

A Cognitive Hierarchy Theory of One-shot Games

Colin F. Camerer¹

California Institute of Technology
Pasadena, CA 91125
camerer@hss.caltech.edu

Teck-Hua Ho

Haas School of Business
University of California-Berkeley
Berkeley CA 94720
hoteck@haas.berkeley.edu

Juin-Kuan Chong

The NUS Business School
National University of Singapore
bizcjk@nus.edu.sg

December 23, 2002

¹This research was supported by NSF grants SES-0078911 and SES-0078853. Thanks to C. Mónica Capra, Haitao Cui, Paul Glimcher, and Roger Myerson. Former math nerd Matthew Rabin directed our attention to the golden ratio. Ming Hsu and Brian Rogers provided excellent research assistance. Useful comments were received from seminars at Caltech, Chicago, New York University, Pittsburgh, the Nobel Symposium in Sweden (December 2001), Columbia, and Berkeley.

Abstract

Strategic thinking, best-response, and mutual consistency (equilibrium) are three key modeling principles in noncooperative game theory. This paper relaxes mutual consistency to predict how players are likely to behave in one-shot games before they can learn to equilibrate. We introduce a one-parameter cognitive hierarchy (CH) model to predict behavior in one-shot games, and initial conditions in repeated games. The CH approach assumes that players use k steps of reasoning with frequency $f(k)$. Zero-step players randomize. Players using k (≥ 1) steps best respond given partially rational expectations about what players doing 0 through $k - 1$ steps actually choose. A simple axiom which expresses the intuition that steps of thinking are increasingly constrained by working memory, implies that $f(k)$ has a Poisson distribution (characterized by a mean number of thinking steps τ). The CH model converges to dominance-solvable equilibria when τ is large, predicts monotonic entry in binary entry games for $\tau < 1.25$, and predicts effects of group size which are not predicted by Nash equilibrium. Best-fitting values of τ have an interquartile range of (.98,2.21) and a median of 1.55 across 60 experimental samples of matrix games, entry games and mixed-equilibrium games. The CH model also has economic value because subjects would have raised their earnings substantially if they had best-responded to model forecasts instead of making the choices they did.

1 Introduction

Noncooperative game theory uses three distinct concepts to make precise predictions of how people will, or should, interact strategically: Formation of beliefs based on analysis of what others might do (strategic thinking); choosing a best response given those beliefs (optimization); and adjustment of best responses and beliefs until they are mutually consistent (equilibrium). Standard equilibrium models combine all three features.

The strong assumption of mutual consistency can be reasonably defended on the grounds that *some* modeling device is necessary to ‘close’ the model by specifying a players’ beliefs; forcing beliefs to match likely choices is one reasonable way to close it. Mutual consistency can also be sensibly justified as a mathematical shortcut which represents the result of some unspecified learning or evolutionary adjustment process.² However, the learning or evolutionary justifications logically imply that beliefs and choices will *not* be consistent if players do not have time to learn or evolve. That leaves a large hole in game theory: Viz., how will people behave *before* equilibration to mutual consistency takes place? This question is important because many games occur between unfamiliar rivals, and because the way in which play starts probably influences the long-run path of play when there are multiple equilibria.

This paper introduces a cognitive hierarchy (CH) model which weakens mutual consistency but retains the concepts of strategic thinking (to a limited degree) and optimization. The model is closed by specifying a hierarchy of decision rules and the frequencies with which players stop at different steps of the hierarchy. The model is intended to predict what players do in one-shot games, and to supply initial conditions for dynamic learning models. It is parameterized by one parameter (τ), which is the average number of steps of thinking. Axioms and estimation across four experimental data sets suggest that plausible values of τ are between 1 and 2.

Quantal response equilibrium (QRE) is a complementary approach, which retains mutual consistency but relaxes optimization (Rosenthal, 1989; Chen, Friedman and Thisse, 1996; McKelvey and Palfrey, 1995, 1998; Goeree and Holt, 1999). QRE weakens the best-response property in Nash equilibrium and substitutes statistically-rational expectations in the sense that a player’s beliefs about the distribution of play by others matches the

²Weibull (1995) notes that in Nash’s thesis proposing a concept of equilibrium, Nash himself suggested equilibrium might arise from some “mass action” which adapted over time.

actual distribution. We believe QRE and CH models should be treated as alternative paths and both deserve exploration.

The CH model illustrates how “behavioral game theory” is done (e.g., Camerer, 2003). In behavioral game theories, psychological regularities and empirical data are used to suggest parsimonious ways to weaken assumptions of rationality, equilibrium, and self-interest.

The CH model is guided by the same aesthetic criteria that motivate analytical game theorists— viz., generality, precision, and theoretical usefulness. The theory is *general* because it can be applied to all one-shot games. (Extension to repeated games is a challenge for future work.) The theory is *precise* because it predicts a specific distribution of strategy frequencies once one parameter is specified. (In fact, in games with multiple equilibria it is *more* precise than Nash equilibrium and many equilibrium refinements.) And the theory is simple enough that mathematical analysis can be used to derive some interesting *theoretical* implications.

The CH approach also strives to meet two other criteria which many ideas in analytical game theory do not: It is *cognitive*, and meant to *predict* behavior accurately. That is, the steps of thinking players do in the cognitive hierarchy are meant to be taken seriously as reduced-form outputs of some cognitive mechanism. The theory can therefore be tested with cognitive data such as self-reports, tests of memory, response times, measures of eye gaze and attention (Camerer et al., 1994; Costa-Gomes, Crawford and Broseta, 2001), or even brain imaging (cf. Camerer, Loewenstein, and Prelec, 2002).

Our approach is also heavily disciplined by data. The data reported in this paper are experimental. Because game-theoretic predictions are notoriously sensitive to what players know, when they move, and what their payoffs are, laboratory environments enable good control of these crucial variables (see Crawford, 1997) and hence provide sharp tests of theoretical predictions. As in all sciences with a laboratory component, of course, the research program hones models sharply on lab data, in order to choose good candidate models to eventually explain naturally-occurring field phenomena.

The CH model is designed to be a useful empirical competitor to Nash equilibrium in three ways: First, CH should be able to capture deviations when equilibrium behavior does not occur. An example is behavior in dominance-solvable games. In experimental studies of these games, most players *do* think strategically, but they do only one or two steps of iterated reasoning and hence do not reach an equilibrium in which choices are

mutually consistent (Camerer, 2003, chapter 5). The CH model accounts reasonably well for deviations like those in dominance-solvable games.

Second, the CH model should reproduce the success of Nash equilibrium in games where Nash fits well. For example, in games with mixed equilibria, Nash equilibrium approximates some aspects of behavior surprisingly well, even in one-shot games with no opportunities to learn. In these games, it appears as if a population mixture of players using different pure strategies ("purification") can roughly approximate Nash equilibrium. Since the equilibrium model works mysteriously well in these games, the goal of CH is to offer a clue to a cognitive process that creates purification and instant near-equilibration.

Third, in many interesting games (perhaps most) there are multiple Nash equilibria. Less plausible equilibria are typically 'refined' away by positing additional restrictions (such as subgame or trembling-hand perfection, and selection principles like risk- or payoff-dominance). The CH model is another solution to the problem of refinement. Since it always yields an unique solution, it solves the multiplicity problem. The key insight is that multiplicity of equilibria arise *because* of the assumption of mutual consistency. Since the CH model does not impose mutual consistency, it does not lead to multiplicity—in effect, a model of the process of thinking acts as a statistical selection principle (cf. Harsanyi and Selten, 1988). Ironically, in strategic situations a model with less (mutual) rationality can be *more* precise (cf. Lucas, 1986³). In extensive-form games, refinement of the Nash concept is needed to eliminate equilibria which rest on incredible threats (hence subgame perfection) and odd beliefs after surprising events (hence trembling-hand perfection). In the CH model, every strategy is chosen with positive probability. So incredible threats and odd beliefs never arise.

The paper is organized as follows. The next section describes the CH model, and discusses both precursors and alternative specifications. Section III collects some theoretical results. Section IV reports estimation of the τ parameter from four classes of games. Section V explores the prescriptive economic value of the CH theory (and some other theories), by calculating whether subjects would have earned more money if they had used the CH model to forecast, rather than making their own choices. Section VI

³Lucas (1986) makes the same point in macroeconomic models. Rational expectations often yields indeterminacy whereas adaptive expectations pins down a dynamic path. Importantly, Lucas also calls for experiments as a way to supplement intuition about which dynamics are likely to occur, and help explain why.

notes how the CH model can account for cognitive details. Section VII concludes and points out directions for further research.

2 The cognitive hierarchy (CH) model

First, notation. Players are indexed by i and strategies by j and j' . Player i has m_i strategies denoted by s_i^j . Other players are denoted by $-i$. Denote other players' strategies by $s_{-i}^{j'}$, and player i 's payoffs by $\pi_i(s_i^j, s_{-i}^{j'})$.

We will denote a player's position in the cognitive hierarchy (the number of steps or steps of thinking she does) by k and a k -step player's expected payoffs (given her beliefs) by $E_k(s_i^j)$. Denote the actual frequency of k step players by $f(k)$.

A precise thinking steps theory needs two components: Decision rules for what players using each step of thinking k in the cognitive hierarchy will do, and a distribution of thinking steps $f(k)$.

2.1 Decision rules for different thinking steps

We assume that 0 step players are not thinking strategically at all; they randomize equally across all strategies. Other simple rules could be used to start the cognitive hierarchy process off, but equal randomization has some empirical and theoretical advantages.⁴ Zero-step thinkers may also be “unlucky” rather than “dumb”. Players who start to analyze the game carefully but get confused or make an error might make a choice that appears random and far from equilibrium (much as a small algebra slip in a long proof can lead to a bizarre result). Denote the choice probability of step k for strategy s_{-i}^j by $P_k(s_{-i}^j)$. So, we have $P_0(s_{-i}^j) = \frac{1}{m_{-i}}$ in a 2-player game.⁵

⁴Equal randomization implies that all strategies are chosen with positive probability. This is helpful for empirical work because it means all strategies will have positive predicted probabilities, so there is no zero likelihood problem when using maximum likelihood estimation. This also liberates us to assume best response by players using more steps of thinking (rather than stochastic response). For theoretical work, having all strategies chosen with positive probability solves two familiar problems— eliminating incredible threats (since all threats are “tested”) as subgame perfection does; and eliminating *ad hoc* rules for Bayesian updating after zero probability events (since there are no such events).

⁵In a n -player game, the step-0 probability is a $n - 1$ multinomial expression.

Players doing one or more steps of thinking are assumed to not realize that others are thinking as ‘hard’ as they are (or harder), but they have an accurate guess about the relative proportions of players using fewer steps than they do. Formally, players at step k know the true proportions $f(0), f(1), \dots, f(k-1)$. Since these proportions do not add to one, they normalize them by dividing by their sum. That is, step- k players believe the proportions of players doing h steps of thinking are $g_k(h) = f(h) / \sum_{l=0}^{k-1} f(l), \forall h < k$ and $g_k(h) = 0, \forall h \geq k$.

Given these beliefs, the expected payoff to a k -step thinker from strategy s_i^j is $E_k(\pi_i(s_i^j)) = \sum_{j'=1}^{m-i} \pi_i(s_i^j, s_{-i}^{j'}) \{ \sum_{h=0}^{k-1} g_k(h) \cdot P_h(s_{-i}^{j'}) \}$. For simplicity, assume players best-respond (or randomize equally if two or more strategies have identical expected payoffs).

The normalized-beliefs assumption $g_k(h) = f(h) / \sum_{l=0}^{k-1} f(l)$ has an interesting property we call “increasingly rational expectations”. To see what this means, first note that the absolute total deviation of step- k ’s beliefs and true frequencies is $D(k) = \sum_{h=0}^{\infty} |f(h) - g_k(h)|$. Then consider how large this total deviation is for players at different levels of the cognitive hierarchy.

Zero-step thinkers have no beliefs at all. One-step thinkers believe everyone is doing 0 steps of thinking (i.e., $g_1(0) = 1$); since only $f(0)$ are doing 0 steps of thinking the one-step beliefs are wrong by a total absolute deviation of $D(1) = 1 - f(0) + \sum_{h=1}^{\infty} (f(h) - 0) = 2 \cdot (1 - f(0))$. Two-step thinkers believe $g_2(0) = \frac{f(0)}{f(0)+f(1)}$ and $g_2(1) = \frac{f(1)}{f(0)+f(1)}$. Since the actual frequencies are $f(0)$ and $f(1)$ the sum of the deviations between their beliefs and the true frequencies is $D(2) = g_2(0) - f(0) + g_2(1) - f(1) + \sum_{h=2}^{\infty} (f(h) - 0)$. A little algebra shows that this total deviation is $D(2) = 2 \cdot (1 - f(0) - f(1))$, which is smaller than the size of the 1-step thinkers’ belief error, $D(1)$. In fact, it is easy to show⁶ that the total deviation $D(k)$ falls monotonically as k increases. The reason is that the “missing” belief $\sum_{h=k}^{\infty} f(h)$ which is reallocated by the step- k thinker to the lower-step types shrinks as k grows large. The k -step thinkers’ beliefs gradually come closer and closer to the truth.

⁶Use the sum of the absolute deviations to measure the distance of the normalized distributions from the true distribution. The total absolute deviation for $k \geq 1$ is:

$$D(k) = \sum_{h=0}^{k-1} \left[\frac{f(h)}{\sum_{l=0}^{k-1} f(l)} - f(h) \right] + \sum_{h=k}^{\infty} [f(h) - 0] \quad (2.1)$$

Algebra shows that this is $D(k) = 2 \cdot [1 - \sum_{h=0}^{k-1} f(h)]$. $D(k)$ is decreasing in k - so beliefs get closer and closer to the truth- and $\lim_{k \rightarrow \infty} D(k) = 0$ because $\sum_{h=0}^{\infty} f(h) = 1$.

The fact that beliefs converge as k grows large has another important implication: As the missing belief grows small, players who are doing k and $k + 1$ steps of thinking will have approximately the same beliefs, and will therefore have approximately the same expected payoffs. That is, it doesn't pay to think too hard, because doing k steps and $k+1$ steps yields roughly the same expected payoff.

2.2 Principles underlying distributions $f(k)$

So far we have said nothing about the distribution $f(k)$. Denote the average number of steps of thinking by τ . By definition $\sum_{h=0}^{\infty} f(h) = 1$ and $\sum_{h=0}^{\infty} h \cdot f(h) = \tau$. Our approach is to derive a parsimonious distribution from axioms, and use both further axioms and empirical estimation to pin the distribution's mean down further. A more empirical approach is to allow $f(0), f(1), \dots, f(k)$ to be free parameters up to some reasonable k , and estimate each one separately (cf. Nagel, 1995; Stahl and Wilson, 1995; Ho, Camerer and Weigelt, 1998; Nagel et al, 2002). We proceed by stating three sensible principles $f(k)$ should satisfy and seeing what distributions satisfy them. They are:

1. Discreteness: Because the steps of reasoning are discrete, it is convenient if the distribution is discrete too (i.e., it only puts probability mass on integer values; cf. Stahl, 1998).
2. Unimodality: It is likely that most players are doing *some* degree of strategic thinking (so zero is not the mode), but constraints on working memory will constrain players from doing many steps of thinking (and may be unprofitable at the margin). The first two principles imply $\frac{f(k)}{f(k-1)}$ should be *greater* than one for low k and *less* than one for high k .
3. Convexity: Let k^* be the most common type. The convexity principle requires that for $k > k^*$, $\frac{f(k+2)}{f(k+1)} < \frac{f(k+1)}{f(k)}$ (upper convexity) and for $k < k^*$, $\frac{f(k-1)}{f(k-2)} > \frac{f(k)}{f(k-1)}$ (lower convexity). The upper convexity condition implies that the distribution $f(k)$ drops off rapidly for high k . (In Keynes's famous passage on the stock market as a beauty contest, he guesses that "there are some, I believe, who practise the fourth, fifth and higher degrees [of reasoning about reasoning]" (p156); his wording— "some"— suggests Keynes thinks that not many investors do that much thinking.) Rapid drop-off also means computations can be truncated at a modest number of steps

of thinking (e.g., 8), and the results normalized, with a tiny loss in precision. The lower convexity condition is useful for theorizing when k^* is large. It creates a kind of separability: Players doing k steps will believe (almost) all players are just one step below them, which means they best-respond to a single strategy rather than a mixture of strategies across steps which depends on τ .

Unimodality and convexity are both satisfied by $f(k)/f(k-1) \propto 1/k \rightarrow f(k)/f(k-1) = \tau/k$. Among discrete distributions, this property holds if and only if the distribution $f(k)$ is Poisson⁷, $f(k) = e^{-\tau} \cdot \frac{\tau^k}{k!}$. In addition, the Poisson distribution has only one parameter, τ , which is its mean *and* its variance. This simplicity has obvious advantages in estimation.

One way to test the convexity condition is to compare fits from the Poisson model and a general CH model in which the frequencies of k -step thinkers, $f(k)$, are not constrained to satisfy the convexity condition while keeping the discreteness and unimodality conditions. We facilitate estimation, we truncate the model at a maximum of seven steps resulting in a 7-parameter model. Except for the Stahl-Wilson data, imposing the convexity condition as in the Poisson model degrades fit very little compared to the general CH model.⁸

2.3 Plausible values of τ

What are reasonable values of τ ? Our approach is deduce some values from principles, and also estimate them from data, and hope the deduced and estimated values are not too far apart.

Since the Poisson distribution has only one parameter, restrictions on proportions $f(k)$ can pin down the value of τ exactly. This enables us to directly links intuitions about $f(k)$ to values of τ .

⁷Since $\frac{f(k)}{f(k-1)} = \frac{\tau}{k}$, $\frac{f(k+1)}{f(k)} = \frac{\tau}{k+1}$, and k^* is the largest integer that is lower or equal to τ , the upper and lower convexity conditions follow naturally.

⁸Across the five data sets reported below, the reduction in LL is only 40, 0, 1, 14, and 0 points for the Poisson model. In addition, the fractions of players estimated to use each level in the general specification are reasonably close to those constrained by the Poisson distribution. See Camerer, Ho and Chong, 2002b, for details.

For example, suppose k is Poisson-distributed, and 1-step thinking is the most common. Then it follows that $\tau \in (1, 2)$. Now suppose the frequencies of zero- and two-step thinking are equal ($f(0) = f(2)$). Then $\tau = \sqrt{2}$ again.

Two other interesting restrictions are

$$f(0) + f(1) = \sum_{j=2}^{\infty} f(j) \quad (2.2)$$

$$f(2) = \sum_{j=3}^{\infty} f(j) \quad (2.3)$$

The first restriction says that the amount of nonstrategic (step 0) or not-very-strategic (step 1) thinking is equal to the amount of truly strategic thinking (step 2 and above). The second restriction says that two steps of thinking and the sum of all higher steps are equally common. If k is Poisson-distributed, the two properties together imply that τ equals $\frac{\sqrt{5}+1}{2} \approx 1.618$, a remarkable constant known as the “golden ratio” (usually denoted Φ)⁹. The golden ratio is equal to the limit of the ratios of adjacent numbers in the Fibonacci sequence, and is often used in architecture because rectangles with golden ratio proportions are aesthetically pleasing.

The other way to pick a value of τ is to estimate it from many data sets. Camerer (2003, chapter 5) surveyed experiments on dominance-solvable games and suggested that 1-2 steps of thinking are typical¹⁰. Section III reports formal estimation from a wide variety of one-shot games (60 games in total). Most estimates are between 1 and 2; the median across all 60 games is 1.55.

⁹Condition (2.2) implies that

$$1 + \tau = \sum_{j=2}^{\infty} \frac{\tau^j}{j!} = e^{\tau} - (1 + \tau) \quad (2.4)$$

or equivalently, $1 + \tau = \frac{e^{\tau}}{2}$, which gives $\tau = 1.68$. Condition (2.3) implies that

$$\frac{\tau^2}{2} = \sum_{j=3}^{\infty} \frac{\tau^j}{j!} = e^{\tau} - (1 + \tau + \tau^2/2) \quad (2.5)$$

or equivalently, $e^{\tau} = (1 + \tau + \tau^2)$, which gives $\tau = 1.8$. The two conditions together imply $f(0) + f(1) = f(2) + f(2)$ (since $\sum_{j=2}^{\infty} f(j) = f(2) + \sum_{j=3}^{\infty} f(j)$), which is equivalent to $1 + \tau = \tau^2$ which gives $\tau = \Phi$.

¹⁰See also Nagel, 1995; Stahl and Wilson, 1995; Ho, Camerer, and Weigelt, 1998; Costa-Gomes, Crawford, and Broseta, 2001.

2.4 Early models of limited thinking

The CH approach is a natural outgrowth of earlier efforts. Brown (1951) and Robinson (1951) suggested a kind of “fictitious play” as a model of the sort of mental tatonnement or iterative algorithm that could lead to Nash equilibrium.¹¹ In their model, a player starts with a prior belief about what others will choose, and best-responds to that belief. Players then take into account their own reasoning and best-respond to a mixture of prior belief and the behavior generated by their earlier response at the first step. This process iterates to convergence. (See also Harsanyi, 1975 for the so-called “tracing procedure”.)

In our terminology, the original fictitious play model is equivalent to one in which $f(k) = 1/N$ for N steps of thinking, and $N \rightarrow \infty$. Fictitious play was reinterpreted as a real-time learning model by Fudenberg and Kreps (1990) (and later Fudenberg and Levine, 1998) by mapping steps of iteration in a single player’s reasoning into actual periods of play in repetitions of a stage game. Our approach is a return to the original interpretation of fictitious play, except that instead of a single player iterating repeatedly until a fixed point is reached, and taking his earlier tentative decisions as pseudo-data, we posit a population of players in which a fraction $f(k)$ of players stop cold after k steps of thinking. Much like a fast-moving film can be slowed down to show its individual frames “frozen” one by one, the cognitive hierarchy approach assumes that different players freeze at different (finite) steps in the iteration process, rather than assuming that the full reasoning process occurs in all players’ brains before the first period of play.

In 1984, Bernheim (1984) and Pearce (1984) relaxed the requirement of mutual consistency by introducing a coarsening of Nash equilibrium, *rationalizability*. Strategies are rationalizable if they are best responses given some beliefs, and beliefs must respect rationalizability by others, which eliminates strategies that are iteratedly dominated. As in fictitious play, they implicitly assume rationalizability is a process that occurs within a single players’ beliefs. But because they put little structure on where the reasoning process stops, rationalizability does not yield a precise prediction in many games.

The second step in the cognitive hierarchy is the idea that many players respond to

¹¹The fictitious play algorithm always converges to Nash equilibrium in 2x2 games (Robinson, 1951), zero-sum games (Miyasawa, 1961), games solvable by strict dominance (Nachbar, 1990) and some games with strategic complements (Krishna and Sjoström, 1998). Shapley (1964) upset the hope for fictitious play as a general cognitive underpinning for equilibrium with a 3x3 example in which fictitious play cycles around a mixed-strategy equilibrium.

a diffuse, or ‘ignorance’ prior about what others might do. This principle can be traced at least to Laplace. The appeal of 1-step decision rules in games was noted by Camerer (1990), who hypothesized that players in games treat their choices as decisions, and do not reason very strategically about what other players would do (see also Kadane and Larkey, 1982). Banks, Camerer and Porter (1994) also focused on the 1-step rule in trying to explain departures from equilibrium in signaling games. Haruvy and Stahl (1998) found that the 1-step rule is a more robust and useful prediction of behavior in one-shot games than other rules like minimax, maximax and (Nash) equilibrium.

Truncating iterations beyond the a small finite step was suggested by Binmore (1988) and Stahl (1993). Selten (1998, p. 421) argued that

the natural way of looking at game situations...is not based on circular concepts, but rather on a step-by-step reasoning procedure.

The first modern applications to experimental data were done by Nagel (1995) and Stahl and Wilson (1995). Nagel used the simple $k - 1$ model, in which all players think others are using one fewer steps of reasoning than they themselves are. She classified players into thinking steps using the absolute distance of choices from the nearest spike of data in dominance-solvable “ p -beauty contest games” (which are explored further below). Ho, Camerer and Weigelt (1998) used a more sophisticated procedure to classify players, allowing stochastic response and explored a wider range of data, and basically corroborated Nagel’s finding that only a couple of steps of thinking were being used. Stahl and Wilson (1995) posited a mixture of steps of thinking along with other types (e.g., Nash equilibrium types and “worldly” types who best respond to the distribution of all other types— these are equivalent to our highest-step types), using a total of 12 parameters.¹²

¹²Two other recent models are also related. Capra (1999) proposed a model of thinking steps in which players imagine cycles consisting of a move, an opponent’s best response, and their own best response to the opponent’s best response which coincides with the initial posited move. But responses are actually stochastic; so the model produces probabilities of each possible cycle; summing over them gives predicted probabilities of each considered move. The model is cognitively appealing and fits experimental data well (see Cabrera, Capra, and Gomez, 2002) but is difficult to compute in large games. Goeree and Holt (2002) propose a two-parameter model of ‘noisy introspection’ in which choices are stochastic best responses to iterations of thinking which are increasingly noisy. One parameter expresses the increase in noise across iterations (when there is no increase the model reduces to QRE), and the other expresses the overall level of stochastic response.

Our approach attempts to broaden the scope of application of these ideas to many games, while simultaneously adding precision to Nagel's scheme and economizing on the many parameters used by Stahl and Wilson. The idea is to see how far one can get with a distribution of types that is characterized by only one parameter (τ), by best-response (eliminating the need for response sensitivity parameters used in Stahl and Wilson), and by sharp restrictions on what the various types do.

3 Theoretical properties of the CH model

The combination of optimizing decision rules and the one-parameter Poisson structure makes the CH model relatively easy to work with theoretically. This section illustrates the model's theoretical properties in several classes of games.

3.1 Dominance-solvable games

As noted earlier, when τ is large, the relative proportions of adjacent types, which is $f(k-1)/f(k-2) = \tau/(k-1)$, puts overwhelming weight on the higher-step types. Iterating, this means that when τ is large, a k -step thinker acts as if almost all others are using $k-1$ steps. One-step thinkers will never violate dominance. Two-step thinkers will never choose strategies which are dominated when dominated strategies are deleted (since they think they are playing one-steppers who don't violate dominance.) The same logic can be iterated indefinitely when τ is large (i.e., for any finite number of iterations of deletion, a large enough value of τ exists which yields decision rules that correspond to that amount of deletion). So when τ is sufficiently large, the CH model converges to Nash equilibrium in games that are solved by repeated deletion of weakly dominated strategies.

This relation ties the CH idea closely to Nash equilibrium in dominance-solvable games: If you believe players will choose equilibrium strategies in dominance-solvable games, then you must also believe the CH model with large τ is an equally-good model of behavior in those games. The relation between the CH and equilibrium approaches also highlights where the Nash approach is likely to go wrong. Since large values of τ are needed to reach dominance-solvable equilibrium in games that are only solved by

deletion of very many (iteratively) dominated strategies, if thinking is limited then only partial movement toward equilibrium will occur.

The CH model makes an interesting prediction that Nash equilibrium does not make. In p -beauty contest games two or more players all choose numbers in some interval (say $[0,100]$) and the player whose number is closest to $p < 1$ times the average in absolute value wins a fixed prize (see Nagel, 1995; Ho, Camerer and Weigelt, 1998; Nagel, 1999). The game is dominance-solvable and the unique Nash equilibrium is zero (the number which is equal to p times itself).

There is an interesting behavioral effect on group size however. In three-person games with $p = 2/3$, players tend to choose higher numbers than in 2-person games (see Grosskopf and Nagel, 2001, and below). The 2-person game is special because it can be solved by weak dominance. In the 2-person game, one player will always be high and one low, and for any $p < 1$, p times the average will be closer to the lower player's number. Therefore, rational players want to choose the lowest number possible- 0. In fact, in the CH model all players using one or more thinking steps will choose zero. This is not true in the 3-player game; a smart player wants to choose a number between the other two numbers if they are sufficiently far apart.

3.2 Market entry games

In the market entry games we studied experimentally in section IV below, N entrants simultaneously decide whether to enter (1) or not enter (0) a market. Denote capacity by c (expressed as a fraction of number of potential entrants). If c or fewer players enter, the entrants all earn a payoff of 1; if more than c enter, the entrants earn zero. Not entering yields a payoff of 0.5. For theoretical simplicity, assume there are infinitely many atomistic entrants. (In our empirical estimation we drop this assumption.) If entrants are atomistic and risk-neutral, they only care about whether the fraction of others entering is above c or not (if not, they enter; if so, they stay out). Denote the entry function of step k players for capacity c by $e(k, c) : c \rightarrow [0, 1]$.

We are interested in the conditions under which actual entry is monotonic in c . Denote the normalized cumulative entry function for all steps up to and including k by $E(k, c) : c \rightarrow [0, 1]$.

We have:

$$e(0, c) = \frac{1}{2}, \quad \forall c$$

$$E(k, c) = \frac{\sum_{j=0}^k f(j) \cdot e(j, c)}{\sum_{j=0}^k f(j)} = \frac{\sum_{j=0}^k f(j) \cdot e(j, c)}{F(j)}, \text{ where } F(j) \equiv \sum_{j=0}^k f(j)$$

In general, for $k \geq 1$

$$e(k, c) = \begin{cases} 0 & \text{if } E(k-1, c) > c \\ 1 & \text{if } E(k-1, c) < c \end{cases}$$

In general, $E(k, c)$ is a step function with the following cutpoint values (at which steps begin or end) with increasing c for $c < 1/2$

$$\frac{\frac{1}{2}f(0)}{F(k)}, \frac{\frac{1}{2}f(0)+f(k)}{F(k)}, \frac{\frac{1}{2}f(0)+f(k-1)}{F(k)}, \frac{\frac{1}{2}f(0)+f(k-1)+f(k)}{F(k)}, \dots, \frac{\frac{1}{2}f(0)+f(2)+\dots+f(k)}{F(k)}$$

The cutpoint values for $c > 1/2$ are

$$\frac{\frac{1}{2}f(0)+f(1)}{F(k)}, \frac{\frac{1}{2}f(0)+f(1)+f(2)}{F(k)}, \dots, \frac{\frac{1}{2}f(0)+f(1)+f(2)+\dots+f(k)}{F(k)}$$

(For $c = 1/2$ atomistic entrants are all indifferent and randomize so $E(k, .5) = .5 \forall k$.)

These cutpoints imply two properties: The cutpoints are always (weakly) monotonically increasing in c for the $c < 1/2$ segment as long as $f(k-1) > f(k)$, $\forall k \geq 2$. For a Poisson $f(k)$, this is equivalent to $\tau \leq 2$. Furthermore, the last cutpoint for the $c < 1/2$ segment is greater than the first cutpoint of the $c > 1/2$ segment iff $\frac{1}{2}f(0) + f(2) + f(3) + \dots + f(k-1) + f(k) \leq \frac{1}{2}f(0) + f(1)$. This is equivalent to $f(1) \geq f(2) + f(3) + \dots + f(k)$, which implies $f(1) \geq 1 - f(0) - f(1)$. For Poisson this implies $(1 + 2\tau) \geq e^\tau$ or $\tau \leq 1.25$. Thus, $\tau \leq 1.25$ implies weak monotonicity throughout both the left ($c < 1/2$) and right ($c > 1/2$) segments of the entry function $E(k, c)$ (since $\tau < 1.25$ satisfies the $\tau < 2$ condition *and* ensures monotonicity across the crossover from the left to right halves of $e(k, c)$).

Figure 1 shows $e(0, c)$, $e(1, c)$, $e(2, c)$, and $E(2, c)$ for $\tau = 1.6$. Note that the higher level types “smoothen” the cumulative entry function. As we will discuss further below, in one-shot entry games like these, the entry rate *is* usually monotonic in capacity c .

But how? Daniel Kahneman (1988) wrote that “to a psychologist, it is like magic”. There are many pure-strategy equilibria in which c entrants enter, but how do subjects who play once without talking coordinate on one equilibrium? There is a symmetric mixed-strategy equilibrium, but it cannot account for the regular empirical fact that too many players enter at low values of c and too few enter at high values of c . The proof above shows that the CH model, with $\tau \leq 1.25$, can explain how monotonic entry rates arise from a simple cognitive process. As we will see below, the model also explains the deviations from Nash behavior at low and high capacities.

3.3 Nash demand games

An interesting property of the CH model is that it can produce behavior which corresponds to fair or focal outcomes in bargaining games. The easiest example is the Nash demand game. Two players divide a one unit prize by demanding x_1, x_2 simultaneously. They earn what they demanded iff $x_1 + x_2 \leq 1$. In the CH model, zero-step players randomize over $[0,1]$. A one step player who demands x expects to earn $x(1 - x)$, which is maximized by demanding half ($x = .5$). Higher-step players also demand half since that is a best response to any mixture of random and .5 demands. The model therefore predicts that $1 - f(0)$ players will demand half (around 80% if $\tau = 1.5$) and other demands will be sprinkled throughout the $[0,1]$ interval.

When one player has an outside option, the CH model approximates the “split the difference” equilibrium in which players demand half the surplus beyond the option.¹³ Binmore et al. (1985) found that in early periods of their experiment many demands were consistent with the split-the-difference solution, though after learning over rounds of bargaining most players converged to the perfect equilibrium demands of about a half.

3.4 Stag hunt games

The CH model also produces an interesting effect of group size in stag hunt games. Imagine a stag hunt game in which each of n players choose either H or L. Players earn

¹³Suppose level 0 has an option of y and randomizes over demands in the interval $[y,1]$. Now the one-step player’s demand of x is accepted with probability $\max(0, (1 - x - y)/(1 - y))$. The expected payoff is $x(1 - x - y)/(1 - y)$ which is maximized at $(1 - y)/2$, dividing the surplus.

1 if they choose H and everyone else does, 0 if they choose H and anybody else chooses L, and x if they choose L (regardless of what others do).

In the two-player game, 0-step thinkers randomize so 1-step thinkers (and all higher-step thinkers) choose H if $x \leq 1/2$ and choose L if $x \geq 1/2$ (the higher-step behavior corresponds to the risk-dominance refinement). In the three-player game, however, a 1-step player thinks she is facing two 0-step players who randomize independently; so the chance of at least one L is .75. As a result, the 1-step player (and higher-level players) choose H iff $x \leq .25$. Thus, for values $.25 \leq x \leq .5$, there will be mostly H play in 2-player games and mostly L-play in 3-player games. This is a simple way of expressing the idea that there is more strategic uncertainty in games with more players, and corresponds to the empirical fact that choices are lower in stag hunt (or ‘weak-link’) game experiments as the number of players rises (e.g., Camerer, 2003, chapter 7).

4 Estimation and model comparison

This section estimates best-fitting values of τ in the Poisson CH model and compares it to other models. Our philosophy is that exploring a wide range of games and models is especially useful in the early stage of a research program. Models which sound appealing (perhaps because they are conventional) may fit badly. Fitting a wide range of games often turns up clues about where models fail and how to improve them.

Since the cognitive hierarchy model is designed to be general, it is particularly important to check its robustness across different types of games and see how regular the best-fitting values of τ are. Once the mean number of thinking steps τ is specified, the model’s predictions about the distribution of choices can be easily derived. We then use maximum likelihood (MLE) techniques to estimate best-fitting values of τ and their precision. The MLE procedure can be shown to estimate τ reliably with samples of 50 or so.¹⁴

¹⁴Simulations were used to see how well the MLE procedure recovers the true value of τ when the CH model actually creates the data. In the simulations, pseudosamples of size N (where $N = 20, 48, 100$) were created using the CH value with a known τ , for three games from the mixed-strategy game sample (one $K \times K$ game for $K \in 2, 4, 6$). For each pseudosample, the value of τ was then estimated using MLE to see whether the procedure can recover the correct (known) value of τ which actually created the data. There is no systematic bias in recovered τ , and the precision of recovered estimates is reasonably good—bootstrapped confidence intervals are around $(\tau - .3, \tau + .3)$ —except when samples are small ($N = 20$)

We fit five data sets: 33 matrix games with 2-4 strategies from three data sets, 22 games with mixed equilibria (new data) and the binary entry game described above (new data). The matrix games are 12 games from Stahl and Wilson (1995), 8 games from Cooper and Van Huyck (2001) (used to compare normal- and extensive-form play), and 13 games from Costa-Gomes, Crawford and Broseta (2001). All these games were played only once with feedback, with sample sizes large enough to permit reliable estimation. The 22 games with mixed-equilibria are taken from those reviewed by Camerer (2003, chapter 3), with payoffs rescaled so subjects win or lose about \$1 in each game.¹⁵ These games were run in four experimental sessions of 12 subjects each, using the “playing in the dark” software developed by McKelvey and Palfrey. Two sessions used undergraduates from Caltech and two used undergraduates from Pasadena City College (PCC), which is near Caltech.

The binary entry game is the one described above. In the four experimental sessions, each of 12 players simultaneously decides whether to enter a market with announced capacity c . If c or fewer players enter the entrants earn \$1; if more than c enter they earn nothing. Not entering earns \$.5. In this simple structure, risk-neutral players care only about whether the expected number of entrants will be less than $c - 1$.¹⁶ Subjects were shown five capacities $c = 2, 4, 6, 8, 10$ in a fixed random order, with no feedback.

The estimation aims to answer two questions: Is the estimated value of τ reasonably regular across games with very different structures? How does the CH Poisson specification compare to Nash equilibrium?

in 2×2 games. For detailed output, see Camerer, Ho and Chong, 2002b.

¹⁵The 22 mixed games are (in order of presentation to the subjects): Ochs (1995), (matching pennies plus games 1-3); Bloomfield (1994); Binmore et al. (2001) Game 4; Rapoport and Amaldoss (2000); Binmore et al (2001), games 1-3; Tang (2001), games 1-3; Goeree, Holt, and Palfrey (2000), games 2-3; Mookerjee and Sopher (1997), games 1-2; Rapoport and Boebel (1992); Messick (1965); Lieberman (1962); O’Neill (1987); Goeree, Holt, and Palfrey (2000), game 1. Four games were perturbed from the original payoffs: The row upper left payoff in Ochs’s original game 1 was changed to 2; the Rapoport and Amaldoss (2000) game was computed for $r=15$; the middle row payoff in Binmore et al (2001) game 2 was 30 rather than -30; and the lower left row payoff in Goeree, Holt and Palfrey’s (2000) game 3 was 16 rather than 37. Original payoffs in games were multiplied by the following conversion factors: 10, 10, 10, 10, 0.5, 10, 5, 10, 10, 10, 1, 1, 1, 0.25, 0.1, 30, 30, 30, 5, 3, 10, 0.25. Currency units were then equal to \$.10.

¹⁶This structure suppresses the effect of overconfidence actual business entrants might have in a game in which more skilled entrants earn more (e.g., Camerer and Lovallo, 1999).

4.1 How regular is τ ?

Table 1 shows game-by-game estimates of τ in the Poisson CH model, and estimates when τ is constrained to be common across games within each data set. Five of 60 game-specific τ estimates are high (4 or more) and some are zero. The interquartile range across 60 estimates is (.98,2.21) and the median is 1.55. An Appendix Table shows bootstrapped 95% confidence intervals for the τ estimates. Most of the intervals have a range of about one, which means τ is estimated fairly precisely. The common τ estimates are roughly 1-2. This highly regular τ suggests that the CH model can be used to reliably predict behaviors in new games.

4.2 Which models fit best?

Table 2 shows log likelihoods (LL) and mean-squared deviations for several model estimated game-by-game or with common parameters across games in a dataset.¹⁷ This table answers several questions. Focusing first on the CH Poisson model, moving from game-specific estimates of τ to common within-column estimates only degrades fit badly in the Stahl-Wilson data; in the other samples imposing a common τ fits about as well as letting τ vary in each game.

The CH Poisson model also fits substantially better than Nash. This result suggests that relaxing mutual consistency is a fruitful approach to building a descriptive theory of games. A graphical comparison of predicted and actual strategy frequencies helps give a clearer image of how accurate the CH and Nash approaches are. Each point in Figures 2a-b represents a distinct strategy in each of the 33 matrix games (Figure 2a) and 22 mixed games (Figure 2b), comparing actual strategy frequencies which CH predictions using a single common τ within each data set (i.e., one τ per figure). The R^2 's are both around .80 for the CH model. Figures 3a-b show the corresponding figures comparing actual frequencies and the Nash predictions. In Figure 3a there are many strategies which are predicted to be always chosen (probability one) or never chosen (probability zero), so the fit is not visually impressive and the R^2 is modest (.32). Figure 3b is a fairer test because most of the Nash predictions are in the interior, but there is still wide dispersion and R^2 rises to about a half. Comparing Figures 2a-b with 3a-b shows that the CH

¹⁷When the Stahl-Wilson games 2, 6, 8 are included the common τ is zero because these games swamp the other 10. We therefore excluded these games in estimating the common τ .

model is able to tighten up the fit dramatically for matrix games, and substantially for mixed games, using only one parameter in each data set.

4.3 Predicting across games

Good theories should predict behavior in new situations. A simple way to see whether the CH model can do this, within a large sample of games, is to estimate the value of τ on $n-1$ games and forecast behavior in each holdout game separately. (This is a roundabout way to test how stable τ appears to be across games, and also whether small variations in estimated τ create large or small differences in predicted choice frequencies.) The bottom panel of Table 2 reports the result of this sort of cross-game estimation. The CH Poisson model fits cross-game a little less accurately than when estimates are common within games. This suggests that the CH model can be used to predict behaviors in new games.

5 Economic value of theories

One way to use theories of strategic thinking is to give advice to players. Camerer and Ho (2001) introduced the idea of judging theories by their *economic value*. Economic value is computed by using a theory to predict what other players will do, choosing a best response based on that prediction, and comparing whether the best response actually would have earned more money than the response a subject actually chose.

Economic value is also an indirect way to measure how well behavior is equilibrated. If players are mutually consistent, then their beliefs already match likely choices so no theory can have economic value. Therefore, if Nash equilibrium is predictively accurate, then it cannot have economic value. Similarly, if players are in equilibrium then models which assume they are not in equilibrium (such as the CH model) will have negative economic value.

Table 3 reports the profits players would have earned if they used the CH model to forecast likely behavior and chose best responses. The economic value of a theory is the difference between these hypothetical profits and the actual profits players earned. (The

payoffs from predicting perfectly, using the actual distribution of strategies chosen by others, are also reported because these represent an upper bound on economic value).

The top panel shows economic value when common parameters are estimated within each set of games. The CH approach adds value in all data sets, typically 30-50% of the maximum possible economic value. Nash equilibrium adds a little less value, and subtracts value in two data sets. The bottom panel shows economic value when parameters are estimated on $n - 1$ data sets and used to forecast the remaining data set. The results are basically the same.

6 Economic implications of limited strategic thinking

Models of iterated thinking can be applied to several interesting problems in economics, including asset pricing, speculation, competition neglect in business entry, incentive contracts, and macroeconomics.

Asset pricing: As Keynes pointed out (and many commentators since him; e.g., Tirole 1985; Shleifer and Vishny, 1990), if investors in stocks are not sure that others are rational (or will price assets rationally in the future) then asset prices will not necessarily equal fundamental or intrinsic values.¹⁸ A precise model of limited strategic thinking might therefore be used to explain the existence and crashes of price bubbles.

Speculation: The “Groucho Marx theorem” says that traders who are risk-averse should not speculate by trading with each other even if they have private information (since the only person who will trade with you may be better-informed). But this theorem rests on unrealistic assumptions of common knowledge of rationality and is violated constantly by massive speculative trading volume and other kinds of betting, as well as in experiments.¹⁹ Speculation will occur in CH models because 1- and higher-step players think they are sometimes betting against random (0-step) bettors who make mistakes.

¹⁸Besides historical examples like Dutch tulip bulbs and the \$5 trillion tech-stock bubble in the 1990s, experiments have shown such bubbles even in environments in which the asset’s fundamental value is controlled and commonly-known. See Smith, Suchanek and Williams, 1988; Camerer and Weigelt, 1993; and Lei, Noussair and Plott, 2001.

¹⁹See Sonsino, Erev and Gilat, 2000; Sovik, 2000.

Competition neglect and business entry: Players who do limited iterated thinking, or believe others are not as smart as themselves, will neglect competition in business entry, which may help explain why the failure rate of new businesses is so high (see Camerer and Lovo, 1999; Huberman and Rubinstein, 2000). Simple entry games are studied below. Theory and estimates from experimental data show that the CH model can explain why the amount of entry is monotonic in market capacity, but too many players enter when capacity is low. Managerial hubris, overconfidence, and self-serving biases which are correlated with costly delay and labor strikes in the lab (Babcock et al., 1995) and in the field (Babcock and Loewenstein, 1997) can also be interpreted as players not believing others always behave rationally.

Incentives: In his review of empirical evidence on incentive contracts in organizations, Prendergast (1999) notes that workers typically react to simple incentives as standard models predict. However, firms usually do not implement complex contracts which *should* elicit higher effort and improve efficiency. This might be explained as the result of firms thinking strategically, but not believing that workers will respond rationally.

Macroeconomics: Woodford (2001) notes that in Phelps-Lucas “islands” models, nominal shocks can have real effects, but their predicted persistence is too short compared to actual effects in data. He shows that imperfect information about *higher-order* nominal GDP estimates—beliefs about beliefs, and higher-order iterations—can cause longer persistence which matches the data, and Svensson (2001) notes that iterated beliefs are probably constrained by computational capacity. In CH models, players’ beliefs are not mutually consistent so there is higher-order belief inconsistency which might explain the longer persistence of shocks that Woodford noted.

7 Conclusion

This paper introduced a parsimonious one-parameter cognitive hierarchy (CH) model of limited reasoning in games. The model is designed to be as general and precise as Nash equilibrium (in fact, it refines implausible Nash equilibria and selects one of multiple Nash equilibria). One innovation is to use axioms and estimation to restrict the frequencies of players who stop thinking at various levels. The idea that most players do *some* strategic thinking, but the amount of strategic thinking is sharply constrained by working memory, is consistent with a simple axiom which implies a Poisson distribution of thinking steps

that can be characterized by one parameter τ (the mean number of thinking steps, and the variance). Plausible restrictions and estimates from many experimental data sets suggest that the mean amount of thinking τ is between one and two. The value $\tau = 1.5$ is a good omnibus guess which makes the CH theory parameter-free.

The other innovation in this paper is to show that the same model can explain limited equilibration in dominance-solvable games (like p -beauty contests) and also to explain why behavior in one-shot games with mixed equilibria is surprisingly well-approximated by Nash equilibrium. A useful example is simultaneous binary entry games in which players choose whether to enter a capacity-constrained market. In one-shot games with no communication, the rate of entry in these games is ‘magically’ monotonic in the capacity c , but there is reliable over-entry at low values of c and under-entry at high values of c . The CH approach predicts monotonicity (it is guaranteed when $\tau \leq 1.25$) and also explains over- and under-entry. Furthermore, in the CH approach most players use a pure strategy, which creates a kind of endogenous purification that can explain how a population mixture of players who use pure strategies (and perhaps regard mixing as nonsensical) can approximate a mixed equilibrium.

Because players do not appear to be mutually consistent in one-shot games where there is no opportunity to learn, it is possible that a theory of how others are likely to play has economic value— i.e, players would earn more if they used the model to recommend choices, compared to how much they actually earn. In fact, economic value is always positive for the CH model, whether τ is estimated within a data set or across data sets. (Economic value is about 1/3 to 1/2 of the maximum possible economic value.) The Nash approach adds less economic value, and sometimes subtract economic value (e.g., in p -beauty contests players are better choosing on their own than picking the Nash recommendation).

There are many challenges in future research. An obvious one is to endogenize the mean number of thinking steps τ , presumably from some kind of cost-benefit analysis in which players weigh the marginal benefits of thinking further against cognitive constraint (cf. Gabaix and Laibson, 2000). It is also likely that a more nuanced model of what 0-step players are doing would improve model fits in some types of games.

Since CH model makes a prediction about the kinds of algorithms that players use in thinking about games, cognitive data other than choices can - like belief-prompting, response times, information lookups, or even brain imaging - can be used to test the

model. It is an interesting direction to pursue.

The model is easily adapted to incomplete information games because the 0-step players make choices which reach every information set, which eliminates the need to impose delicate refinements to make predictions. Explaining behavior in signaling games and other extensive-form games with incomplete information is therefore workable and a high priority for future work. (Brandts and Holt, 1992, and Banks, Camerer, and Porter, 1994, suggest that mixtures of decision rules in the first period, and learning in subsequent periods, can explain the path of equilibration in signaling games; the CH approach may add some bite to these ideas.)

Another important challenge is repeated games. The CH approach will generally underestimate the amount of strategic foresight observed in these games (e.g., players using more than one step of thinking will choose supergame strategies which always defect in repeated prisoners' dilemmas). An important step is to draw a sensible parametric analogy between steps of strategic foresight and steps of iterated thinking is necessary to explain observed behavior in such games (cf. Camerer, Ho and Chong, 2002a; Camerer et al., 2002).

Finally, the ultimate goal of the laboratory honing of simple models is to explain behavior in the economy. Field phenomena which seem to involve limits on iterated thinking include speculation in zero-sum betting games, price bubbles in asset markets, contract structure and behavior, and macroeconomic applications involving limits on iterated expectations.

While we have endogenized the mean number of thinking steps, this convergence property is a clue about why players will do only a few steps of thinking— If the number of steps of thinking is endogenized by some kind of comparison of benefits (marginal expected payoffs from thinking more) and cognitive costs, the fact that the expected payoffs of higher-step thinkers converges will lead to a natural truncation which limits the amount of thinking.

References

- [1] Babcock, Linda, and George Loewenstein, “Explaining Bargaining Impasses: The Role of Self-serving Biases,” *Journal of Economic Perspectives*, 11 (1997), 109–26.
- [2] Babcock, Linda, George Loewenstein, Samuel Issacharoff, and Colin F. Camerer, “Biased Judgments of Fairness in Bargaining,” *American Economic Review*, 85 (1995), 1337-43.
- [3] Banks, Jeffrey, Colin F. Camerer, and David Porter, “Experimental Tests of Nash Refinements in Signaling Games,” *Games and Economic Behavior*, 6 (1994), 1-31.
- [4] Bernheim, Douglas, “Rationalizable Strategic Behavior,” *Econometrica*, 52 (1984), 1007-1028.
- [5] Binmore, Kenneth, “Modeling rational players: Part II,” *Economics and Philosophy*, 4 (1988), 9-55.
- [6] Binmore, Kenneth, Avner Shaked, and John Sutton, “Testing Noncooperative Bargaining Theory: A Preliminary Study,” *American Economic Review*, 75 (1985), 1178-80.
- [7] Binmore, Kenneth, Joe Swierzbinski and Chris Proulx, “Does Maximin Work? An Experimental Study,” *Economic Journal*, 111 (2001), 445-464.
- [8] Bloomfield, Robert, “Learning a Mixed Strategy Equilibrium in the Laboratory,” *Journal of Economic Behavior and Organization*, 25 (1994), 411-436.
- [9] Brandts, Jordi and Charles Holt, “An Experimental Test of Equilibrium Dominance in Signaling Games,” *American Economic Review*, 82 (1992), 1350-65.
- [10] Brown, George, “Iterative Solution of Games by Fictitious Play,” in *Activity Analysis of Production and Allocation*, New York: Wiley, 1951.
- [11] Cabrera, Susana, C. Mónica Capra, and Rosario Gomez, “The Effects of Common Advice on One-shot Traveler’s Dilemma Games: Explaining Behavior Through an Introspective Model with Errors,” Washington and Lee University working paper, 2002.
- [12] Camerer, Colin F., “Behavioral Game Theory,” in R. Hogarth (Ed.), *Insights in Decision Making: A tribute to Hillel J. Einhorn*. University of Chicago Press, 1990.

- [13] Camerer, Colin F., *Behavioral Game Theory: Experiments on Strategic Interaction*, Princeton:Princeton University Press, 2003.
- [14] Camerer, Colin F. and Ho Teck-Hua, “Strategic Learning and Teaching,” in S.J. Hoch and H.C. Kunreuther (Eds.), *Wharton on Making Decisions*, New York: John Wiley, 2001, 159-176.
- [15] Camerer, Colin F., Teck-Hua Ho and Juin-Kuan Chong, “Sophisticated EWA Learning and Strategic Teaching in Repeated Games,” *Journal of Economic Theory*, 104 (2002), 137-188.
- [16] Camerer, Colin F., Teck-Hua Ho and Juin-Kuan Chong, “A Cognitive Hierarchy Theory of One-shot Games: Some Preliminary Results,” CalTech working paper, 2002.
- [17] Camerer, Colin F., Teck-Hua Ho, Juin-Kuan Chong and Keith Weigelt, “Strategic Teaching and Equilibrium Models of Repeated Trust and Entry Games,” CalTech Working Paper, 2002. <http://www.hss.caltech.edu/camerer/camerer.html>
- [18] Camerer, Colin F., Eric J. Johnson, Talia Rymon, and Sankar Sen, “Cognition and Framing in Sequential Bargaining for Gains and Losses,” in K. Binmore, A. Kirman, and P. Tani (Eds.), *Frontiers of Game Theory*, Cambridge: MIT Press, 1994, 27-47.
- [19] Camerer, Colin F., George Loewenstein, and Drazen Prelec, “Gray Matters: How Neuroscience can Inform Economics,” under revision for *Journal of Economic Perspectives*, 2002.
- [20] Camerer, Colin F., and Daniel Lovallo, “Overconfidence and Excess Entry: An Experimental Approach,” *American Economic Review*, 89 (1999), 306-18.
- [21] Camerer, Colin F. and Keith Weigelt, “Convergence in Experimental Double Auctions for Stochastically Lived Assets,” in D. Friedman and J. Rust (Eds.), *The Double Auction Market: Theories, Institutions and Experimental Evaluations*, Redwood City, CA: Addison–Wesley, 1993, 355–396.
- [22] Capra, C. Mónica, *Noisy Expectation Formation in One-shot Games*, Ph.D. thesis, 1999.

- [23] Chen, Hsiao-Chi, James W. Friedman, and Jacques-Francois Thisse, "Boundedly Rational Nash Equilibrium: A Probabilistic Choice Approach," *Games and Economic Behavior*, 18 (1996), 1832-54.
- [24] Cooper, David and J. Van Huyck, "Evidence on the Equivalence of the Strategic and Extensive Form Representation of Games," Texas A&M University Department of Economics, 2001. http://econlab10.tamu.edu/JVH_gtee/
- [25] Costa-Gomes, M. and V. Crawford and B. Broseta, "Cognition and Behavior in Normal-form Games: An Experimental Study," *Econometrica*, 69 (2001), 1193-1235.
- [26] Crawford, V., "Theory and Experiment in the Analysis of Strategic Interactions," in D. Kreps and K. Wallis (Eds.), *Advances in Economics and Econometrics: Theory and Applications*, Seventh World Congress, Volume I. Cambridge: Cambridge University Press, 1997.
- [27] Fudenberg, D. and D. Kreps, "A Theory of Learning, Experimentation and Equilibrium in Games," Stanford University working paper, 1990.
- [28] Fudenberg, Drew and David Levine, *The Theory of Learning in Games*. Cambridge: MIT Press, 1998.
- [29] Gabaix, Xavier and David Laibson, "Bounded Rationality and Directed Cognition," Harvard University working paper, 2000.
- [30] Goeree, Jacob K. and Charles A. Holt, "Stochastic Game Theory: For Playing Games, Not Just for Doing Theory," *Proceedings of the National Academy of Sciences*, 96 (1999), 10564-10567.
- [31] Goeree, Jacob K. and Charles A. Holt, "Stochastic Game: A Theory of Noisy Introspection," *Games and Economic Behavior*, (2002), forthcoming.
- [32] Goeree, Jacob; Holt, Charles A. and Thomas R. Palfrey, "Risk Aversion in Games with Mixed Strategies," University of Virginia Department of Economics, May 2000.
- [33] Grosskopf, Brit and Rosemarie Nagel, "Rational Reasoning or Adaptive Behavior? Evidence from Two-Person Beauty Contest Games," Harvard NOM Research Paper No. 01-09, 2001.

- [34] Harsanyi, John, "The Tracing Procedure: A Bayesian Approach to Defining a Solution for n-person Noncooperative Games," *International Journal of Game Theory*, 4 (1975), 61-94.
- [35] Harsanyi, John C., and Reinhard Selten, *A General Theory of Equilibrium Selection in Games*. Cambridge, MA: MIT Press, 1988.
- [36] Haruvy, Ernan and Dale O. Stahl, "An Empirical Model of Equilibrium Selection in Symmetric Normal-form Games," University of Texas Department of Economics Working Paper, 1998.
- [37] Ho, Teck-Hua, Colin Camerer, and Keith Weigelt, "Iterated Dominance and Iterated Best-response in p -Beauty Contests," *American Economic Review*, 88 (1998), 947-969.
- [38] G. Huberman and A. Rubinstein, "Correct Belief, Wrong Action and a Puzzling Gender Difference," working paper, 2000.
- [39] Kadane, Jay and Pat Larkey, "Subjective Probability and the Theory of Games, and Reply," *Management Science*, 28 (1982), 113-120 & 124.
- [40] D. Kahneman, "Experimental Economics: A Psychological Perspective," in R. Tietz, W. Albers, and R. Selten (Eds.) *Bounded Rational Behavior in Experimental Games and Markets*, New York: Springer-Verlag, 1988, 11-18.
- [41] Krishna, Vijay and Sjöström, Tomas, "On the Convergence of Fictitious Play," *Mathematics of Operations Research*, 23 (1998), 479-511.
- [42] Lei, Vivian, Charles Noussair and Charles Plott, "Nonspeculative Bubbles in Experimental Asset Markets: Lack of Common Knowledge of Rationality vs. Actual Irrationality," *Econometrica*, 69 (2001), 831-859.
- [43] Lieberman, Bernhardt, "Experimental Studies of Conflict in Some Two-person and Three-person Games," in J. H. Criswell, H. Solomon, and P. Suppes (Eds.), *Mathematical Models in Small Group Processes*. Stanford: Stanford University Press, 1962, 203-220.
- [44] Lucas, R.G., "Adaptive Behavior and Economic Theory," *Journal of Business*, 59 (1986), S401-S426.

- [45] McKelvey, Richard D. and Thomas R. Palfrey, "Quantal Response Equilibria for Normal-form Games," *Games and Economic Behavior*, 7 (1995), 6-38.
- [46] McKelvey, Richard D. and Thomas R. Palfrey, "Quantal Response Equilibria for Extensive Form Games," *Experimental Economics*, 1 (1998), 9-41.
- [47] Messick, David M., "Interdependent Decision Strategies in Zero-sum Games: A Computer-controlled Study," *Behavioral Science*, 12 (1967), 33-48.
- [48] Miyasawa, "On the Convergence of Learning Processes in a 2×2 Non-zero Two-person Game," Research memo 33, Princeton University, 1961.
- [49] Mookerjee, Dilip and Barry Sopher, "Learning and Decision Costs in Experimental Constant-sum Games," *Games and Economic Behavior*, 19 (1997), 97-132.
- [50] Nachbar, John, "Evolutionary Selection Dynamics in Games: Convergence and Limit Properties," *International Journal of Game Theory*, 19 (1990), 59-89.
- [51] Nagel, R., "Unraveling in Guessing Games: An Experimental Study," *The American Economic Review*, 85 (1995), 1313-1326.
- [52] Nagel, R., "A Survey of Research on Experimental Beauty-contest Games," in D. Budescu, I. Erev, and R. Zwick (Eds.), *Games and Human Behavior: Essays in Honor of Amnon Rapoport*, Mahwah, NJ: Lawrence Erlbaum, 1999.
- [53] Nagel, Rosemarie, Antoni Bosch-Domnech, Albert Satorra, and Jos Garca-Montalvo, "One, Two, (Three), Infinity: Newspaper and Lab Beauty-contest Experiments," *American Economic Review*, forthcoming, 2002. <http://www.econ.upf.es/leex/working.htm>.
- [54] Ochs, Jack, "Games with Unique, Mixed Strategy Equilibria: An Experimental Study," *Games and Economic Behavior*, 10 (1995), 202-217.
- [55] O'Neill, Barry, "Nonmetric Test of the Minimax Theory of Two-person Zero-sum Games," *Proceedings of the National Academy of Sciences*, 84 (1987), 2106-2109.
- [56] Pearce, David G., "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, 52 (1984), 1029-1050.
- [57] Prendergast, Canice, "The Provision of Incentives in Firms," *Journal of Economic Literature*, 37 (1999), 7-63.

- [58] Rapoport, Amnon and Amaldoss, Wilfred, "Mixed Strategies and Iterative Elimination of Strongly Dominated Strategies: An Experimental Investigation of States of Knowledge," *Journal of Economic Behavior and Organization*, 42 (2000), 483-521.
- [59] Rapoport, Amnon and Boebel, Richard B., "Mixed Strategies in Strictly Competitive Games: A Further Test of the Minimax Hypothesis," *Games and Economic Behavior*, 4 (1992), 261-283.
- [60] Robinson, Julia, "An Iterative Method of Solving a Game," *Annals of Mathematics*, 54(2), (1951), 296-301.
- [61] Rosenthal, Robert W., "A bounded rationality approach to the study of noncooperative games," *International Journal of Game Theory*, 18(3), (1989), 273-92.
- [62] Selten, Reinhard, "Features of Experimentally Observed Bounded Rationality," *European Economic Review*, 42 (1998), 413-436.
- [63] Shapley, Lloyd S., "Some Topics in Two-person Games," in M. Dresher, L. S. Shapley, and A. W. Tucker (Eds.), *Advances in Game Theory*, Princeton, NJ: Princeton University Press, 1964, 1-28.
- [64] Shleifer, Andrei and Robert W. Vishny, "Equilibrium Short Horizons of Investors and Firms," *American Economic Review Papers and Proceedings*, 80 (May 1990), 148-153.
- [65] Smith, Vernon L., Gerry Suchanek and A. Williams, "Bubbles, Crashes and Endogenous Expectations in Experimental Spot Asset Markets," *Econometrica*, 56 (1988), 1119-1151.
- [66] Sonsino, D., I. Erev and S. Gilat, "On the Likelihood of Repeated Zero-sum Betting by Adaptive (Human) Agents," Technion, Israel Institute of Technology, 2000.
- [67] Sovik, Ylva, "Impossible Bets: An Experimental Study," University of Oslo Department of Economics, 2000.
- [68] Stahl, Dale O., "Evolution of $smart_n$ Players," *Games and Economic Behavior*, 5 (1993), 604-17.
- [69] Stahl, D. O., "Is Step- j Thinking an Arbitrary Modeling Restriction or a Fact of Human Nature?" *Journal of Economic Behavior and Organization*, 37 (1998), 33-51.

- [70] Stahl, D. O. and P. Wilson, "On Players Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior*, 10 (1995), 213-54.
- [71] Svensson, L.E.O., "Comments on Michael Woodford paper," presented at Knowledge, Information and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps. Columbia University, October 5-6, 2001.
- [72] Tang, Fang-Fang, "Anticipatory Learning in Two-person Games: Some Experimental Results," *Journal of Economic Behavior and Organization*, 44 (2001), 221-32.
- [73] Tirole, J., "Asset Bubbles and Overlapping Generations," *Econometrica*, 53 (1985), 1071-1100 (reprinted 1499-1528).
- [74] Weibull, *Evolutionary Game Theory*, Cambridge, MA: MIT Press, 1995.
- [75] Woodford, M., "Imperfect common knowledge and the effects of monetary policy," presented at Knowledge, Information and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps. Columbia University, October 5-6, 2001.

Table 1: Parameter Estimate τ for Cognitive Hierarchy Models

Data set	Stahl & Wilson (1995)	Cooper & Van Huyck	Costa-Gomes et al.	Mixed	Entry
<u>Game-specific τ</u>					
Game 1	2.93	16.02	2.16	0.98	0.69
Game 2	0.00	1.04	2.05	1.71	0.83
Game 3	1.35	0.18	2.29	0.86	-
Game 4	2.34	1.22	1.31	3.85	0.73
Game 5	2.01	0.50	1.71	1.08	0.69
Game 6	0.00	0.78	1.52	1.13	
Game 7	5.37	0.98	0.85	3.29	
Game 8	0.00	1.42	1.99	1.84	
Game 9	1.35		1.91	1.06	
Game 10	11.33		2.30	2.26	
Game 11	6.48		1.23	0.87	
Game 12	1.71		0.98	2.06	
Game 13			2.40	1.88	
Game 14				9.07	
Game 15				3.49	
Game 16				2.07	
Game 17				1.14	
Game 18				1.14	
Game 19				1.55	
Game 20				1.95	
Game 21				1.68	
Game 22				3.06	
Median τ	1.86	1.01	1.91	1.77	0.71
<u>Common τ</u>	1.54	0.80	1.69	1.48	0.73

Table 2: Model Fit (Log Likelihood LL and Mean-squared Deviation MSD)

Data set	Stahl & Wilson (1995)	Cooper & Van Huyck	Costa-Gomes et al.	Mixed	Entry
<u>Within-dataset Forecasting</u>					
<u>Cognitive Hierarchy (Game-specific τ)</u> ¹					
LL	-721	-1690	-540	-824	-150
MSD	0.0074	0.0079	0.0034	0.0097	0.0004
<u>Cognitive Hierarchy (Common τ)</u>					
LL	-918	-1743	-560	-872	-150
MSD	0.0327	0.0136	0.0100	0.0179	0.0005
<u>Cross-dataset Forecasting</u>					
<u>Cognitive Hierarchy (Common τ)</u>					
LL	-941	-1929	-599	-884	-153
MSD	0.0425	0.0328	0.0257	0.0216	0.0034
<u>Nash Equilibrium</u> ²					
LL	-3657	-10921	-3684	-1641	-154
MSD	0.0882	0.2040	0.1367	0.0521	0.0049

Note 1: The scale sensitivity parameter λ for the Cognitive Hierarchy models is set to infinity. The results reported in Camerer, Ho & Chong(2001) presented at the Nobel Symposium 2001 are for models where λ is estimated.

Note 2: The Nash Equilibrium result is derived by allowing a non-zero mass of 0.0001 on non-equilibrium strategies.

Table 3: Economic Value for Cognitive Hierarchy and Nash Equilibrium

Data set	Stahl & Wilson (1995)	Cooper & Van Huyck	Costa-Gomes et al.	Mixed	Entry
<u>Total Payoff (% Improvement)</u>					
Actual Subject Choices	384	1169	530	328	118
Ex-post Maximum	685	1322	615	708	176
	79%	13%	16%	116%	49%
<u>Within-dataset Estimation</u>					
Cognitive Hierarchy (Game-specific τ)	401	1277	573	471	128
	4%	9%	8%	43%	8%
Cognitive Hierarchy (Common τ)	418	1277	573	471	128
	9%	9%	8%	43%	8%
<u>Cross-dataset Estimation</u>					
Cognitive Hierarchy (Common τ)	418	1277	573	460	128
	9%	9%	8%	40%	8%
Nash Equilibrium	398	1230	556	274	112
	4%	5%	5%	-16%	-5%

Note 1: The economic value is the total value (in USD) of all rounds that a "hypothetical" subject will earn using the respective model to predict other's behavior and best responds with the strategy that yields the highest expected payoff in each round.

Table A1: 95% Confidence Interval for the Parameter Estimate τ of Cognitive Hierarchy Models

Data set	Stahl & Wilson (1995)		Cooper & Van Huyck		Costa-Gomes et al.		Mixed		Entry	
	Lower	Upper	Lower	Upper	Lower	Upper	Lower	Upper	Lower	Upper
<u>Game-specific τ</u>										
Game 1	2.40	3.65	15.40	16.71	1.58	3.04	0.67	1.22	0.21	1.43
Game 2	0.00	0.00	0.83	1.27	1.44	2.80	0.98	2.37	0.73	0.88
Game 3	0.75	1.73	0.11	0.30	1.66	3.18	0.57	1.37	-	-
Game 4	2.34	2.45	1.01	1.48	0.91	1.84	2.65	4.26	0.56	1.09
Game 5	1.61	2.45	0.36	0.67	1.22	2.30	0.70	1.62	0.26	1.58
Game 6	0.00	0.00	0.64	0.94	0.89	2.26	0.87	1.77		
Game 7	5.20	5.62	0.75	1.23	0.40	1.41	2.45	3.85		
Game 8	0.00	0.00	1.16	1.72	1.48	2.67	1.21	2.09		
Game 9	1.06	1.69			1.28	2.68	0.62	1.64		
Game 10	11.29	11.37			1.67	3.06	1.34	3.58		
Game 11	5.81	7.56			0.75	1.85	0.64	1.23		
Game 12	1.49	2.02			0.55	1.46	1.40	2.35		
Game 13					1.75	3.16	1.64	2.15		
Game 14							6.61	10.84		
Game 15							2.46	5.25		
Game 16							1.45	2.64		
Game 17							0.82	1.52		
Game 18							0.78	1.60		
Game 19							1.00	2.15		
Game 20							1.28	2.59		
Game 21							0.95	2.21		
Game 22							1.70	3.63		
<u>Common τ</u>	1.39	1.67	0.74	0.87	1.53	2.13	1.30	1.78	0.42	1.07

Figure 1: Entry Functions for $\tau = 1.6$




Figure 2a: Predicted Frequencies of Cognitive Hierarchy Models for Matrix Games (common τ)

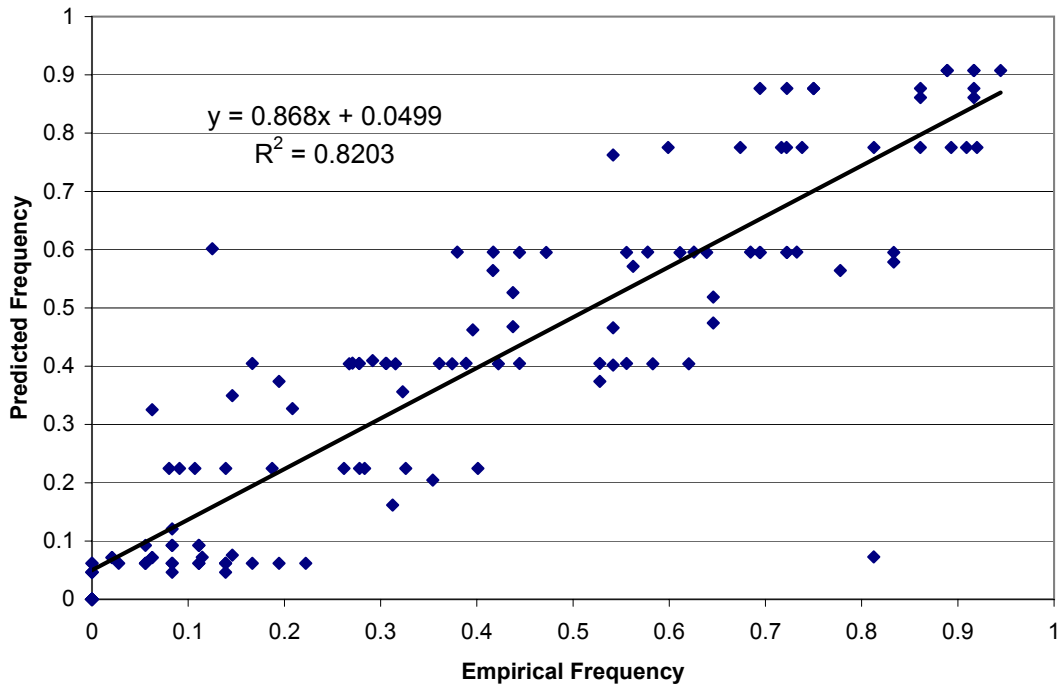


Figure 2b: Predicted Frequencies of Cognitive Hierarchy Models for Entry and Mixed Games (common τ)

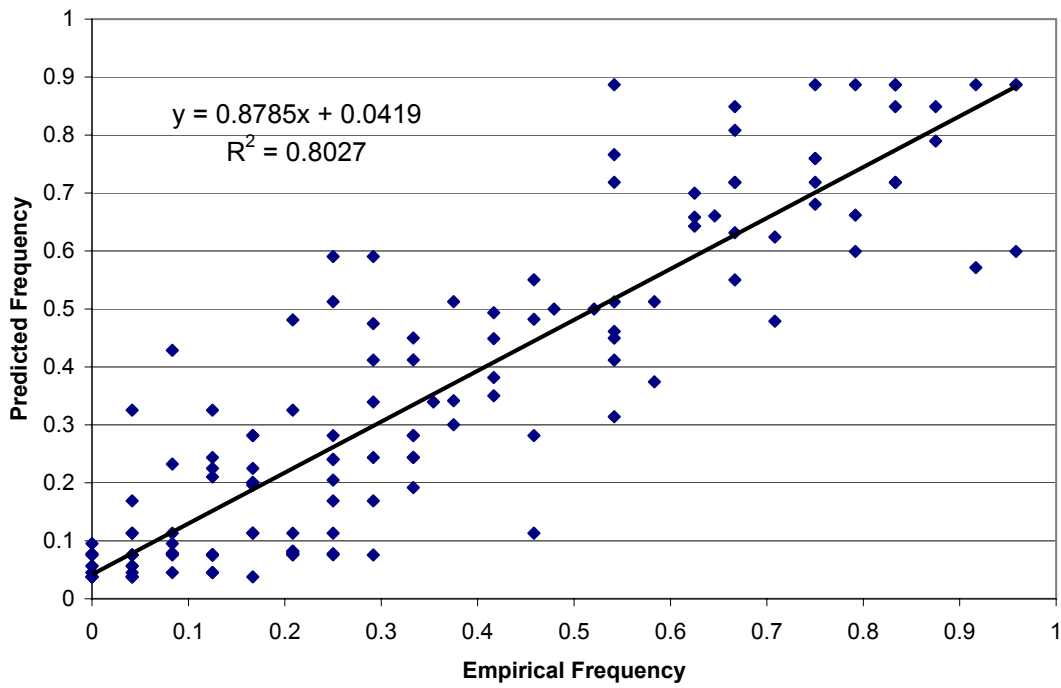


Figure 3a: Predicted Frequencies of Nash Equilibrium for Matrix Games




Figure 3b: Predicted Frequencies of Nash Equilibrium for Entry and Mixed Games

