

A cognitive hierarchy theory of one-shot games: Some preliminary results

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 6, 2002

¹This research was supported by NSF grant SES-0078911. *This draft is not intended for publication.* It simply summarizes a variety of results in a format accessible to interested readers through web-based archiving. 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 modelling 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 (.98,2.40) and a median of 1.65 across 80 experimental samples of matrix games, entry games, mixed-equilibrium games, and dominance-solvable p-beauty contests. 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* modelling 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 has taken 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.

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

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.

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. 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 notes how the CH model can account for cognitive details and also sketches how CH and the QRE approach can be compared. Section VII concludes and points out directions for further research.

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.

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 s_i^j . Denote other players' (denoted $-i$) 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^j by $P_k(S^j)$. So, we have $P_0(S^j) = \frac{1}{m_i}$.

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

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).

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 $g_k(h) = f(h) / \sum_{i=0}^{k-1} f(i)$, $\forall h < k$ and $g_k(h) = 0$, $\forall h \geq k$.

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

The normalized-beliefs assumption $g_k(h) = f(h) / \sum_{i=0}^{k-1} f(i)$ 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 $D(1) = 1 - f(0) + \sum_{h=1}^{\infty} (f(h) - 0) = 2 - 2e^{-\tau}$. 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 $D(2) = 2 - 2e^{-\tau}(1 + \tau)$, which is smaller than the size of the 1-step thinkers' belief error, $D(1)$. In fact, it is easy to show⁵ 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. (In Stahl and Wilson's (1995) terminology the highest-step thinkers become “worldly”.)

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. While we have endogenized the mean number of thinking steps, this

⁵Use the sum of the absolute deviations to measure the distance of the normalized distributions from the true distribution. The total absolute deviation is

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

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

convergence property is a clue about why players will do only a few steps of thinking— it doesn't pay to think too hard, because doing k steps and $k+1$ steps yields roughly the same expected payoff. 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.

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). The results of this sort of estimation are reported below. A good way to proceed is to posit some sensible principles $f(k)$ should satisfy and see what distributions satisfy them. We propose three principles:

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).⁶ 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 $k < k^*$, $\frac{f(k-1)}{f(k)} > \frac{f(k-2)}{f(k-1)}$ (lower convexity). The upper convexity condition implies that the distribution $f(k)$ drops

⁶An alternative principle is that $f(k) = \phi \cdot f(k-1), \forall k$ and $\phi = \frac{\sqrt{5}-1}{2}$, so that $f(k)$ is always declining and $\tau = \Phi = \frac{1+\sqrt{5}}{2}$, the golden ratio. This yields the Boltzmann-like distribution $f(k) = (1-\phi) \cdot \phi^k, \forall k$, which fits the statistical distribution of particles with different energy quantum levels.

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 dropoff 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, $f(k) = e^{-\tau} \cdot \frac{\tau^k}{k!}$.

The Poisson distribution can also be derived as the steady state of a Markov transition process in which players can either remain in k steps of thinking or move one step forward or backward to become $k-1$ or $k+1$ with equal probability given by $P_{k,k-1} = P_{k,k+1} = \frac{k!}{\tau^k}(1 - P_{0,0})$, $\forall k > 0$ and any given $P_{0,0}$.

The Poisson distribution has only one parameter, τ , which is its mean *and* its variance. This simplicity has obvious advantages in estimation.

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.⁸ This enables us to directly links intuitions

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

⁸This one-step precision property is not true of most distributions, of course. For example, if a distribution is Gaussian then forcing the mode or mean or median to be a specific value does not restrict the variance. Forcing the coefficient of variation to have a particular value does not imply a particular mean or variance.

about $f(k)$ to values of τ .

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 $f(1)$ is most common and is as large as possible relative to the neighboring frequencies $f(0)$ and $f(2)$ ⁹. Then $\tau = \sqrt{2} = 1.4142$. Or 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)¹⁰. 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

⁹i.e., τ maximizes the minimum $\{\frac{f(1)}{f(0)}, \frac{f(1)}{f(2)}\}$ or minimum $\{f(1) - f(0), f(1) - f(2)\}$.

¹⁰Condition (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 $\tau = \Phi$.

1-2 steps of thinking are typical¹¹. Section III reports formal estimation from a wide variety of one-shot games (80 games in total). Most estimates are between 1 and 2; the median across all 80 games is 1.65.

Note that one should not expect the average amount of thinking τ to have a universally constant value. It is like a risk-aversion parameter or a discount factor. Values of those parameters are typically not derived from first principles and are not expected to be constant. Discounting and risk-aversion vary across people (and even across a person's life; children are measurably more impatient than adults) and situations; τ probably does too. The hope is simply to find a range of τ values which are plausible, and regular enough to permit us to make guesses about behavior in new games with some confidence. And because the CH model has cognitive detail, τ *should* change in response to certain kinds of treatment effects. For example, people who are more analytically skilled or have special training in game theory will probably exhibit higher values of τ (more strategic thinking), just as people who are treated for fear of flying act as if a parameter characterizing their aversion to flight risk was changed by therapy.

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.¹² 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's, 19??, 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

¹¹See also Nagel, 1995; Stahl and Wilson, 1995; Costa-Gomes, Crawford, and Broseta, 2001.

¹²The 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 Sjostrom, 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.

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 a diffuse, or 'ignorance' prior about what others might. 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 focussed 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 second 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. Two other models are close to ours in style. 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, economizing on the many parameters used by Stahl and Wilson, and going in a different direction than Capra and Goeree and Holt. 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.

2.5 Alternative specifications

Once the mutual consistency of choices and beliefs is relaxed, there are many ways to specify choices and beliefs that are *not* consistent. The Poisson model in which k -step thinkers believe everyone else does $k - 1$ or fewer steps is one specification, but others spring to mind.

One alternative specification is to assume that step- k thinkers believe everyone else

is doing $k - 1$ steps (i.e., $g_k(h) = I(k - 1, h)$ where $I(x, y)$ is an identity function). Call this the “k-1” specification. Preliminary estimates showed that the k-1 model fits about as accurately as our specification in three sets of matrix games.¹³

However, the $k - 1$ model has some unfortunate properties. The $k - 1$ model is a freeze-frame version of Cournot dynamics, in which a player always believes others will repeat their choices in the most recent period and best responds to that belief. Since it is possible for play to cycle endlessly in Cournot, the $k - 1$ model can cycle too.

Furthermore, in the $k - 1$ model, players doing more and more steps of thinking *do not* become more worldly— in fact, their beliefs *diverge* from rational expectations as k increases.¹⁴ The belief deviations in the $k - 1$ model are also larger than in our specification; in a sense, the $k - 1$ thinkers have “less rational” expectations than in our approach.¹⁵ Furthermore, because k -step thinkers’ beliefs do not converge to the correct distribution in the $k - 1$ model, their beliefs embody a double dose of overconfidence. Two-step thinkers, for example, think that all players are doing one step of thinking, and think that all the one-step thinkers are completely deluded in thinking there are 0-step thinkers.

It is also easy to find games in which the $k - 1$ specification fits data very poorly. In the market entry games discussed in section IV below, the $k - 1$ model predicts a step function— the rate of entry into a capacity-constrained market will depend only on whether the capacity is less than half of the number of entrants, or more than half. But the data are surprisingly monotonic in the capacity, so the predicted step function is a poor approximation.¹⁶ Another example is asymmetric matching pennies, shown in Table 1.

¹³The $k - 1$ assumption is easy to work with theoretically because the sequence of predicted choices can be computed by working up the hierarchy without using any information about the true distribution $f(k)$.

¹⁴The total absolute deviation for the $k - 1$ model is

$$D_{k-1}(k) = \sum_{h=0}^{k-2} f(h) + 1 - f(k-1) + \sum_{h=k}^{\infty} f(h) = 2(1 - f(k-1)) \quad (2.6)$$

. This figure falls as k approaches the distribution mode (τ) then rises again, which means the beliefs of the highest step thinkers (beyond τ) are *furthest* from the truth.

¹⁵Recall that in our approach, the sum of absolute belief deviations is $D(k) = 2(1 - \sum_{h=0}^{k-1} f(h))$. This is smaller than $D_{k-1}(k)$ for any $k > 0$ because $2(1 - \sum_{h=0}^{k-1} f(h)) < 2(1 - f(k-1))$.

¹⁶In the entry games, as you increase k the $k - 1$ model decision rules alternate back and forth between entering at low c (i.e., c less than half the number of entrants) and staying out at high c , and

Table 1: Asymmetric matching pennies

	L	R
T	$x, 0$	$0, 1$
B	$0, 1$	$1, 0$

For $x > 1$ the model in which players think everyone is one step below them makes the same prediction for every value of x . But row players actually choose T more often when x is larger, a fact which is anomalous for the $k - 1$ specification but is predicted by the CH model specification.¹⁷

In the CH model, players who do k steps of thinking are not aware that others might be thinking like they do (or even thinking harder). An alternative approach is to make players “self-aware” so that k -step players’ beliefs include the possibility that there are others doing k steps like themselves (e.g., $g_k(c) = f(c)/(\sum_{c=0}^k f(c))$ for $0 \leq c \leq k$ and $g_k(c) = 0$ otherwise).

Selten (1998) argues that the “circular” reasoning implicit in self-awareness is less cognitively plausible than a purely sequential procedure. Self-awareness is deliberately excluded from the CH model for that reason, and several others. One reason is that overconfident players will doubt others are thinking as much as they themselves are. If players all think they are ‘smarter’ (or harder-thinking) than others then they *will* neglect the possibility that others are thinking at their step. This sort of overconfidence about relative skill is well-documented in some economic domains (e.g., Roll, 1984; Camerer and Lovallo, 1999; Benabou and Tirole, 2002).

Including self-awareness also leads to a model which is very much like a noisy equilibrium or quantal response equilibrium (QRE) model, because including self-awareness reintroduces an element of the mutual consistency which is precisely what the CH approach jettisons. To see this, first note that the relative proportion of 0 and 1-step

the opposite pattern. Aggregating these decision rules will produce a step function in which the rate of entry is constant for $c < .5N$ then switches to a higher rate, which does not look at all like the monotonicity in the data.

¹⁷For column players, the 0- and 1-step thinkers randomize equally over L and R, 2- and 3-step thinkers choose R, 4- and 5-step thinkers choose L, and so forth in a two-step cycle. For row players, 0-step thinkers randomize equally over T and B, 1- and 2-step thinkers choose T, 3- and 4-step thinkers choose B, and so forth in a two-step cycle. These best-response cycles do not depend on $f(k)$ or on x .

thinkers, conditional on thinking only up to 1 step, is $f(0)/(f(0) + f(1)) = 1/(1 + \tau)$. For large τ this fraction will be small, which means 1-step thinkers believe most others are 1-step thinkers too (with a small fraction of 0-step randomizers thrown in the mixture). The 1-step thinkers' optimal choices will be the solution to a recursive equation which requires (approximate) mutual consistency—bringing us back near Nash equilibrium. But the point of the cognitive hierarchy approach was to *improve* on the predictive accuracy of Nash equilibrium for one-shot games; circling back towards equilibrium defeats the purpose of creating a different approach.

Self-awareness also adds computational difficulty, compared to the CH specification, because it requires solving for fixed points. This is especially cumbersome in games with large strategy spaces or many players. Finally, and perhaps most importantly, in earlier work we *did* compare CH models with and without self-awareness. Adding self-awareness always reduced explanatory power, often substantially.¹⁸

3 Theoretical properties of the Poisson 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 mentions a few simple properties that can be derived from it.

3.1 Convergence to equilibrium in 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. (That's just the $k-1$ specification mentioned in the previous section on alternative approaches.) 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

¹⁸The log-likelihood values for the CH models with and without self-awareness are -1265 vs -1115 for Stahl and Wilson (1995), -1802 vs -1740 for Cooper and Van Huyck (2001) and -570 vs -555 for Costa-Gomes et al. (2001).

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.

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 a half 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. Denote the normalized cumulative entry function for all steps up to and including k by $E(k, c) : c \rightarrow [0, 1]$.

We have:

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

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

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 $c < 1/2$

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

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 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 $e(k, c)$).

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 Bargaining: Nash demand games and ultimatum 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.¹⁹ 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.

Behavior in ultimatum games is similar. In ultimatum games, a Proposer offers a division of a unit pie and a Responder accepts or rejects; if she rejects both earn nothing. Suppose Responders choose a rejection threshold between 0 and .5²⁰ and reject all offers above their threshold. Since the probability of accepting an offer of size x is $2x$, and the Proposer earns $1 - x$ if the offer is accepted, a 1-step thinker’s expected payoff of $2x(1 - x)$ is maximized by offering half. Since a Responder doing one or more steps of thinking accepts any offer, a 2-step Proposer faces a perceived acceptance probability of $(f(1) + 2xf(0))/(f(1) + f(0)) = (\tau + 2x)/(\tau + 1)$. Expected payoffs are maximized at $x_2^* = (2 - \tau)/4$, which is .125 for $\tau = 1.5$. If we mix together random offers by 0-step Proposers, offers of .5 from 1-step Proposers, and offers of $(2 - \tau)/4$ from 2-step proposers, the expected offer is .38 when $\tau = 1.5$ ²¹, which is close to the empirical average

¹⁹Suppose player 1 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.

²⁰This assumption is admittedly ad hoc. If we assume Responders choose rejection thresholds between 0 and 1 the 2-step player offers are $x_2^* = (1 - \tau)/2$ which go to zero “too fast” compared to data (in which zero offers are rare). We take this and other examples to suggest that a more refined and general theory of 0-step behavior is an important challenge for future work.

²¹The expected offer is $.5(g(0) + g(1)) + g(2)(2 - \tau)/5 = (4 + 8\tau + 2\tau^2 - \tau^3)/(8 + 8\tau + 4\tau^2)$.

observed in many countries and populations (e.g., Camerer, 2003, chapter 2; Henrich et al., 2002).²² This simple exercise is not meant to be a full model of ultimatum bargaining, which surely contains some element of social preference or repeated-game instinct. The point is just that even without a concept of social preferences, simple structural forces can lead to bargaining offers which are closer to equal splits than standard models predict, simply because of limited reasoning²³.

3.4 Effects of group size dominance-solvable and stag hunt games

Some interesting effects of group size emerge from the thinking-steps model. These effects are plausible but are not predicted by Nash equilibrium.

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 also an interesting behavioral effect of group size. 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.

Another example is stag hunt. Imagine a stag-hunt game in which each of n players choose either H or L. Players earn 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).

²²Including higher-step thinkers will lower the expected offer, but only by a little (to around .30) if $\tau \approx 1.5$ because there are so few players using three or more steps of thinking.

²³See also Johnson et al, 2002; Gale, Binmore, and Samuelson, 1995; Roth and Erev, 1995; McKelvey and Palfrey, 1998).

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 CH Poisson 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.²⁴

We fit four data sets: 33 matrix games with 2-4 strategies from three data sets; 22

²⁴Simulations 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 K were created using the CH value with a known τ , for six games from the mixed-strategy game sample (one $K \times K$ game for $K \in 2, 3, 4, 5, 6$). 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 ($N=20$) in 2×2 games. See Appendix for details.

games with mixed equilibria (new data); the binary entry game described above (new data); and 24 samples of subjects playing variants of the dominance-solvable ‘p-beauty contest game’.²⁵

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 (see Appendix for details). 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 $c - 1$.²⁶ Subjects were shown five capacities $c = 2, 4, 6, 8, 10$ in a fixed random order, with no feedback.

The 24 p-beauty contest games were taken from previously published results (Nagel, 1995; 1999; Ho, Camerer and Weigelt, 1998), from 2- and 3-player games conducted in the four 12-subject sessions with Caltech and PCC students, from unpublished data collected by Ho, Camerer and Weigelt, and from convenience samples collected by author Camerer with various audiences playing for \$20.

At this exploratory stage, there are three questions: Is the estimated value of τ reasonably regular across games with very different structures? When τ is unusually low or high, is this a clue about how the model might be improved? How does the CH Poisson specification compare to QRE and Nash equilibrium?

²⁵We are also estimating the model on 7 sender-receiver signaling games studied by Banks, Camerer and Porter (1994) to explore which signaling game refinements (intuitive criterion, divinity, and universal divinity, etc.) predict best when there are multiple Nash equilibria. The samples from first-period play are too small for reliable estimation so we are currently collecting more data.

²⁶This 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 Lovo, 1999).

4.1 How regular is τ ?

Table 2 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 56 game-specific τ estimates are high (4 or more) and some are zero. Including the estimates from p -beauty contests reported later, the interquartile range across 80 estimates is (.98,2.40) and the median is 1.65. An Appendix Table shows bootstrapped 95% confidence intervals for 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.

In three of the Stahl-Wilson games (2, 6, 8) the estimated value of τ is zero. These games are clues about when the model can fail badly because $\tau = 0$ is simply random choice. Their game 2 is shown in Table 3. (The Table shows only the row player payoffs since the game is symmetric.)

Let's figure out what goes wrong. One-step thinkers will choose M. Two-step thinkers will choose B for reasonable τ (above .42) since B is a reasonable response against 0-step thinkers and a best-response against 1-step thinkers who choose M. But only 25% and 13% of players actually choose those strategies, M and B. The Nash strategy T is chosen most often (63% of the time). The CH model cannot explain this because it is impossible to 'reach' strategy T by steps of iterated reasoning. The reason is that T is a poor response to a low number of steps of thinking. If some low-step thinkers chose it, then higher-step thinkers would lock in to the equilibrium strategy and choose it also, but low-step thinkers do not gravitate toward it. Put differently, playing M or B are nearly best responses to T, so subjects have to lock sharply in to strategy T for it to be chosen frequently. (As a result, QRE does not fit this game well either.) Choosing T seems to require a leap of faith in which subjects somehow deduce that a very large fraction of other subjects will almost surely choose T.²⁷ Since estimates of $\hat{\tau} = 0$ are so rare in other games, this suggests there is something special about these games, or the Stahl-Wilson procedure, which encouraged an unusual amount of Nash thinking.²⁸ Perhaps

²⁷The set of beliefs in which T is optimal is a small sliver of the P(T), P(M) simplex in which $P(M) > (55/80) - (60/80)P(T)$ and $P(M) < (45/65) - (40/65)P(T)$. Unless P(M) lies in a very small interval when P(T) is low, beliefs only lie in this sliver if P(T) is high.

²⁸In the Costa-Gomes et al. sample, there are also games in which Nash strategies do not correspond to one- or two-step thinking and they are played less often so the CH model does not produce low τ values. In fact, Stahl and Wilson used a mean-matching procedure in which each subject was randomly matched with every other subject and paid according to their mean payoff. This procedure may have

Table 3: Payoff tables and actual frequencies, Stahl and Wilson (1995) game 2

Game 2	payoff table			actual frequency	decision rules selecting strategy	
	T	M	B			
	T	75	40	45	0.63	Nash
	M	70	15	100	0.25	1-step, RE
	B	70	60	0	0.13	2-step

anticipating this, Stahl and Wilson included a special Nash-type player, which helps fit the games 2,6, and 8 that the CH model does not fit well. In general, however, including such a player does not appear to be necessary in most other one-shot games.

4.2 Which models fit best?

Table 4 shows log likelihoods (LL) and mean-squared deviations for several model estimated game-by-game or with common parameters across games in a dataset.²⁹ This Table answers several questions. Focussing 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 QRE (and hence, better than Nash), or about as well, except in the Stahl-Wilson games when common parameters are imposed. This result does *not* mean QRE research (which imposes mutual consistency but relaxes optimization) should be abandoned in favor of the CH approach (which does the opposite, relaxing consistency and retaining optimization); our view is that both approaches should be explored further. But the relative success of CH in many games is an indication that mutual consistency is not *necessary* to produce a model that fits data from one-shot games reasonably well.

Table 4 also reports fits from a general CH model in which the frequencies of k -step encouraged players to form strategically-thoughtful beliefs and catalyzed a large fraction of Nash play.

²⁹When 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 , which penalizes the resulting LL a bit.

thinkers, $f(k)$, are not constrained to satisfy the Poisson distribution (and truncated at six steps).³⁰ Except for the Stahl-Wilson data (once again), imposing this 6-parameter general specification degrades fit very little compared to the Poisson distribution. Table 5 shows the fractions of players estimated to use each level in the general specification; these fractions are reasonably close to those constrained by the Poisson distribution.

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 1a-b represents a distinct strategy in each of the 33 matrix games (Figure 1a) and 22 mixed games (Figure 1b), 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 2a-b show the corresponding figures comparing actual frequencies and the Nash predictions. In Figure 2a 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 2b 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 1a-b with 2a-b shows that the CH 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.

Games with many strategies create a special challenge for estimating models of limited thinking. When these models assume best response, if the number of steps of thinking is limited to n empirically plausible numbers (e.g., 6), then the model only predicts

³⁰The frequencies $f(k)$ are constrained in a small way to improve identification, which is essentially harmless in terms of fit. The constraint imposed is that the $f(k)$ function should be inverted-U shaped in k . That is, if $f(k) < f(k-1)$ for a particular k —that is, the distribution function turns downward—then $f(k+x) < f(k)$ for any positive integer x —i.e., once the $f(k)$ distribution turns downward it cannot rise up again. This constraint is necessary because when the estimates are not constrained in this way, it is possible to have, say, a large fraction of 0,1, and 2 subjects, but no 3-step subjects. But 4-step subjects who have (normalized) beliefs about this distribution will simply ignore the 3-step types. As a result, they will choose the same strategy a 3-step type would choose. So the unconstrained estimation can place zero $f(k)$ values anywhere in the distribution and produce precisely the same pattern of best responses (and hence, fit) as an alternative specification in which the zero is removed. In econometric language, there is a severe identification problem. One way to eliminate the possibility of these unidentified insertions of zero $f(k)$ types is to force the distribution to *not* wake up again after a zero $f(k)$ and produce positive values of $f(k+1)$. Happily, imposing this no-inverted-U constraint degrades LL very little. Across the four data sets, the reduction in LL is only 40, 0, 1, 14, and 0 points so the constraint is essentially harmless.

$n - 1$ separate choices from the higher-step thinkers (i.e., those using one step or more). In games with many strategies—like pricing or location games—this means the high-step thinking model will only predict a small fraction of the strategies that are actually chosen. Maximizing the likelihood of the observed choices is a poor method for assessing how well such a model fits because most observed choices will have predicted likelihoods of zero, regardless of the value of τ . Likelihood estimation throws away information because a model which predicts strategies very close to those which are chosen, but is not exactly right, will have the same likelihood as a model which predicts strategies which are far away from the data.³¹

An instructive example is “ p -beauty contest games”. In this class of games players choose numbers from a bounded interval, say $[0,100]$. Each player’s payoff depends on how close their number is to some multiple p of a summary statistic (typically the average) of all the choices. In a frequently-studied game, the statistic is the average, $p = 2/3$, and the player whose number is closest earns a fixed sum of money (see Nagel, 1999, for a recent review). In this game the 0-step thinkers randomize so 1-step players will best-respond to an expected average of 50, choosing 33. Two-step players best-respond to a mixture of 50 and 33, choosing $(50f(0) + 33f(1))/(f(0) + f(1))$; if $f(k)$ is Poisson their choice will be $(50 + 33\tau)/(1 + \tau)$. Note that this model predicts that a fraction $f(1)$ of players will choose exactly 33. In experiments, however, number choices near 33 are common but the exact response 33 is only chosen around 5-10% of the time. In likelihood estimation, the fact that the model is close to the data, but slightly off, is ignored.

Given that likelihood maximization does not adequately characterize how well the model is approximately fitting, it makes sense to choose a different estimation method.³² In the generalized method of moments, parameters are chosen to minimize a weighted average of the mean and higher moments (particularly variance). We use a simple method of moments in which τ is chosen to make the predicted mean as close as possible to the actual mean. We can then see whether the predicted variance is close to the actual

³¹Likelihood-maximization severely penalizes a model which assigns low probability to a choice that is observed (since $\log(\epsilon)$ becomes hugely negative as $\epsilon \rightarrow 0$). In the CH model, when lots of strategies are actually chosen, but only a few are predicted to be chosen, the only way for the model to maximize likelihood is to choose a low value of τ (since the 0-step types smear probability across all strategies).

³²An alternative is to relax the assumption that players are best-responding to allow stochastic response. This means the predicted choices of 33 will then be smoothed *around* 33 and will put more likelihood in the right places (e.g., rounded-off choices of 30 or 35). This is actually computationally more difficult, and adds an extra parameter which is of little interest.

variance; if it is we take that as an out-of-moments clue that the model is capturing some important features of the data.

Table 6 shows estimates of τ in 24 beauty contest games, which were chosen to minimize the (absolute) difference between the predicted and actual mean of chosen numbers. The table is ordered from top to bottom by the mean number chosen. The first seven lines show games in which the equilibrium is not zero; in all the others the equilibrium is zero. The first four columns describe the game or subject pool, the source, group size, and total sample size. The fifth and sixth columns show the Nash equilibrium and the difference between the equilibrium and the average choice. The middle three columns show the mean, standard deviation, and mode in the data. The mean choices are generally far off from the equilibrium; they choose numbers which are too low when the equilibrium is high (first six rows) and numbers which are too high when the equilibrium is low (lower rows). The rightmost six columns show the estimate of τ from the CH Poisson model, and the mean, prediction error, standard deviation, and mode predicted by the best-fitting estimate of τ , and the 90% confidence interval for τ estimated from a randomized resampling (bootstrap) procedure.

There are several interesting patterns in Table 6. The prediction error of the mean (column 13, “error”) are extremely small, less than .6 in all but two cases. Of course, this is no surprise since τ is estimated (separately in each row) to minimize this error. The pleasant surprise is that the predicted standard deviations and modes which result from the error-minimizing estimate of τ are also fairly close (across rows, the correlation of the predicted and actual standard deviation is .72).

The values of τ have a median and mean across rows of 1.30 and 1.61, close to the golden ratio and $\sqrt{2}$ values derived from simple restrictions earlier in this paper. The confidence intervals have a range of about one in samples of reasonable size (above 50 subjects).

Outlying low and high values of τ are instructive. Estimates of τ are quite low (0-.1) when $p > 1$ and, consequently, the equilibrium is at the upper end of the range of possible choices (rows 1-2). In these games, subjects seem to have trouble realizing they should choose very large numbers when $p > 1$ (though they equilibrate rapidly by learning; see Ho, Camerer and Weigelt, 1998). Low τ 's are also estimated among the PCC subjects playing 2- and 3-player games (rows 8 and 10). High values of τ (≈ 3 -5) appear in games where the equilibrium is in the interior, 72, (rows 7-10)– small incremental steps

toward the equilibrium in these games produce high values of τ . High τ values are also estimated in games with an equilibrium of zero when subjects are professional stock market portfolio managers (row 19), Caltech students (row 20), game theorists (row 24), and subjects self-selecting to enter newspaper contests (row 25). The latter subject pools show that in highly analytical and educated subject pools (especially with self-selection) τ can be much higher than in other subject pools.

A sensible intuition is that when stakes are higher, subjects will use more steps of reasoning (and may think others will think harder too). Rows 3 and 6 compare low stakes (\$1 per person per period) and high stakes (\$4) in games with an interior equilibrium of 72. When stakes are higher τ is estimated to be twice as large (5.01 versus 2.51), which is a clue that some sort of cost-benefit analysis may underlie steps of reasoning.

Notwithstanding these interesting outliers, there is also substantial regularity across very diverse subject pools. About half the samples have confidence intervals which include $\tau = 1.5$. Subsamples of corporate CEOs (row 13), high-functioning 70-year old spouses of memory-impaired patients (row 15), and high school students (row 16) all have τ values from 1.1-1.7.

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 sample 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 4 reports the result of this sort of cross-game estimation. Both the CH Poisson and QRE models fit cross-game a little less accurately than when estimates are common within games. QRE degrades particularly badly in the Costa-Gomes et al. and mixed-equilibrium games. This is not surprising since the free parameter in the QRE model is a response sensitivity which is sensitive to changes in payoff scales (e.g., McKelvey, Palfrey and Weber, 2000).

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 7 reports the profits players would have earned if they used the CH or QRE models 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. QRE and Nash equilibrium add a little less value, and both subtract 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 Cognitive implications and comparisons with QRE

6.1 Comparing QRE and thinking-steps models

The CH model retains optimization but relaxes mutual consistency. Quantal response equilibrium 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 actual distribution.

In empirical applications, QRE and CH will usually predict deviations in the same direction from Nash equilibrium, and they should be treated as alternative paths which both deserve exploration. However, keep in mind that in the results above, QRE fit and predicted a little less accurately than CH (except in the Stahl-Wilson matrix games).

Because they will often predict similar deviations, it is useful to carefully distinguish how they differ. QRE will generally make different predictions when games are subject to “inessential transformations” (see Dalkey, 1953; Ho and Weigelt, 1996). For example, in QRE “cloning” strategies (adding precisely equivalent strategies) will generally increase the frequency of play of the set of cloned strategies (because players who noisily best-respond will play these strategies equally often).³³ In CH, in contrast, cloning strategies will only increase how often the cloned strategy (set) is played for 0-step thinkers. If the cloned strategy has the highest expected payoffs, then higher-step thinkers are assumed to randomize across the set of equally-good (and best) responses so they will play a *set* of best responses just as often as if one strategy was a uniquely best response. A similar property arises if strategies are amalgamated rather than cloned. Mookerjee and Sopher (1997) found that amalgamating strategies did not change how frequently they were played, which goes against QRE and is more consistent with CH.

Another subtle contrast between the two models is when some strategies are nearly dominated. For example, if one strategy yields ϵ less than another strategy, than as $\epsilon \rightarrow 0$ the QRE frequencies of the two strategies will become equal. Since the CH approach assumes best responses, for any $\epsilon > 0$ the predicted frequency of the dominated strategy will be lower than the predicted frequency of the strategy which dominates it. This contrast sharpens the difference between the two approaches. If subjects spot dominated strategies (regardless of the degree of dominance) and never play them, the CH approach will predict better than the QRE approach; oppositely, if subjects do not notice or care about small degrees of dominance then they will play nearly-dominated strategies relatively often consistent with QRE.

³³This is because the logit has the *Independence from Irrelevant Alternatives* property. One could presumably develop a hierarchical equivalent that does not exhibit this property.

Another difference is that the CH model naturally generates heterogeneity— “spikes” which can potentially match spikes in data. (The p -beauty contest is an example.) QRE, in contrast, predicts a smooth statistical distribution with no spikes (the same is true of Capra’s, 1999, and Goeree and Holt’s , 2002, models). In the same way, CH can easily explain endogeneous purification but the simplest form of QRE cannot (in QRE each player mixes with the same statistical distribution across strategies).

Finally, our analysis of mixed games shows that when the true model is CH, MLE estimation recovers the correct τ parameters in modest samples (around 50) (see Appendix). However, when samples are small, sampling error is ‘accurately’ fit by QRE with a low response sensitivity λ . So we suspect that MLE and other techniques will generally underestimate the true value of λ (i.e., estimates are biased downward) in small or medium samples.

6.2 Cognitive measures

The CH model should be taken seriously as a prediction about the kinds of algorithms that players use in thinking about games. This means that cognitive data other than choices can— like belief-prompting, response times, information lookups, or even brain imaging — can, in principle, be used to test the model.³⁴

Several studies show that prompting players for beliefs about what others will do actually changes their choices, typically moving them closer to equilibrium. A simple example was first demonstrated by Giovanna Devetag and Eldar Shafir and replicated by Warglien, Devetag and Legrenzi (1998). Their game and results are shown in Table 8. If players think others are step 0 (randomizing), choosing X yields an (expected)

³⁴See Camerer, Loewenstein, and Prelec, 2002. The use of brain imaging will sound farfetched to most economists. But Glimcher, 2002, chapter 13, reports the existence of ‘equilibrating’ neurons in monkeys which fire in rough proportion to expected payoffs of strategies, as the monkeys play a 2x2 ‘work-or-shirk’ game with a mixed equilibrium against a computerized opponent. When play is out of equilibrium, neurons fire more actively when the monkey plays the strategy with the higher expected payoff. This activation guides the monkey to play the ‘better’ strategy more often, which eventually produces equilibration. After equilibration, when both strategies have equal expected payoffs, the neurons fire at the same rate after each of the two strategies are played, so that the brain is ‘recognizing’ equilibration. Since the human brain is essentially the monkey brain with some extra cortex, which is used largely for planning and understanding social structure and language, it is likely that humans have a similar neural circuitry which encodes expected payoffs and guides equilibration in simple games.

Table 8: How belief-prompting promotes dominance-solvable choices by row players (Warglien, Devetag and Legrenzi, 1998)

row move	column player		without belief	with belief
	L	R	prompting	prompting
X	60,20	60,10	.70	.30
Y	80,20	10,10	.30	.70

payoff of 60 rather than 45 from choosing Y. When players simply choose (with financial incentives) 70% of the row players choose X. When subjects are prompted to articulate a belief about what the column players will do before they choose, 70% then choose the dominance-solvable equilibrium choice Y (see also Croson, 2000; and Hoffman et al, 2000).³⁵ In the CH model, the fractions of X play are fit perfectly by $\tau = .58$ without belief-prompting and $\tau = 2.20$ after belief-promoting. This suggests that the effect of belief prompting is to encourage strategic thinking among the 0-step players and shift the entire distribution up by about a step and a half of thinking.

If the algorithmic reasoning in the CH model is taken seriously as a model of human cognition, then the model can be tested by jointly estimating both choices *and* cognitive variables. An easy variable to measure is response time. We are currently investigating whether response times are longer for choices that correspond to higher steps of thinking.

Another alternative is to directly measure the information subjects acquire in a game by forcing subjects to “look up” payoffs in games (as in Camerer et al (1994), Costa-Gomes, Crawford, and Broseta (2001) and Johnson et al (2002)). Information lookups are another cognitive measure which we expect to be correlated to thinking steps. Johnson et al show that how much players look ahead to future “pie sizes” in alternating-offer bargaining is correlated with the offers they make. Costa-Gomes et al show that lookup

³⁵Schotter et al (1994) found a similar effect of display and timing in games with Nash equilibria which are not subgame-perfect. In the simultaneous matrix form more players chose the Nash equilibrium, as if they did not reason through what others would do. Note that these display effects can be interpreted as focussing players’ attention in different ways, altering the number of thinking steps they are doing or what players think at different steps. We also observed a belief-prompting effect in beauty contest games (unpublished). When players simply made choices, 25% chose numbers above 50 in the first period. When forced to guess what the average choice would be, this figure fell to 15%. The samples were small so the effect is not significant but it goes in the same direction as the effects above.

patterns are clearly correlated with choices that result from various (unobserved) decision rules. These patterns are not proof that models based on steps of thinking are correct, but they do illustrate a fresh prediction that results from these models.

7 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.³⁶ 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.³⁷ 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.

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

³⁶Besides 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.

³⁷See Sonsino, Erev and Gilat, 2000; Sovik, 2000.

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 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.

8 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 overentry at low values of c and underentry 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 QRE and Nash approaches add 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 or QRE 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.

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,b).

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.

References

- [1] Babcock, Linda, George Loewenstein, Samuel Issacharoff, and Colin F. Camerer (1995): “Biased Judgments of Fairness in Bargaining,” *American Economic Review*, 85, 1337-43.
- [2] Babcock, Linda, and George Loewenstein. 1997. Explaining bargaining impasses: The role of self-serving biases. *Journal of Economic Perspectives*, 11, 109–26.
- [3] Banks, Jeffrey, Colin F. Camerer, and David Porter (1994): “Experimental Tests of Nash Refinements in Signaling Games,” *Games and Economic Behavior*, 6, 1-31.
- [4] Benabou, Roland and Jean Tirole (2002): Self-Confidence and Personal Motivation. *Quarterly Journal of Economics*, August 2002, 117, 871-915.
- [5] Bernheim, Douglas. (1984). Rationalizable strategic behavior. *Econometrica*, 52, 1007-1028.
- [6] Binmore, Kenneth. 1988. Modelling rational players: Part II. *Economics and Philosophy*, 4, 9-55.
- [7] Binmore, Kenneth, Avner Shaked, and John Sutton (1985): “Testing Noncooperative Bargaining Theory: A Preliminary Study,” *American Economic Review*, 75, 1178-80.
- [8] Binmore, Kenneth; Swierzbinski, Joe; and Chris Proulx. Does maximin work? An experimental study. *Economic Journal*, July 2001, 111, 445-464.
- [9] Bloomfield, Robert. Learning a mixed strategy equilibrium in the laboratory. *Journal of Economic Behavior and Organization*, 25, 1994, 411-436.
- [10] Brandts, Jordi and Charles Holt (1992): “An Experimental Test of Equilibrium Dominance in Signaling Games,” *American Economic Review*, 82, 1350-65.
- [11] Brown, George (1951): “Iterative Solution of Games by Fictitious Play,” In *Activity Analysis of Production and Allocation*, New York: Wiley.
- [12] Cabrera, Susana; C. Mónica Capra, and Rosario Gomez. (2002). 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, October.

- [13] Camerer, Colin F. 1990. Behavioral game theory. In R. Hogarth (Ed.), *Insights in Decision Making: A tribute to Hillel J. Einhorn*. University of Chicago Press.
- [14] Camerer, Colin F. (1995): “Individual Decision Making,” in *the Handbook of Experimental Economics* ed. by J. Kagel and A. Roth. Princeton: Princeton University Press, 587-703.
- [15] Camerer, Colin F. (2003): *Behavioral Game Theory: Experiments on Strategic Interaction*, Princeton:Princeton University Press.
- [16] Camerer, Colin F. and Ho Teck-Hua (2001): Strategic learning and teaching. In S. J. Hoch and H. C. Kunreuther (Eds.), *Wharton on Making Decisions*, New York: John Wiley, 159-176.
- [17] Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong (2002a): “Strategic Teaching and Equilibrium Models of Repeated Trust and Entry Games,” CalTech Working Paper, <http://www.hss.caltech.edu/camerer/camerer.html>
- [18] Camerer, Colin F., Teck-Hua Ho and Juin-Kuan Chong (2002b): “Sophisticated EWA Learning and Strategic Teaching in Repeated Games,” *Journal of Economic Theory*, 104, 137-188.
- [19] Camerer, Colin F., Eric J. Johnson, Talia Rymon, and Sankar Sen. (1994): 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, 27-47.
- [20] Camerer, Colin F., George Loewenstein, and Drazen Prelec (2002): “Gray Matters: How Neuroscience can Inform Economics,” under revision for *Journal of Economic Perspectives*.
- [21] Camerer, Colin F., and Daniel Lovallo (1999): “Overconfidence and Excess Entry: An Experimental Approach,” *American Economic Review*, 89, 306-18.
- [22] Camerer, Colin F. and Keith Weigelt (1993): 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, 355–396.

- [23] Capra, C. Mónica (1999): *Noisy Expectation Formation in One-shot Games*, Ph.D. thesis.
- [24] Chen, Hsiao-Chi, James W. Friedman, and Jacques-Francois Thisse (1996): “Boundedly Rational Nash Equilibrium: A Probabilistic Choice Approach,” *Games and Economic Behavior*, 18, 1832-54.
- [25] Cooper, David and J. Van Huyck (2001): “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/>
- [26] Costa-Gomes, M. and V. Crawford and B. Broseta (2001): “Cognition and Behavior in Normal-form Games: An Experimental Study,” *Econometrica*, 69, 1193-1235.
- [27] V. Crawford, “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.
- [28] Croson, Rachel T. A. (2000): “Thinking Like A Game Theorist: Factors Affecting the Frequency of Equilibrium Play,” *Journal Of Economic Behavior And Organization*, 41(3), 299-314.
- [29] Dalkey, Norman (1953): “Equivalence of information patterns and essentially determinate games” Contributions to the Theory of Games II, *Annals of Mathematical Studies*, 28, Princeton, NJ: Princeton University Press.
- [30] Fudenberg, D. and D. Kreps (1990): “A Theory of Learning, Experimentation and Equilibrium in Games,” Stanford University working paper.
- [31] Fudenberg, Drew and David Levine (1998): *The Theory of Learning in Games*. Cambridge: MIT Press.
- [32] Gabaix, Xavier and David Laibson (2000). Bounded rationality and directed cognition. Harvard University working paper.
- [33] Gale, John, Kenneth Binmore, and Larry Samuelson (1995): “Learning to be Imperfect: The Ultimatum Game,” *Games and Economic Behavior*, 8, 56–90.
- [34] Glimcher, Paul (2002): *Decisions, Uncertainty, and the Brain: The Science of Neuroeconomics*. MIT Press.

- [35] Goeree, Jacob K. and Charles A. Holt (1999): "Stochastic Game Theory: For Playing Games, Not Just for Doing Theory," *Proceedings of the National Academy of Sciences*, 96, 10564-10567.
- [36] Goeree, Jacob K. and Charles A. Holt (2002): "Stochastic Game: A theory of noisy introspection," *Games and Economic Behavior*, forthcoming.
- [37] Goeree, Jacob; Holt, Charles A. and Thomas R. Palfrey. Risk aversion in games with mixed strategies. University of Virginia Department of Economics, May 2000.
- [38] Grosskopf, Brit and Rosemarie Nagel (2001): Rational Reasoning or Adaptive Behavior? Evidence from Two-Person Beauty Contest Games, Harvard NOM Research Paper No. 01-09.
- [39] Harsanyi, John C., and Reinhard Selten (1988): *A General Theory of Equilibrium Selection in Games*. Cambridge, MA: MIT Press.
- [40] Haruvy, Ernan and Stahl, Dale O. (1998): "An Empirical Model of Equilibrium Selection in Symmetric Normal-form Games," University of Texas Department of Economics Working Paper.
- [41] Henrich, Joseph, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr, Herbert Gintis, Richard McElreath, Michael Alvard, Abigail Barr, Jean Ensminger, Kim Hill, Francisco Gil-White, Michael Gurven, Frank Marlowe, John Q. Patton, Natalie Smith, and David Tracer. (2002): "Economic man" in cross-cultural perspective: Behavioral experiments in 15 small-scale societies, Working paper, <http://webuser.bus.umich.edu/henrich/gameproject.htm>.
- [42] Ho, Teck-Hua, Colin Camerer, and Keith Weigelt (1998): "Iterated Dominance and Iterated Best-response in p -Beauty Contests," *American Economic Review*, 88, 947-969.
- [43] Ho, Teck-Hua and Keith Weigelt (1996): "Task Complexity, Equilibrium Selection, and Learning: An Experimental Study," *Management Science*, 42, 659-679.
- [44] Hoffman, Elizabeth, Kevin McCabe, and Vernon L. Smith (2000): "The impact of exchange context on the activation of equity in ultimatum games," *Experimental Economics*, 3, 5-9.

- [45] G. Huberman and A. Rubinstein, "Correct Belief, Wrong Action and a Puzzling Gender Difference," working paper 2000.
- [46] Johnson, E., C. Camerer, S. Sen and T. Rymon (2002): "Detecting Backward Induction in Sequential Bargaining Experiments," *Journal of Economic Theory*, 104, 16-47.
- [47] Kadane, Jay and Larkey, Pat (1982): Subjective probability and the theory of games, and reply. *Management Science*, 28, 113-120, 124.
- [48] 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.
- [49] Krishna, Vijay and Sjöström, Tomas (1998): On the Convergence of Fictitious Play. *Mathematics of Operations Research*, 23, pp. 479-511.
- [50] Lei, Vivian, Charles Noussair and Charles Plott (2001): "Nonspeculative Bubbles in Experimental Asset Markets: Lack of Common Knowledge of Rationality vs. Actual Irrationality," *Econometrica*, 69(4), 831-859.
- [51] Lieberman, Bernhardt. 1962. 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, pp. 203-220.
- [52] Lucas, R.G. (1986) "Adaptive Behavior and Economic Theory," *Journal of Business*, 59, S401-S426.
- [53] McKelvey, Richard D. and Thomas R. Palfrey (1995) "Quantal Response Equilibria for Normal-form Games," *Games and Economic Behavior*, 7, 6-38.
- [54] McKelvey, Richard D. and Thomas R. Palfrey (1998): "Quantal Response Equilibria for Extensive Form Games," *Experimental Economics*, 1, 9-41.
- [55] McKelvey, Richard, Thomas Palfrey and Roberto Weber (2000): "The Effects of Payoff Magnitude and Heterogeneity on Behavior in 2x2 Games with Unique Mixed Strategy Equilibria," *Journal of Economic Behavior and Organization*, 42(4), pp.523-548.

- [56] Messick, David M. Interdependent decision strategies in zero-sum games; A computer-controlled study. *Behavioral Science*, 12, 33-48.
- [57] Miyasawa (1961): On the convergence of learning processes in a 2x2 non-zero two-person game. Research memo 33, Princeton University.
- [58] Mookerjee, Dilip and Barry Sopher (1994): "Learning Behavior in an Experimental Matching Pennies Game," *Games and Economic Behavior*, 7, 62-91.
- [59] Mookerjee, Dilip and Barry Sopher (1997): "Learning and Decision Costs in Experimental Constant-sum Games," *Games and Economic Behavior*, 19, 97-132.
- [60] Nachbar, John (1990): "Evolutionary" selection dynamics in games: Convergence and limit properties. *International Journal of Game Theory*, 19, 59-89.
- [61] Nagel, R (1995): "Unraveling in Guessing Games: An Experimental Study," *The American Economic Review*, 85, (1995), 1313-1326.
- [62] Nagel, Rosemarie (1999). 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.
- [63] Nagel, Rosemarie, Antoni Bosch-Domnech, Albert Satorra, and Jos Garca-Montalvo. In press. One, two, (three), infinity: Newspaper and lab beauty-contest experiments. *American Economic Review*, <http://www.econ.upf.es/leex/working.htm>.
- [64] Ochs, Jack. Games with unique, mixed strategy equilibria: An experimental study. *Games and Economic Behavior*, 10, 1995, 202-217.
- [65] O'Neill, Barry. Nonmetric test of the minimax theory of two-person zerosum games. *Proceedings of the National Academy of Sciences*, 84, 1987, 2106-2109.
- [66] Pearce, David G. (1984): Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52, 1029-1050.
- [67] Prendergast, Canice. 1999. "The Provision of Incentives in Firms," *Journal of Economic Literature*, 37, 7-63.
- [68] Rapoport, Amnon and Amaldoss, Wilfred. 2000. Mixed strategies and iterative elimination of strongly dominated strategies: An experimental investigation of states of knowledge. *Journal of Economic Behavior and Organization*, 42, 483-521.

- [69] Rapoport, Amnon and Boebel, Richard B. 1992. Mixed strategies in strictly competitive games: A further test of the minimax hypothesis. *Games and Economic Behavior*, 4, 261-283.
- [70] Robinson, Julia (1951): "An Iterative Method of Solving a Game," *Annals of Mathematics*, 54(2), pp. 296-301.
- [71] Roll, R. (1984): "Orange Juice and Weather," *American Economic Review*, 74, 861-880.
- [72] Rosenthal, Robert W. (1989): "A bounded rationality approach to the study of noncooperative games," *International Journal of Game Theory*, 18(3), 273-92.
- [73] Roth, A.E. and I. Erev (1995): "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," *Games and Economic Behavior*, 8, 164-212.
- [74] Schotter, A., K. Weigelt and C. Wilson (1994): "A Laboratory Investigation of Multiperson Rationality and Presentation Effects," *Games and Economic Behavior*, 6, 445-468.
- [75] Selten, Reinhard. 1998. Features of experimentally observed bounded rationality. *European Economic Review*, 42, 413-436.
- [76] Shapley, Lloyd S (1964): "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, 1-28.
- [77] Shleifer, Andrei and Robert W. Vishny (1990): Equilibrium Short Horizons of Investors and Firms. *American Economic Review Papers and Proceedings*, May, 80, 148-153.
- [78] Smith, Vernon L., Gerry Suchanek and A. Williams (1988) "Bubbles, Crashes and Endogenous Expectations in Experimental Spot Asset Markets," *Econometrica*, 56, 1119-1151.
- [79] D. Soutsino, Ido Erev and S. Gilat (2000): "On the Likelihood of Repeated Zero-sum Betting by Adaptive (Human) Agents," Technion, Israel Institute of Technology.
- [80] Ylva Sovik (2000): "Impossible Bets: An Experimental Study," University of Oslo Department of Economics.

- [81] Stahl, Dale O. 1993. Evolution of $smart_n$ players. *Games and Economic Behavior*, 5, 604–17.
- [82] Stahl, D. O. (1998): Is Step-j Thinking an Arbitrary Modeling Restriction or a Fact of Human Nature? *Journal of Economic Behavior and Organization*, 37, 33-51.
- [83] Stahl, D. O. and P. Wilson (1995): "On Players Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior*, 10, 213-54.
- [84] Svensson, L.E.O. (2001): "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.
- [85] Tang, Fang-Fang. 2001. Anticipatory learning in two-person games: Some experimental results. *Journal of Economic Behavior and Organization*, 44, 221–32.
- [86] Tirole, J. (1985): "Asset Bubbles and Overlapping Generations," *Econometrica*, 53, 1071-1100 (reprinted 1499-1528).
- [87] M. Warglien, M. G. Devetag and P. Legrenzi (1998) "Mental Models and Nave Play in Normal Form Games," Universita Ca' Foscari di Venezia working paper.
- [88] Weibull, Jörgen (1995): *Evolutionary Game Theory*, Cambridge, MA: MIT Press.
- [89] Woodford, M. (2001) "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.

9 Appendix: Model recovery, mixed games, and bootstrapped standard errors

The 22 mixed games were taken from the review in Camerer (2003, chapter 3). They 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 Almadoss (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 Almadoss (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.

Table 9 below shows estimates of τ recovered from simulated data, created using the CH model, to see how well the estimation procedure recovers τ when the actual value is known. Each line shows a different value of "true" τ , for different sizes of simulated samples (n , either 20, 48, or 100), across a 2x2, 4x4, and 6x6 game, and then averaged over the three games. There is little bias in recovering the actual τ values (except for a slight upward bias when τ is small), although the 95% confidence intervals are rather wide when samples are of size 20, and for the 2x2 game. The key lesson is that small samples do not have much power, and 2x2 games are not very useful for estimating CH models. The problem is that each level of thinking, above 0, picks a distinct strategy; so when there are only two strategies several different levels all pick the same strategy, which means it is hard to identify how many levels are being used.

also insert table with bootstrapped standard errors (from coghi1102.xls)

Table 9: Estimates of τ from model recovery simulations

true		2×2		4×4		6×6		average across games	
τ	n	mean	90% CI	mean	90% CI	mean	90% CI	mean	90% CI
.5	20	0.62	(.05, 2)	0.53	(.2, .9)	0.51	(.3, .9)	0.56	(.19, 1.23)
	48	0.52	(.15, .95)	0.51	(.3, .75)	0.5	(.35, .65)	0.51	(.28, .77)
	100	0.52	(.25, .8)	0.51	(.35, .7)	0.5	(.4, .6)	0.51	(.34, .68)
1	20	1.08	(.5, 2)	1.02	(.55, 1.55)	0.97	(.6, 1.45)	1.04	(.57, 1.6)
	48	1.04	(.6, 1.6)	1.01	(.7, 1.35)	0.96	(.7, 1.2)	1.01	(.69, 1.37)
	100	1.01	(.7, 1.3)	1	(.8, 1.2)	0.96	(.75, 1.05)	1	(.78, 1.21)
1.5	20	1.55	(.95, 2)	1.52	(1, 2.15)	1.45	(1.05, 1.55)	1.48	(.96, 1.83)
	48	1.55	(1.1, 2)	1.51	(1.15, 1.85)	1.48	(1.4, 1.5)	1.5	(1.22, 1.75)
	100	1.54	(1.2, 1.95)	1.5	(1.25, 1.75)	1.5	(1.5, 1.5)	1.5	(1.3, 1.7)
2	20	1.96	(.85, 2.85)	2.05	(1.45, 2.7)	2.06	(1.55, 2.65)	2.01	(1.31, 2.68)
	48	2.02	(1.25, 2.5)	2.02	(1.65, 2.5)	2.01	(1.65, 2.4)	2.01	(1.59, 2.43)
	100	2.05	(1.8, 2.35)	2.01	(1.75, 2.3)	2.01	(1.75, 2.3)	2.01	(1.77, 2.3)
2.5	20	2.39	(1.2, 3)	2.45	(1.9, 2.9)	2.52	(1.9, 3)	2.46	(1.74, 2.98)
	48	2.49	(2, 3)	2.47	(2.05, 2.7)	2.53	(2.05, 3)	2.49	(2.06, 2.91)
	100	2.51	(2.15, 2.9)	2.49	(2.2, 2.65)	2.51	(2.2, 2.85)	2.5	(2.21, 2.8)

Table 2: 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 4: 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 t)¹</u>					
LL	-721	-1690	-540	-824	-150
MSD	0.0074	0.0079	0.0034	0.0097	0.0004
<u>Quantal Response (Game-specific I)</u>					
LL	-792	-1750	-574	-897	-150
MSD	0.0153	0.0148	0.0108	0.0203	0.0001
<u>Cognitive Hierarchy (Common General Distribution $f(k)$)</u>					
LL	-777	-1741	-554	-866	-150
MSD	0.0125	0.0132	0.0086	0.0175	0.0004
<u>Cognitive Hierarchy (Common t)</u>					
LL	-918	-1743	-560	-872	-150
MSD	0.0327	0.0136	0.0100	0.0179	0.0005
<u>Quantal Response (Common I)</u>					
LL	-843	-1838	-596	-1005	-151
MSD	0.0248	0.0269	0.0178	0.0456	0.0022
<u>Cross-dataset Forecasting</u>					
<u>Cognitive Hierarchy (Common t)</u>					
LL	-941	-1929	-599	-884	-153
MSD	0.0425	0.0328	0.0257	0.0216	0.0034
<u>Quantal Response (Common I)</u>					
LL	-862	-1980	-748	-1429	-166
MSD	0.0275	0.0500	0.0697	0.0501	0.0216
<u>Nash Equilibrium²</u>					
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 5: Probability Distribution of Thinking Levels for the General Cognitive Hierarchy Models

Data set	Stahl & Wilson (1995)	Cooper & Van Huyck	Costa-Gomes et al.	Mixed	Entry
----------	-----------------------	--------------------	--------------------	-------	-------

Frequency Estimates of the General Cognitive Hierarchy Models with Constraints ¹

Thinking Levels	Stahl & Wilson (1995)	Cooper & Van Huyck	Costa-Gomes et al.	Mixed	Entry
0	0.25	0.43	0.22	0.20	0.50
1	0.12	0.43	0.22	0.38	0.40
2	0.12	0.11	0.26	0.23	0.08
3	0.12	0.03	0.18	0.08	0.01
4	0.12	0.00	0.08	0.04	0.00
5	0.12	0.00	0.03	0.04	0.00
6	0.12	0.00	0.01	0.01	0.00
7 and Higher	0.00	0.00	0.00	0.00	0.00

Frequency of the Poisson Cognitive Hierarchy Models

Thinking Levels	Stahl & Wilson (1995)	Cooper & Van Huyck	Costa-Gomes et al.	Mixed	Entry
0	0.21	0.45	0.19	0.23	0.48
1	0.33	0.36	0.31	0.34	0.35
2	0.25	0.14	0.26	0.25	0.13
3	0.13	0.04	0.15	0.12	0.03
4	0.05	0.01	0.06	0.05	0.01
5	0.02	0.00	0.02	0.01	0.00
6	0.00	0.00	0.01	0.00	0.00
7 and Higher	0.00	0.00	0.00	0.00	0.00

Note 1: The constraints imposed are: (1) if $f(k) < f(k-1)$ then $f(k+x) \leq f(k)$ where $f(k)$ is the frequency estimate for thinking level k ; (2) $f(6) \geq f(7 \text{ and higher})$.

Table 6: Data and CH estimates of τ in various p-beauty contest games

subject pool or game	source ¹	group size	sample size	Nash equil'm	pred'n error	data			fit of CH model				bootstrapped 90% c.i.	
						mean	std dev	mode	τ	mean	error	std dev		mode
p=1.1	HCW (98)	7	69	200	47.9	152.1	23.7	150	0.10	151.6	-0.5	28.0	165	(0.0,0.5)
p=1.3	HCW (98)	7	71	200	50.0	150.0	25.9	150	0.00	150.4	0.5	29.4	195	(0.0,0.1)
high \$	CHW	7	14	72	11.0	61.0	8.4	55	4.90	59.4	-1.6	3.8	61	(3.4,4.9)
male	CHW	7	17	72	14.4	57.6	9.7	54	3.70	57.6	0.1	5.5	58	(1.0,4.3)
female	CHW	7	46	72	16.3	55.7	12.1	56	2.40	55.7	0.0	9.3	58	(1.6,4.9)
low \$	CHW	7	49	72	17.2	54.8	11.9	54	2.00	54.7	-0.1	11.1	56	(0.7,3.8)
.7(M+18)	Nagel (98)	7	34	42	-7.5	49.5	7.7	48	0.20	49.4	-0.1	26.4	48	(0.0,1.0)
PCC	CHC (new	2	24	0	-54.2	54.2	29.2	50	0.00	49.5	-4.7	29.5	0	(0.0,0.1)
p=0.9	HCW (98)	7	67	0	-49.4	49.4	24.3	50	0.10	49.5	0.0	27.7	45	(0.1,1.5)
PCC	CHC (new	3	24	0	-47.5	47.5	29.0	50	0.10	47.5	0.0	28.6	26	(0.1,0.8)
Caltech board	Camerer	73	73	0	-42.6	42.6	23.4	33	0.50	43.1	0.4	24.3	34	(0.1,0.9)
p=0.7	HCW (98)	7	69	0	-38.9	38.9	24.7	35	1.00	38.8	-0.2	19.8	35	(0.5,1.6)
CEOs	Camerer	20	20	0	-37.9	37.9	18.8	33	1.00	37.7	-0.1	20.2	34	(0.3,1.8)
German students	Nagel (95)	14-16	66	0	-37.2	37.2	20.0	25	1.10	36.9	-0.2	19.4	34	(0.7,1.5)
70 yr olds	Kovalchik	33	33	0	-37.0	37.0	17.5	27	1.10	36.9	-0.1	19.4	34	(0.6,1.7)
US high school	Camerer	20-32	52	0	-32.5	32.5	18.6	33	1.60	32.7	0.2	16.3	34	(1.1,2.2)
econ PhDs	Camerer	16	16	0	-27.4	27.4	18.7	N/A	2.30	27.5	0.0	13.1	21	(1.4,3.5)
1/2 mean	Nagel (98)	15-17	48	0	-26.7	26.7	19.9	25	1.50	26.5	-0.2	19.1	25	(1.1,1.9)
portfolio mgrs	Camerer	26	26	0	-24.3	24.3	16.1	22	2.80	24.4	0.1	11.4	26	(2.0,3.7)
Caltech students	Camerer	17-25	42	0	-23.0	23.0	11.1	35	3.00	23.0	0.1	10.6	24	(2.7,3.8)
newspaper	Nagel (98)	3696, 1460, 2728	7884	0	-23.0	23.0	20.2	1	3.00	23.0	0.0	10.6	24	(3.0,3.1)
Caltech	CHC (new	2	24	0	-21.7	21.7	29.9	0	0.80	22.2	0.6	31.6	0	(4.0,1.4)
Caltech	CHC (new	3	24	0	-21.5	21.5	25.7	0	1.80	21.5	0.1	18.6	26	(1.1,3.1)
game theorists	Nagel (98)	27-54	136	0	-19.1	19.1	21.8	0	3.70	19.1	0.0	9.2	16	(2.8,4.7)
									mean	1.30				
									median	1.61				

Note 1: HCW (98) is Ho, Camerer, Weigelt AER 98; CHC are new data from Camerer, Ho, and Chong; CHW is Camerer, Ho, Weigelt (unpublished); Kovalchik is unpublished data collected by Stephanie Kovalchik

Table 7: Economic Value of Various Theories

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%
Quantal Response (Game-specific λ)	418	1277	573	371	128
	9%	9%	8%	13%	8%
Cognitive Hierarchy (Common General Distribution $f(k)$)	421	1277	561	472	128
	10%	9%	6%	44%	8%
Cognitive Hierarchy (Common τ)	418	1277	573	471	128
	9%	9%	8%	43%	8%
Quantal Response (Common λ)	389	1277	561	324	128
	1%	9%	6%	-1%	8%
<u>Cross-dataset Estimation</u>					
Cognitive Hierarchy (Common τ)	418	1277	573	460	128
	9%	9%	8%	40%	8%
Quantal Response (Common λ)	379	1277	484	427	120
	-1%	9%	-9%	30%	2%
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	0.17	1.64
Game 2	0.00	0.00	0.83	1.27	1.44	2.80	0.98	2.37	0.71	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	0.57	1.10
Game 5	1.61	2.45	0.36	0.67	1.22	2.30	0.70	1.62	0.26	1.59
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	0.41	1.03

Figure 1a: Predicted Frequencies of Cognitive Hierarchy Models for Matrix Games (common t)




Figure 1b: Predicted Frequencies of Cognitive Hierarchy Models for Entry and Mixed Games (common t)




Figure 2a: Predicted Frequencies of Nash Equilibrium for Matrix Games




Figure 2b: Predicted Frequencies of Nash Equilibrium for Entry and Mixed Games




Figure A1a: Predicted Frequencies of Cognitive Hierarchy Models for Matrix Games (individual t)




Figure A1b: Predicted Frequencies of Cognitive Hierarchy Models for Entry and Mixed Games (individual t)

