe a CLPA of 57.6%, indicating substantially stronger predictive power than if types were purely random, though still far from perfectly accurate. In a Monte Carlo simulation of 10,000 samples of independently-drawn levels, none have a CLPA this large. Thus, we reject the null hypothesis of random levels with a p -value of less than 0.0001.

The results are quite different in the guessing games, where Level-1 acts as an absorbing state. Little difference is seen between the rows of Table 5, suggesting no correlation across games. The realized prediction accuracy (CLPA) is 34.7%. The expected CLPA under randomly-drawn levels is 31.1%. Our Monte Carlo simulation of randomly-generated levels does reject the null hypothesis with a p -value of 0.0030, though the absolute magnitude of the difference (34.7% versus 31.1%) implies little real gain in predictive accuracy over the assumption of random levels.²⁹

We conclude that estimated levels can reasonably be modeled as constant within certain families of similar games, but not within other families. This suggests that Level- k thinking may be applied robustly in some settings, but not in others. Little guidance is currently available as to which families of games will trigger Level- k reasoning and which will not. In short, using a player's level in one game to predict her action in another may be a futile exercise without further information about the factors that determine whether Level- k reasoning is triggered.

Result 3: Persistence of relative levels

To examine the frequency with which the ordinal ranking of players' levels changes between the two families of games, we consider each possible pair of two players and measure the frequency with which the strictly higher-level player in one game becomes the strictly lower-level player in another ($k_i(\gamma, \tau) > k_j(\gamma, \tau)$ but $k_i(\gamma', \tau) < k_j(\gamma', \tau)$). We refer to this as the "switch frequency." This is compared against the "non-switch frequency," or the frequency with which the same player has a strictly higher level in both games ($k_i(\gamma, \tau) > k_j(\gamma, \tau)$ and $k_i(\gamma', \tau) > k_j(\gamma', \tau)$). Pairs whose levels are the same in at least one game are excluded, so the switch and non-switch frequencies often do not sum to one. The "switch ratio" is the switch frequency divided by the non-switch frequency; this has an expected value of one under the null hypothesis of independently-drawn levels. Under the Level- k model with stable relative levels, the ratio will equal zero.³⁰

The switch frequency, non-switch frequency, and switch ratio when comparing the pooled undercutting games to the pooled guessing games are reported in Table 6. The table also reports these statistics for the four undercutting games, and the six guessing games. The last column shows the predicted values under the null hypothesis of independently-drawn levels.

For the comparison between game families, switching actually occurs more frequently than non-switching. In other words, if Anne exhibits a higher level than Bob in the undercutting games, then Bob is more likely to have a higher level in the guessing games. Our 10,000-sample Monte Carlo simulation actually rejects the null hypothesis in favor of *negatively* correlated levels, with a p -value of 0.0230. This is consistent with our earlier observation that absolute levels are negatively correlated across families of games.

²⁹ Our analysis of the CGC06 data (in the online appendix) reveals a CLPA of 41.9%, which is between that of our guessing games and our undercutting games.

³⁰ In practice, the switch ratio would not exactly equal zero since some Level-0 players would be incorrectly identified as higher-level players. Our simulations suggest the actual switch ratio would be around 0.09 using our overall level distributions.

Table 6

Observed frequency with which two players' levels strictly switch their ordering, compared to the expected frequency under independent, randomly-drawn levels.

	Data	Null hyp.
Pooled UGs vs. pooled GGs		
Switch frequency:	25.0%	24.9%
Non-switch frequency:	22.7%	24.9%
Switch ratio:	1.10	1.00
Undercutting games		
Switch frequency:	13.2%	27.1%
Non-switch frequency:	45.3%	27.1%
Switch ratio:	0.29	1.00
Guessing games		
Switch frequency:	19.9%	23.8%
Non-switch frequency:	22.2%	23.8%
Switch ratio:	0.89	1.00

Since absolute levels within the undercutting games are fairly stable, we expect similar persistence in subjects' relative levels. This is the case: Non-switching pairs are observed more than three times more frequently than switching pairs, giving a switch ratio of 0.29. None of the 10,000 simulated samples have a switch ratio this low, indicating a clear rejection of the null hypothesis at the 0.0001 level.

In the guessing games, however, switching occurs nearly as frequently as non-switching, with a switch ratio of 0.89. The Monte Carlo simulation yields a marginal p -value of exactly 0.05. Thus, while there is some stability of relative levels within the guessing games, it is much weaker than the stability we observe in undercutting games, both in magnitude and statistical significance.

Overall, we conclude that little to no extra predictive power is gained by considering relative levels instead of absolute levels. Assuming k_i is constant for each subject performs roughly as well as assuming the ordering of k_i across subjects is constant.

Result 4: Robustness to order of play

In our experiment, every subject played the games in the same order, with the undercutting games (UG1–UG4) first and the guessing games (GG5–GG10) second. We find conformance with the Level- k model in the first set of games, but not in the second. One explanation for this result is that subjects become fatigued or lose attention through the course of the experiment, leading to more random play in later games. To test this hypothesis, we ran two new sessions in which the order of the games was completely reversed. Twenty-eight Ohio State subjects participated in two sessions of fourteen subjects each, using slightly modified versions of our original software and instructions.³¹ Participants took the quizzes, played GG10 through GG5, and then played UG4 through UG1.

³¹ To determine how many subjects were needed, we ran a bootstrapped power calculation for the Monte Carlo test of switch ratios. Specifically, for various values of $\hat{n} < 116$, we created 10,000 simulated data sets by drawing \hat{n} subjects (with replacement) from our original data. Then we generated 10,000 samples of \hat{n} simulated subjects with randomly-chosen levels. The power of the Monte Carlo test at \hat{n} is the fraction of "real" data sets whose switch ratio is less than 95% of the "random" data sets. The usual minimum power requirement of 80% is achieved at $\hat{n} = 9$. To be conservative, we aimed to recruit at least 20 subjects, and actually had 28 participate. Our test power is over 99%.

Recall that in our original data, the CLPA within the family of undercutting games is 57.6%, which is significantly greater than the 26.3% expected under randomly-drawn levels. In our new experiment, where the undercutting games are played second, the CLPA rises to 62.5%, indicating a slightly stronger conformance with the Level- k model. Again we reject the null hypothesis ($p < 0.0001$). The switch ratio, however, rises from 0.29 in the original data to 0.37 in the new data, indicating slightly weaker conformance with the Level- k model. But again we reject the null hypothesis of random levels with $p < 0.0001$. So the overall conclusion remains the same: across-game behavior in the undercutting games is fairly consistent with the Level- k model, and certainly far from behavior expected under the hypothesis of randomly-drawn levels, even when these games are played second.

For the guessing games, the CLPA shifts from 34.7% in our original data to 38.3% when these games appear first. This suggests slightly stronger conformance with the Level- k model. Again we reject the null ($p = 0.0070$), but note that the predictive accuracy is only 6.4 percentage points higher than under randomly-generated levels. The switch ratio, however, rises from 0.89 in the original data to 0.93 in the reversed treatment, indicating less conformance with the Level- k model. Again we cannot reject the null hypothesis ($p = 0.281$).

In summary, we do not find that overall conformance with the Level- k model is significantly improved in the guessing games when they appear first, or worsened in the undercutting games when they appear second.³² Thus, we reject the hypothesis that order effects explain our results.

Result 5: Using quizzes to predict levels

We next consider whether the five quizzes we administered as potential independent measures of strategic sophistication—the IQ quiz, the Eye Gaze quiz, a memory quiz, the Cognitive Reflection Test (CRT), and a one-player Takeover Game—predict behavior in the games. Because levels in the guessing games are very unstable, we do not expect them to be predictable by quiz scores. But levels in the undercutting game are stable, so we are particularly interested in whether these levels can be predicted using the quizzes.

First we examine correlations between scores on the various quizzes. These are surprisingly weak. IQ, memory, and CRT scores all appear to be positively correlated, though their estimated Spearman rank correlation coefficients achieve only marginal significance. No other correlations are statistically significant. One might conjecture that our subjects did not exert sufficient effort on the quizzes, leading to noisier scores, but absolute performance seems in line with previous studies for all quizzes except the Takeover Game.³³ Thus, we do not believe the lack of correla-

³² Game-by-game level distributions look similar to the original data, though the reversed data exhibit slightly more Level-1 and Level-2 play and less Level-3 and Nash play. The distributions are fairly stable in the undercutting games and highly variable across guessing games. The between-family CLPA rises to 41.0% in the reversed treatment, but we still cannot reject the null of random levels ($p = 0.112$). The between-family switch ratio drops noticeably from 1.10 to 0.76, but again we cannot reject the hypothesis of random levels ($p = 0.308$).

³³ In the Eye Gaze test, [Baron-Cohen et al. \(2001\)](#) report that the average score in the general population is 81%. Our subjects scored 80% on average. In the Wechsler digit span quiz, the average number of digits correctly recalled before the first failure is 5.63. In clinical applications the test stops after *two* failures; the average score among normal adults is between 5 and 7 ([Cohen et al., 2007](#), p. 416), consistent with our results. In the CRT, the percentage of players scoring (0, 1, 2, 3) (respectively) is (36%, 28%, 22%, 14%), which is very close to the overall average of (33%, 28%, 23%, 17%) reported in [Frederick's \(2005, p. 29\)](#) meta-study. In the Takeover Game, our subjects performed worse than in past studies: the mean bid was 94.3 for our subjects, while most studies report mean bids around 50 (e.g. [Grosskopf et al., 2007](#)). We do not have comparable data for our Mensa IQ test scores.

Table 7
Regression of expected earnings on the five quiz scores.

	Const.	IQ	EyeGaze	Memory	CRT	Takeover
Coefficient	39.179	0.620	0.250	−0.275	0.679	−0.229
p-value	(<0.001)	(0.204)	(0.318)	(0.215)	(0.002)	(0.261)

tion is caused by unusually poor performance or lack of effort. Instead, we believe these quizzes measure relatively orthogonal traits.

Next we ask whether quiz scores predict overall earnings. To reduce randomness in the earnings measure, we calculate what each subject's expected earnings would be in each game if they played against the empirical distribution of actions of all other subjects. The correlation between subjects' total expected earnings and the sum of their five quiz scores is positive, but not statistically significant (Spearman correlation of 0.172 with p -value 0.064). Regressing total expected earnings on each quiz (Table 7) reveals that only the Cognitive Reflection Test (CRT) score is significantly correlated with expected earnings.

Intuitively, players using higher levels should be more sophisticated. But they do not earn more money. Indeed, Level-2 is the most profitable type, since most subjects are estimated to be Level-1 types. When looking at correlations between quiz scores and estimated levels, we therefore do not restrict ourselves to a linear relationship, as Level-2 types may actually score the highest on the quizzes.

We focus on predicting levels estimated from the pooled families of games. For each type of quiz, we first perform a Kruskal–Wallis test of the null hypothesis that all five levels' quiz scores are drawn from the same distribution. If this null hypothesis is rejected for some type of quiz, then that quiz is diagnostic of at least one of the five estimated levels. In that case, we perform a multinomial logistic regression of levels on that particular quiz score to see which levels have significantly different quiz scores. Since multinomial logistic regressions require an omitted level against which all others are compared, one single regression is not useful in analyzing all possible comparisons. We therefore report the coefficient estimates from all five possible regressions, where each regression omits a different level.³⁴

Fig. 6 shows a box plot of the distribution of each quiz score for each of the five estimated levels in the pooled undercutting games. The p -values of the Kruskal–Wallis tests for each quiz type appear in parentheses at the top of the graph. We find significant differences across levels only for the Eye Gaze quiz, where Levels 0 and 1 appear to perform worse. The multinomial logistic regression results (Table 8) confirm that Level-0 Eye Gaze scores are significantly lower than those of Levels 2, 3, and Nash, and that Level-1 scores are significantly lower than Level-2 or Nash scores.

The Eye Gaze correlation with Level-0 and Level-1 play has intuitive appeal: Poor performance on the Eye Gaze quiz is diagnostic of adult autism (Baron-Cohen et al., 1997). And autism is often characterized the absence of “theory of mind” (Baron-Cohen, 1990), or an inability to recognize that others behave in response to conscious thought. This suggests that some of

³⁴ These five regressions are not meant to be treated as independent tests; rather, reporting them all provides a better view of what is essentially one regression. Using multinomial regression does control for the multiple comparisons within the regression (*i.e.*, within each column).

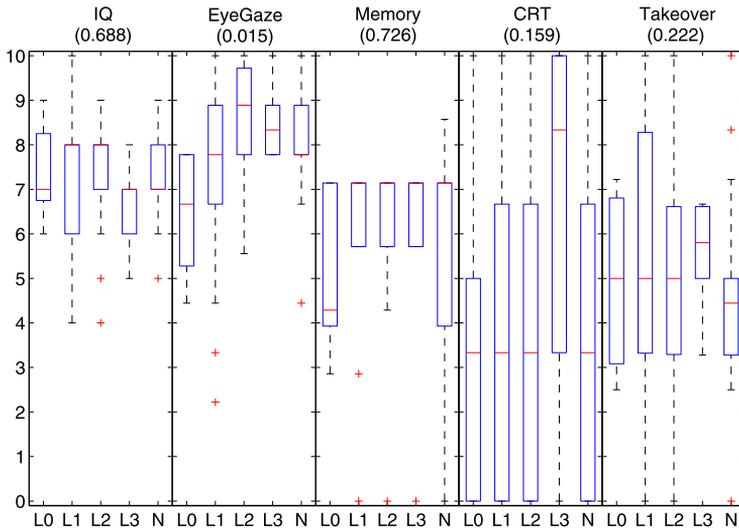


Fig. 6. Box plots of quiz scores for each estimated level in the pooled undercutting games. *p*-values in parentheses are for Kruskal–Wallis tests that all levels generate the same distribution of quiz scores.

Table 8

Multinomial logit regression coefficient estimates of Eye Gaze quiz scores on pooled undercutting game levels. Each column represents a regression with a different omitted category.

Eye Gaze score	vs. L-0	vs. L-1	vs. L-2	vs. L-3	vs. N
Level-0 (<i>n</i> = 5)	–	–0.311 (0.223)	–0.821 (0.006)	–0.865 (0.048)	–0.643 (0.020)
Level-1 (<i>n</i> = 33)	0.311 (0.223)	–	–0.510 (0.011)	–0.553 (0.143)	–0.332 (0.048)
Level-2 (<i>n</i> = 31)	0.821 (0.006)	0.510 (0.011)	–	–0.044 (0.909)	0.178 (0.361)
Level-3 (<i>n</i> = 6)	0.865 (0.048)	0.553 (0.143)	0.044 (0.909)	–	0.222 (0.554)
Nash (<i>n</i> = 41)	0.643 (0.020)	0.332 (0.048)	–0.178 (0.361)	–0.222 (0.554)	–

the Level-0 and Level-1 types are less able to consider others’ beliefs and strategies in games, leading them to play more low-level actions.³⁵

Fig. 7 reports the score distributions for levels estimated from the six pooled guessing games. The Kruskal–Wallis tests indicate that the CRT has some power in predicting subjects’ levels. Specifically, the multinomial logistic regressions (Table 9) indicate that Level-2 can be distinguished from the two higher levels, but not from the two lower levels.

³⁵ In *p*-beauty contest games, Coricelli and Nagel (2009) find that higher-level players exhibit greater neural activation in the medial prefrontal cortex (mPFC). Theory-of-mind experiments also find activation in this region (among others). These results are roughly consistent with our Eye Gaze finding, and also suggest more predictable heterogeneity in *p*-beauty contest games.

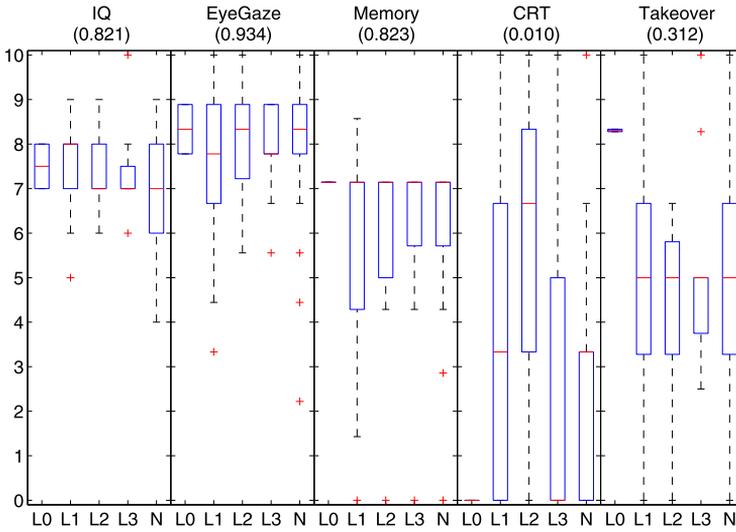


Fig. 7. Box plots of quiz scores for each estimated level in the pooled guessing games. *p*-values in parentheses are for Kruskal–Wallis tests that all levels generate the same distribution of quiz scores.

Table 9

Multinomial logit regression coefficient estimates of CRT quiz scores on pooled guessing game levels. Each column represents a regression with a different omitted category.

CRT score	vs. L-0	vs. L-1	vs. L-2	vs. L-3	vs. N
Level-0 (<i>n</i> = 2)	–	–2.730 (0.824)	–2.883 (0.815)	–2.574 (0.834)	–2.607 (0.832)
Level-1 (<i>n</i> = 58)	2.730 (0.824)	–	–0.153 (0.098)	0.156 (0.133)	0.123 (0.072)
Level-2 (<i>n</i> = 12)	2.883 (0.815)	0.153 (0.098)	–	0.310 (0.018)	0.277 (0.008)
Level-3 (<i>n</i> = 12)	2.574 (0.834)	–0.156 (0.133)	–0.310 (0.018)	–	–0.033 (0.766)
Nash (<i>n</i> = 32)	2.607 (0.832)	–0.123 (0.072)	–0.277 (0.008)	0.033 (0.766)	–

If levels in the guessing games are unstable, how can the CRT quiz be predictive of Level-2? Because Level-2 is also a proxy for higher earnings. We already know that the CRT quiz predicts earnings (Table 7), and those who earn more are more likely to be classified as Level-2, so the correlation between CRT quizzes and Level-2 appears spurious.

We perform similar analyses for game-by-game levels, and the results are consistent with the pooled-game results. In the undercutting games, players that play the Level-1 action in at least three of four games have lower Eye Gaze scores than Levels 0, 2, and Nash. They also have higher Takeover Game scores than Levels 0 and Nash. In the guessing games none of the quizzes are diagnostic of levels; the Kruskal–Wallis *p*-value for the CRT is 0.082 (with subjects estimated to be Level-2 in a majority of games scoring the highest), and is greater than 0.15 for all other quizzes.

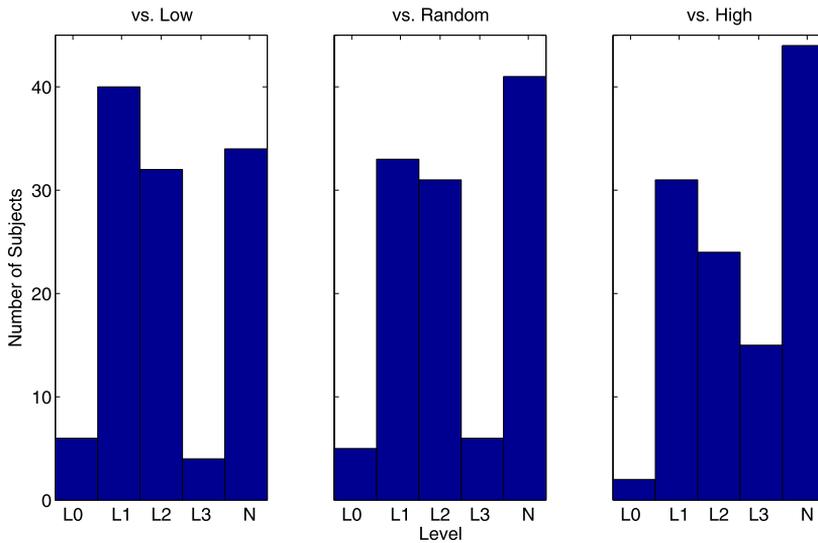


Fig. 8. Level distributions by opponent in the pooled undercutting games.

Result 6: Responsiveness to signals about opponents

In each game each subject is asked to choose a strategy against a randomly-selected opponent, against the opponent with the highest total quiz score, and against the opponent with the lowest total quiz score. Although quiz scores are not strongly related to levels of play—and the relationship certainly is not linear—they are correlated with total earnings, so we hypothesize that subjects might treat quiz scores as proxies for strategic sophistication.³⁶ Thus, how subjects respond to their opponents' characteristics may provide another testable prediction for the Level- k model.

Fig. 8 shows the histogram of estimated levels in the pooled undercutting games for each of the three types of opponent. Subjects appear to increase their level of reasoning against stronger opponents. In particular, Level-1 and Level-2 types become less frequent—and Nash types more frequent—when playing against opponents with higher quiz scores. χ^2 tests confirm that the distribution of levels is significantly different between the low-score and high-score opponent (p -value of 0.018), though not significantly different between the low-score and random opponents (p -value of 0.767) or between the random and high-score opponents (p -value of 0.185).

While the above differences in behavior by opponent are interesting, we are concerned with whether any information can predict this adjustment. That is, can we predict which subjects have a high enough “capacity” to be able to adjust their behavior in response to information about opponents? We therefore ask whether quiz scores predict the magnitude of adjustment. Using the pooled undercutting games, we measure for each subject the difference between their estimated level against a high-scoring opponent and their estimated level against a low-scoring opponent. This difference is then regressed on the five quiz scores. No regression coefficients are found to be significantly different from zero. Thus, quizzes fail to measure the propensity to adjust play against stronger opponents.

³⁶ Many subjects' responses to a debriefing questionnaire confirm this hypothesis.

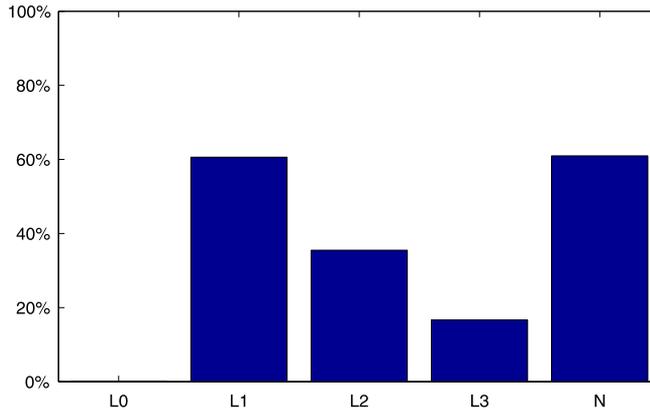


Fig. 9. For each k , the percentage of Level- k subjects in pooled undercutting games who do not change levels in response to opponent types.

Looking at which subjects do *not* shift strategies yields more informative results. For each k , we calculate the fraction of Level- k players in the pooled undercutting games against random opponents whose levels do not shift in response to high- or low-scoring opponents. We refer to these as *stable* players. The percentage of stable players for each k are shown in Fig. 9. If players' levels are constrained by their capacities, we should expect that low-level players are more likely to have low capacities, and therefore are more likely to appear as stable players. The data is consistent with this hypothesis for Level-1 through Level-3. Nash types, however, are the most stable. For the capacity-constrained Level- k model to hold, it must be that these players all have high enough capacities so that their chosen level is always greater than four. Such high levels are rarely observed in the literature, suggesting that these players are more likely “stubborn Nash” types who play Nash equilibrium strategies regardless of the opponent. Thus, there may exist heterogeneity amongst players beyond the number of best responses they perform.

Similar analyses in the pooled guessing games (Fig. 10) yields no significance differences in the low-vs.-random and low-vs.-high comparisons (χ^2 p -values of 0.185 and 0.769, respectively). We do find a significant difference in level distributions between random opponents and high-scoring opponents (p -value of 0.035), but the mean level against high-scoring opponents is actually lower than against random opponents (1.957 versus 2.121), and a Wilcoxon–Mann–Whitney test reveals no stochastic dominance of these level distributions (p -value 0.541), so we cannot claim that players' levels unambiguously shift up (or down) against stronger opponents.³⁷

In summary, we do see some subjects adjusting their realized levels against different opponents, particularly in the undercutting games. This indicates some responsiveness to signals about opponents, but neither the observed levels nor the quiz scores are useful in predicting *which* subjects will make this adjustment.

Result 7: The persistence of players' ordering of games

An alternative identifying restriction one might impose on the Level- k model is that the ranking of games be consistent between players. Formally, this would require that if $k_i(\gamma, \tau) \geq$

³⁷ Lehmann (1951) shows that this test can be viewed as a test of stochastic dominance, even with discrete distributions. Fay and Proschan (2010) provide an excellent survey of valid perspectives for this test.

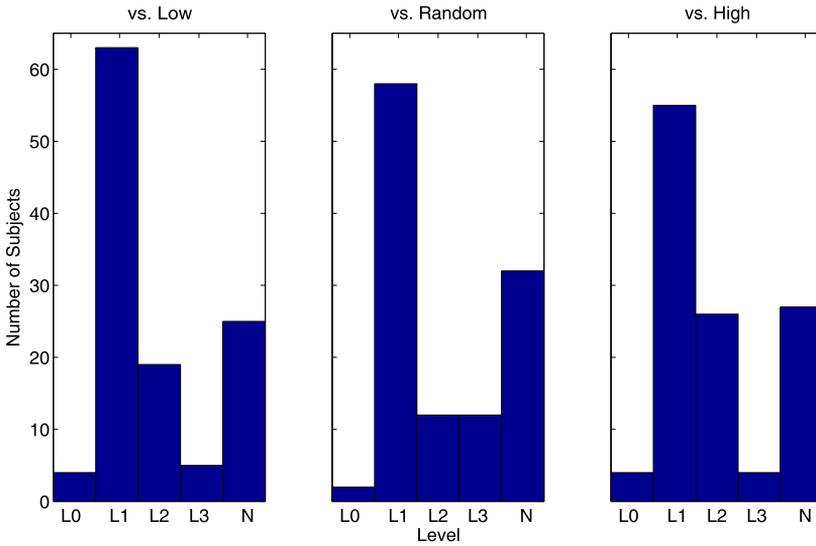


Fig. 10. Level distributions by opponent in the pooled guessing games.

Table 10

Observed frequency of game-rank switching among random pairs of subjects between randomly-drawn games, compared to the expected frequency under independently-drawn (i.i.d.) types.

	Frequency	i.i.d. prob.
Pooled UGs vs. pooled GGs		
Both change in same direction:	27.0%	24.9%
Both change in opposite directions:	29.1%	24.9%
Opposite/same ratio:	1.078	1.00
Undercutting games		
Both change in same direction:	9.6%	27.1%
Both change in opposite directions:	8.5%	27.1%
Opposite/same ratio:	0.884	1.00
Guessing games		
Both change in same direction:	26.6%	23.8%
Both change in opposite directions:	16.5%	23.8%
Opposite/same ratio:	0.618	1.00

$k_i(\gamma', \tau)$ for some i and γ then $k_j(\gamma, \tau) \geq k_j(\gamma', \tau)$ for all j . In this way the Level- k model could be thought of as providing a measure of (relative) game difficulty or complexity.

Table 10 shows the frequency with which a randomly-drawn pair of players changes levels in the same direction when moving between two randomly-chosen games, or in the opposite direction. These frequencies do not sum to one since pairs where at least one player does not switch levels between games are excluded. The reported frequencies are compared against the expected frequencies if levels were drawn independently from the empirical distribution of types.

Comparing the pooled undercutting games with the pooled guessing games, switches occur more often in the opposite direction than in the same direction. The ratio of switch directions is close to 1, which is what one would expect if levels were independently drawn in each game

family. A Monte Carlo simulation with 1000 samples shows that the empirical ratio of switch directions fails to reject the null hypothesis of independently-drawn levels, with a p -value of 0.380. Thus, the two families of games cannot be clearly ranked using estimated levels.

In the undercutting games we find some support for stability of game orderings. It is more likely that players switch levels in the same direction between games, as opposed to in the opposite direction. A Monte Carlo simulation shows that the ratio of switch directions is not consistent with the null hypothesis of independently-drawn levels, with a p -value of 0.026. Although this result is statistically significant, its usefulness is tempered by the fact that the vast majority of pairs have at least one player maintaining the same level between games. Thus, a fairly large sample of behavior would be needed to rank games based on observed levels. Analyzing the game-by-game directions of shifts indicates that UG3 is “easier” than the other three undercutting games. This is also evident from the fact that UG3 has substantially more Nash play than the others. The relative ranks of the other three games is ambiguous. Thus, the ability to rank the undercutting games seems to stem entirely from UG3.

Although the ratio of switch directions is lower in the guessing games, we cannot reject the null hypothesis that the ratio of switch directions was generated by independently-drawn levels—the Monte Carlo simulation yields a p -value of 0.070. This occurs because the level distributions vary more across guessing games, so the variance of switch directions under the null hypothesis is much larger.

8. Discussion

As a broad summary of our findings, the success of the Level- k model is mixed: We find very little cross-game stability when comparing the family of undercutting games with the family of two-person guessing games. We do find reasonably strong cross-game stability within the family of undercutting games, but zero stability in the two-person guessing games. Even in the undercutting games, however, observed levels are hard to predict with our five psychometric measures, except that Level-1 players may have a less keen awareness of others’ emotions and cognition. Finally, it appears that some players “step up” against stronger opponents in undercutting games, but we are unable to predict who makes this adjustment using either psychometric measures or observed levels.

Although ours is the first paper to thoroughly examine cross-game stability of individual levels, our conclusions about the success of the Level- k model are broadly consistent with the past literature. Many papers find strong support for Level- k play in certain games using behavioral data alone (Stahl and Wilson, 1994, 1995; Nagel, 1995; Duffy and Nagel, 1997; Ho et al., 1998; Bosch-Domènech et al., 2002; Camerer et al., 2004) or behavioral data augmented with lookup data (Costa-Gomes et al., 2001; Costa-Gomes and Crawford, 2006) or eye-tracking data (Chen et al., 2009; Wang et al., 2009). For some games, however, the Level- k model does not appear to organize the data well (Ivanov et al., 2009; Ivanov et al., 2010; Crawford and Iriberry, 2007a).³⁸ Shapiro et al. (2009) even find that the model’s fit can vary within a single game when different components of the payoff function are emphasized, with a better fit as the game becomes closer to a standard p -beauty contest and a worse fit as the game approaches the incomplete-information global game of Morris and Shin (2002). The broad con-

³⁸ Crawford and Iriberry (2007a) point out that the Level- k model fails to account for overbidding in second-price auctions.

clusion that emerges from this line of research is that the Level- k approach works well in some games, but not in others.

Camerer et al. (2004, p. 873) argue that “fitting a wide range of games turns up clues about where models fail and how to improve them.” Our research represents one such contribution, by demonstrating the varying individual-level robustness of Level- k models across two families of games. This suggests that the Level- k model may be one of many possible decision processes players employ to select strategies in novel games. Different processes may be triggered unconsciously in different settings, depending on features such as the characteristics of the game and the way in which the game is described. Beauty contests, simple matrix games, and our undercutting games all seem to trigger the Level- k heuristic in a large fraction of subjects, while its use appears infrequent in common-value auctions, global games, and endogenous-timing investment games. In two-person guessing games, Level- k reasoning may not be triggered unless subjects are given sufficient instruction and experience calculating best responses prior to play (see the online appendix).

Understanding the boundaries of the domain of applicability of the Level- k model means understanding when it is used, when it is not, and what factors trigger its use; this, in turn, increases the overall predictive power. At this point, we conjecture that Level- k play is triggered by simple, normal-form games of complete information, as well as in situations where the game’s instructions directly focus attention on calculating best responses, either directly through understanding tests or indirectly through framing effects. In other settings we expect less frequency of Level- k reasoning. These hypotheses give rise to a wide range of open questions that can be addressed in future work.

Given our conclusions, we suggest focusing behavioral research both on identifying distinct decision “heuristics” employed by people playing games *and* exploring their triggers. For example, Ivanov et al. (2009) identify plausible “rules of thumb” to explain their data when Level- k and quantal response equilibrium cannot; to what extent do these heuristics extend beyond the dynamic investment game they study? In our two-person guessing games, we do not identify an alternative heuristic that organizes the data since our analysis focuses on players’ estimated levels and not their actual strategies.

Finally, a multiple-heuristics model of strategic thinking implies that researchers should take care in extrapolating the success of any one model to out-of-sample strategic settings. Instead, future work should focus on understanding which heuristics are widely used and which features of a strategic environment trigger the use of different heuristics. We speculate that experimental protocols, training, and experience all have an impact on the choice of heuristic, and that the presentation of a game in matrix form (as in our undercutting games) is more likely to trigger best-response-based heuristics like the Level- k model.

Acknowledgments

We thank Doug Bernheim, Tilman Börgers, John Kagel, Stephen Leider, John Lightle, Tom Palfrey, Reinhard Selten, Dale Stahl, Joseph Tao-yi Wang, and Matthias Wibral for helpful comments. We are particularly indebted to Colin Camerer and Vince Crawford for their valuable comments and thoughtful suggestions. Financial support for this research was provided by the Ohio State University Department of Economics. The authors have no relevant or material financial interests that relate to the research described in this paper. Institutional Review Board (IRB) approval was obtained for this study.

Appendix A. Supplementary material

Supplementary material related to this article can be found online at <http://dx.doi.org/10.1016/j.jet.2015.07.012>.

References

- Altmann, Steffen, Falk, Armin, 2009. The impact of cooperation defaults on voluntary contributions to public goods. University of Bonn Working Paper.
- Arad, Ayala, Rubinstein, Ariel, 2012a. The 11–20 money request game: a level- k reasoning study. *Am. Econ. Rev.* 102 (7), 3561–3573.
- Arad, Ayala, Rubinstein, Ariel, 2012b. Multi-dimensional iterative reasoning in action: the case of the Colonel Blotto game. *J. Econ. Behav. Organ.* 84, 571–585.
- Ball, S.B., Bazerman, Max H., Carroll, J.S., 1991. An evaluation of learning in the bilateral winner's curse. *Organ. Behav. Hum. Decis. Process.* 48, 1–22.
- Baron-Cohen, Simon, 1990. Autism: a specific cognitive disorder of 'mind-blindness'. *Int. Rev. Psychiatry* 2 (1), 81–90.
- Baron-Cohen, Simon, Jolliffe, Therese, Mortimore, Catherine, Robertson, Mary, 1997. Another advanced test of theory of mind: evidence from very high functioning adults with autism or Asperger syndrome. *J. Child Psychol. Psychiatry* 38, 813–822.
- Baron-Cohen, Simon, Wheelwright, Sally, Hill, Jacqueline, Raste, Yogini, Plumb, Ian, 2001. The "reading the mind in the eyes" test revised version: a study with normal adults, and adults with Asperger syndrome or high-functioning autism. *J. Child Psychol. Psychiatry* 42, 241–251.
- Batzilis, D., Jaffe, S., Levitt, J.A., List, J., Picel, J., 2013. Large-scale analysis of level- k thinking: Facebook's Roshambull. University of Chicago Working Paper.
- Bhatt, Meghana, Camerer, Colin F., 2005. Self-referential thinking and equilibrium as states of mind in games: fMRI evidence. *Games Econ. Behav.* 52 (2), 424–459.
- Bosch-Domènech, Antoni, Garcia-Montalvo, Jose, Nagel, Rosemarie C., Satorra, Albert, 2002. One, two, (three), infinity. . . : newspaper and lab beauty-contest experiments. *Am. Econ. Rev.* 92 (5), 1687–1701.
- Brocas, Isabelle, Camerer, Colin, Carrillo, Juan D., Wang, Stephanie W., 2014. Imperfect choice or imperfect attention? Understanding strategic thinking in private information games. *Rev. Econ. Stud.* 81 (3), 944–970.
- Brown, Alexander L., Camerer, Colin F., Lovo, Dan, 2012. To review or not to review? Limited strategic thinking at the movie box office. *Am. Econ. J. Microecon.* 4, 1–26.
- Bругуier, Antoine J., Quartz, Steven R., Bossaerts, Peter L., 2010. Exploring the nature of "trading intuition". *J. Finance* 65, 1703–1723.
- Burchardi, K., Penczynski, S., 2014. Out of your mind: eliciting individual reasoning in one shot games. *Games Econ. Behav.* 84, 39–57. <http://dx.doi.org/10.1016/j.geb.2013.12.005>.
- Burnham, Terence C., Cesarini, David, Johannesson, Magnus, Lichtenstein, Paul, Wallace, Bjorn, 2009. Higher cognitive ability is associated with lower entries in a p -beauty contest. *J. Econ. Behav. Organ.* 72, 171–175.
- Camerer, Colin F., 2003. *Behavioral Game Theory*. Princeton University Press, Princeton, NJ.
- Camerer, Colin F., Ho, Teck-Hua, 1998. Experience-weighted attraction learning in coordination games: probability rules, heterogeneity and time-variation. *J. Math. Psychol.* 42, 305–326.
- Camerer, Colin F., Ho, Teck-Hua, Chong, Juin-Kuan, 2004. A cognitive hierarchy model of games. *Q. J. Econ.* 119 (3), 861–898.
- Chen, Chun-Ting, Huang, Chen-Ying, Wang, Joseph Tao-yi, 2009. A window of cognition: eyetracking the reasoning process in spatial beauty contest games. National Taiwan University Working Paper.
- Chong, Juin-Kuan, Camerer, Colin F., Ho, Teck-Hua, 2005. Cognitive hierarchy: a limited thinking theory in games. In: Zwick, Rami, Rapoport, Amnon (Eds.), *Experimental Business Research: Marketing, Accounting and Cognitive Perspectives*, vol. III. Springer, The Netherlands, pp. 203–228. chapter 9.
- Cohen, Ronald A., Salloway, Stephen, Sweet, Lawrence H., 2007. Neuropsychiatric aspects of disorder of attention. In: *American Psychiatric Press Textbook of Neuropsychiatry*. American Psychiatric Publishing, Arlington, Virginia, pp. 405–444. chapter 10.
- Coricelli, Giorgio, Nagel, Rose, 2009. Neural correlates of depth of strategic reasoning in medial prefrontal cortex. *Proc. Natl. Acad. Sci. USA*. <http://dx.doi.org/10.1073/pnas.080721106>.
- Costa-Gomes, Miguel, Crawford, Vincent P., 2006. Cognition and behavior in two-person guessing games: an experimental study. *Am. Econ. Rev.* 96 (5), 1737–1768.

- Costa-Gomes, Miguel, Crawford, Vincent P., Broseta, Bruno, 2001. Cognition and behavior in normal-form games: an experimental study. *Econometrica* 69 (5), 1193–1235.
- Crawford, Vincent P., 1995. Adaptive dynamics in coordination games. *Econometrica* 63, 103–143.
- Crawford, Vincent P., Iriberri, Nagore, 2007a. Level- k auctions: can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions? *Econometrica* 75 (6), 1721–1770.
- Crawford, Vincent P., Iriberri, Nagore, 2007b. Fatal attraction: salience, naivete, and sophistication in experimental "hide-and-seek" games. *Am. Econ. Rev.* 97 (5), 1731–1750.
- Crawford, Vincent P., Costa-Gomes, Miguel A., Iriberri, Nagore, 2013. Structural models of nonequilibrium strategic thinking: theory, evidence, and applications. *J. Econ. Lit.* 51, 5–62.
- De Sousa, J., Hollard, G., Terracol, A., 2012. Cognitive ability and learning to play equilibrium: a level- k analysis. Working Paper.
- Devetag, Giovanna, Warglien, Massimo, 2003. Games and phone numbers: do short-term memory bounds affect strategic behavior? *J. Econ. Psychol.* 24, 189–202.
- Duffy, John, Nagel, Rosemarie C., 1997. On the robustness of behavior in experimental 'beauty contest' games. *Econ. J.* 107, 1684–1700.
- Erev, Ido, Roth, Alvin E., 1998. Prediction how people play games: reinforcement learning in games with unique strategy equilibrium. *Am. Econ. Rev.* 88, 848–881.
- Fay, Michael P., Proschan, Michael A., 2010. Wilcoxon–Mann–Whitney or t-test? On assumptions for hypothesis tests and multiple interpretations of decision rules. *Stat. Surv.* 4, 1–39.
- Frederick, Shane, 2005. Cognitive reflection and decision making. *J. Econ. Perspect.* 19, 25–42.
- Georganas, Sotiris, 2011. English auctions with resale: an experimental study. *Games Econ. Behav.* 73, 147–166.
- Gigerenzer, Gerd, 2001. The adaptive toolbox. In: Gigerenzer, Gerd, Selten, Reinhard (Eds.), *Bounded Rationality: The Adaptive Toolbox*. MIT Press, London. chapter 3.
- Gill, David, Prowse, Victoria, 2013. Cognitive ability and learning to play equilibrium: a level- k analysis. Oxford University Working Paper.
- Grosskopf, Brit, Nagel, Rosemarie, 2008. The two-person beauty contest. *Games Econ. Behav.* 62, 93–99.
- Grosskopf, Brit, Bereby-Meyer, Yoella, Bazerman, Max H., 2007. On the robustness of the winner's curse phenomenon. *Theory Decis.* 63, 389–418.
- Harsanyi, John C., 1967. Games with incomplete information played by "Bayesian" players, I–III. Part I: The basic model. *Manag. Sci.* 14, 159–182.
- Ho, Teck-Hua, Su, Xuanming, 2013. A dynamic level- k model in sequential games. *Manag. Sci.* 59, 452–469.
- Ho, Teck-Hua, Camerer, Colin F., Weigelt, Keith, 1998. Iterated dominance and iterated best response in p -beauty contests. *Am. Econ. Rev.* 88, 947–969.
- Ivanov, Asen, Levin, Dan, Peck, James, 2009. Hindsight, foresight, and insight: an experimental study of a small-market investment game with common and private values. *Am. Econ. Rev.* 99 (4), 1484–1507.
- Ivanov, Asen, Levin, Dan, Niederle, Muriel, 2010. Can relaxation of beliefs rationalize the winner's curse?: An experimental study. *Econometrica* 78, 1435–1452.
- Kawagoe, Toshiji, Takizawa, Hirokazu, 2009. Equilibrium refinement vs. level- k analysis: an experimental study of cheap-talk games with private information. *Games Econ. Behav.* 66 (1), 238–255.
- Kawagoe, Toshiji, Takizawa, Hirokazu, 2012. Level- k analysis of experimental centipede games. *J. Econ. Behav. Organ.* 82, 548–566.
- Keynes, John M., 1936. *The General Theory of Interest, Employment and Money*. Macmillan, London.
- Kneeland, Terri, 2012. Coordination under limited depth of reasoning. University of British Columbia Working Paper.
- Lehmann, Erich Leo, 1951. Consistency and unbiasedness of certain nonparametric tests. *Ann. Math. Stat.* 22, 165–179.
- McKelvey, Richard D., Palfrey, Thomas R., 1995. Quantal response equilibria for normal-form games. *Games Econ. Behav.* 10 (1), 6–38.
- McKelvey, Richard D., Palfrey, Thomas R., 1998. Quantal response equilibria for extensive-form games. *Exp. Econ.* 1, 9–41.
- Morris, Stephen, Shin, Hyun Song, 2002. Social value of public information. *Am. Econ. Rev.* 92, 1521–1534.
- Nagel, Rosemarie C., 1993. Experimental results on interactive competitive guessing. Discussion Paper 8-236, Sonderforschungsbereich 303. Universitat Bonn.
- Nagel, Rosemarie C., 1995. Unraveling in guessing games: an experimental study. *Am. Econ. Rev.* 85 (5), 1313–1326.
- Ostling, Robert, Wang, Joseph Tao-yi, Chou, Eileen, Camerer, Colin F., 2011. Testing game theory in the field: Swedish LUPI lottery games. *Am. Econ. J. Microecon.* 3, 1–33.
- Penczynski, Stefan, 2011. Strategic thinking: the influence of the game. Working Paper.
- Rogers, Brian W., Palfrey, Thomas R., Camerer, Colin F., 2009. Heterogeneous quantal response equilibrium and cognitive hierarchies. *J. Econ. Theory* 144, 1440–1467.

- Rubinstein, Ariel, 2014. A typology of players: between instinctive and contemplative. Tel Aviv University Working Paper.
- Samuelson, William F., Bazerman, Max H., 1985. The winner's curse in bilateral negotiations. In: Smith, Vernon L. (Ed.), *Research in Experimental Economics*, vol. 3. JAI Press, Greenwich, CT, pp. 105–137.
- Shapiro, Dmitry, Shi, Xianwen, Zillante, Artie, 2009. Robustness of level- k reasoning in generalized beauty contest games. University of North Carolina Working Paper.
- Siegel, Sidney, Castellan Jr., N. John, 1988. *Nonparametric Statistics for the Behavioral Sciences*, 2nd edition. McGraw-Hill, New York, NY.
- Stahl, Dale O., Wilson, Paul O., 1994. Experimental evidence on players' models of other players. *J. Econ. Behav. Organ.* 25 (3), 309–327.
- Stahl, Dale O., Wilson, Paul W., 1995. On players' models of other players: theory and experimental evidence. *Games Econ. Behav.* 10, 218–254.
- Strzalecki, Tomasz, 2009. Depth of reasoning and higher-order beliefs. Harvard University Working Paper.
- Wang, Joseph Tao-yi, Spezio, Michael, Camerer, Colin F., 2009. Pinocchio's pupil: using eyetracking and pupil dilation to understand truth telling and deception in sender–receiver games. *Am. Econ. Rev.* 100 (3), 984–1007.
- Wechsler, David, 1939. *Measurement of Adult Intelligence*. Williams & Wilkins, Baltimore.