

Heterogeneous Quantal Response Equilibrium and Cognitive Hierarchies¹

Colin F. Camerer, Thomas R. Palfrey, and Brian W. Rogers²

May 12, 2006

¹We are grateful to the National Science Foundation grants SES-0079301 and SES-0450712 for supporting this research.

²Camerer and Rogers: Div HSS 228-77, Caltech, Pasadena CA 91125, Palfrey: Departments of Politics and Economics, Princeton University, Princeton NJ 08544.

Abstract

We explore an equilibrium model of games where players' choice behavior is given by logit response functions, but their payoff responsiveness is heterogeneous. We extend the definition of quantal response equilibrium to this setting, calling it heterogeneous quantal response equilibrium (HQRE), and prove existence under weak conditions. We generalize HQRE to allow for limited insight, in which players can only imagine others with low responsiveness. We identify a formal connection between this new equilibrium concept, called truncated quantal response equilibrium (TQRE), and the Cognitive Hierarchy (CH) model. We show that CH can be approximated arbitrarily closely by TQRE. We report a series of experiments comparing the performance of QRE, HQRE, TQRE and CH. A surprise is that the fit of the models are quite close across a variety of matrix and dominance-solvable asymmetric information betting games. The key link is that in the QRE approaches, strategies with higher expected payoffs are chosen more often than strategies with lower expected payoff. In CH this property is not built into the model, but generally holds in the experimental data,

JEL classification numbers: 024, 026

Key words: experimental economics, game theory, cognitive hierarchy, quantal response equilibrium, bounded rationality

1 Introduction

Finding a disciplined, empirically accurate way to incorporate limits on rationality has been a central challenge in game theory. One approach,¹ “quantal response equilibrium” (QRE) maintains the assumption of equilibrium, in that beliefs are statistically accurate, but relaxes the assumption that players choose best responses.² A different approach, based on “cognitive hierarchy” (CH) theories, relaxes the equilibrium assumption, by assuming that some players do not correctly anticipate what others will do, but retains the assumption of best responding to beliefs (in some versions). QRE and CH approaches both generate statistical predictions in which every strategy is played with some probability (which obviates the need for perfection refinements), and have been used successfully to explain deviations from Nash equilibrium in many types of experiments. This paper introduces a heterogeneous form of QRE, called HQRE, and shows that special forms of HQRE are closely related to some forms of CH, thus establishing a link between the two theories.

This paper makes a theoretical contribution and an empirical contribution. The theoretical contribution is the introduction of HQRE and establishing the link between QRE and CH. In HQRE, players may have different

¹Learning models explore a different kind of rationality limit than the static models considered here (see Camerer (2004), chapter 6).

²A purer interpretation is that players do best respond, but that their expected payoffs include a disturbance term which is unobserved by the econometrician, but whose distribution is commonly known.

sensitivities λ_i to expected payoff differences across actions, parameterized by a distribution across players, $f(\lambda)$. One possibility is that $f(\lambda)$ is common knowledge. A more general possibility, called subjective HQRE, allows players with different values of λ_i to have different beliefs, $f(\lambda|\lambda_i)$, about the distribution of others' parameters. An important special case of subjective HQRE is Truncated HQRE (TQRE), in which a player i with a parameter λ_i truncates the distribution $f(\lambda)$ at an upper bound of $\theta\lambda_i$ (where θ reflects the amount of “imagination”). When $\theta = 1$, players all think that nobody is more responsive than they are, whatever their actual responsiveness. This truncated HQRE theory corresponds quite closely to some types of cognitive hierarchy theories. In particular, a limiting case of a discretized form of truncated HQRE closely approximates the CH theory of Camerer, Ho, and Chong (2004).

The empirical contribution is new experimental data from a variety of games to analyze the differential predictions of QRE, HQRE, TQRE, and CH, comparing their ability to explain our data. Surprisingly, there is very little empirical difference in fit of these approaches across the games we study. We trace this similarity to a property which is built directly into QRE, that strategies which lead to more costly deviations from optimal play are chosen less frequently. This property is not guaranteed by the CH approach but it appears to often hold in these data. Thus, there is a surprising similarity in what the QRE and CH approaches predict about the relation between strategy frequency and expected costs, even though their structures are quite

different (indeed, they are opposite in a certain sense).

While the paper is nominally about two particular limited-rationality theories, the basic questions that are addressed are fundamental. In some areas of game theory it has proved useful to introduce special preference ‘types’ (e.g., Kreps et al, 1982; Fudenberg and Maskin, 1986). The QRE and CH approaches are similar except the heterogeneity in types comes from either imperfect response or limited strategic thinking. Expanding the QRE approach to include heterogeneity creates a unified framework in which to compare these approaches and see their differences and similarity. Furthermore, introducing heterogeneity into QRE allows a concept of ‘gamesmanship’ or skill which is absent in equilibrium analysis (since, in equilibrium, all players are equally accurate and rational). Allowing variation in skill opens up new questions: Why are some people more skilled? Can skill be taught? Do skills develop with experience? These questions are not the focus of our analysis, but are naturally raised by introducing heterogeneity.

The paper proceeds as follows. Section 2 defines HQRE. Section 3 generalizes HQRE to allow for subjectivity and truncation in the beliefs of others’ types. We also introduce the CH approach, highlight its essential features, and formally describe the link between CH and TQRE. Section 4 begins by describing the experimental design. Further subsections report the experimental data and contain an empirical analysis of the fit of the models across various games. Finally, Section 5 concludes.

2 HQRE

We explore a logit QRE model where players' choice behavior follows logit quantal response functions but there is heterogeneity with respect to the responsiveness parameter.³ We now present the model for games in strategic form, following the approach of McKelvey and Palfrey (1995).

Let $\Gamma = [N, \{A_i\}_{i=1}^n, \{u_i\}_{i=1}^n]$ be a game in strategic form, where $N = \{1, \dots, n\}$ is the *set of players*, $A_i = \{a_{i1}, \dots, a_{iJ_i}\}$ is *i's action set* and $u_i : A \rightarrow \mathbb{R}$ is *i's payoff function*, where $A = A_1 \times \dots \times A_n$. Let ΔA_i denote the set of probability distributions over A_i and let $\Delta A = \Delta A_1 \times \dots \times \Delta A_n$ denote the product set of probability distributions over A_i , $i = 1, \dots, n$. If $\alpha \in \Delta A$, then player *i's* expected payoff is denoted by:

$$U_i(\alpha) = \sum_{a \in A} (\prod_{k=1}^n \alpha_k(a_k)) u_i(a).$$

and we denote the expected payoff to player *i* from using action $a_{ij} \in A_i$ by:

$$U_{ij}(\alpha) = \sum_{a_{-i} \in A_{-i}} (\prod_{k \neq i} \alpha_k(a_k)) u_i(a_{ij}, a_{-i}).$$

Each player is independently assigned by nature a response sensitivity, λ_i , drawn from a fixed distribution, $F_i(\lambda_i)$, with smooth density function, f_i , full support on $[0, \infty)$ and finite moments – for example, f_i could be the density function for an exponential or log normal distribution. We call λ_i *i's type*.

³Our approach and main theoretical results would extend to the general framework of regular quantal response equilibrium studied by Goeree, Holt and Palfrey (2005).

Quantal response functions are logit transformations of expected payoffs, so if i has type λ_i and the actions have expected payoffs $U_i = (U_{i1}, \dots, U_{iJ_i})$, then the probability of choosing action j is:

$$p_{ij}(\lambda_i) = \frac{e^{\lambda_i U_{ij}}}{\sum_{k=1}^{J_i} e^{\lambda_i U_{ik}}} \quad (1)$$

We call any measurable function $p_i : [0, \infty) \rightarrow \Delta A_i$ a *strategy* for player i .

The assumption in HQRE is that $F_i(\lambda_i)$ is common knowledge, but i 's type, λ_i , is private information known only to i . Given some fixed profile of expected payoffs to i , $U_i = (U_{i1}, \dots, U_{iJ_i})$, equation (1) implies a choice probability function that depends on λ_i , which we denote by $p_i(\lambda_i) = [p_{i1}(\lambda_i), \dots, p_{iJ_i}(\lambda_i)]$. Therefore, given i 's profile of choice probability functions, $p_i(\cdot)$, the ex ante probability i chooses action j (i.e., before λ_i is drawn) is:

$$\sigma_{ij}(p) = \int_0^\infty p_{ij}(\lambda) f_i(\lambda) d\lambda \quad (2)$$

Following Harsanyi (1973), we call $\sigma_i = (\sigma_{i1}, \dots, \sigma_{iJ_i})$ i 's *induced mixed strategy*. Given σ_{-i} , the induced mixed strategy profile of all players other than i , i 's expected payoffs, $U_i(\sigma_{-i}) = (U_{i1}, \dots, U_{iJ_i})$, can be expressed as:

$$U_{ij}(\sigma) = \sum_{a_{-i} \in A_{-i}} \left(\prod_{k \neq i}^n \sigma_k(a_k) \right) u_i(a_{ij}, a_{-i}). \quad (3)$$

In a heterogeneous quantal response equilibrium with logit response functions, equations (1), (2), and (3) must all be satisfied simultaneously. This

leads to the following

Definition 1 p^* is a *Heterogeneous Logit Equilibrium* if:

$$p_{ij}^*(\lambda_i) = \frac{e^{\lambda_i U_{ij}(\sigma(p^*))}}{\sum_{k=1}^{J_i} e^{\lambda_i U_{ik}(\sigma(p^*))}} \text{ for all } i = 1, \dots, n, j = 1, \dots, J^i \text{ and } \lambda_i \in [0, \infty).$$

This captures the idea that in HQRE players have rational expectations about the distribution of mixed strategies, and these will then be self-fulfilling given the commonly known distribution of profiles of quantal response functions. Therefore, like Nash equilibrium, the solution to the problem is a fixed point of a mapping from choice probabilities to choice probabilities. The Appendix proves existence of HQRE for the logit case, using a fixed point theorem. Note that if each $F_i(\lambda_i)$ has a common single mass point,⁴ then HQRE is the same as QRE with a common response parameter λ for all players.

Theorem 1. *In finite games, a Heterogeneous Logit Equilibrium exists.*

Proof: See appendix.

3 Subjective HQRE

In this section, we consider a more general model which allows expectations about choice probabilities to be inconsistent with the actual choice frequen-

⁴While the assumptions of the F_i above preclude this case, it can be approximated arbitrarily closely by F_i that do satisfy the assumptions.

cies of the other players. Models with this property could prove useful in explaining behavior in one-shot games, or complex games in which learning or other forces have not enabled beliefs to fully equilibrate to actual choices. However, the particular form of inconsistencies allowed in subjective HQRE still permit it to be thought of as an equilibrium model: choice probabilities conditional on type are common knowledge; it is only the perceived distribution of types that varies across players.

We replace the rational expectation assumptions by an assumption of *subjective expectations*. According to this model, the equilibrium strategies⁵ of all players are common knowledge in equilibrium, but players have different beliefs about the type distributions. Denote the conditional subjective beliefs of player i about the type of player k by $F_k^i(\lambda_k|\lambda_i)$. Note that beliefs generally depend on a player's own type.⁶ As we show later in the section, this provides a framework for linking HQRE approaches with cognitive hierarchy approaches, which share a similar feature of belief heterogeneity. This difference in beliefs results in equilibrium strategies (and induced mixed strategies) that in general are different from those in a heterogeneous logit equilibrium.

The notation is similar to that used earlier but the induced mixed strategies are more complicated. Recall that the role of induced mixed strategies is to compute U . That is, induced mixed strategies represent the beliefs

⁵Recall that strategies map types into choice probabilities.

⁶If $F_k^i(\lambda_k|\lambda_i) = F_k(\lambda_k)$ for all $i, k, \lambda_i, \lambda_k$ then subjective HQRE is the same as HQRE.

players other than i have about i 's action choice, without knowing i 's type. Under subjective HQRE, players do not share a common prior about F and therefore do not share identical beliefs about action choices.

As before, for any subjective belief about action profiles, $\hat{\sigma} \in \Delta A$, player i 's expected payoff is given by:

$$U_i(\hat{\sigma}) = \sum_{a \in A} (\Pi_{k=1}^n \hat{\sigma}_k(a_k)) u_i(a).$$

and the (subjective) expected payoff to player i from using action $a_{ij} \in A_i$ is:

$$U_{ij}(\hat{\sigma}) = \sum_{a_{-i} \in A_{-i}} (\Pi_{k \neq i} \hat{\sigma}_k(a_k)) u_i(a_{ij}, a_{-i}).$$

With logit response functions, if i has type λ_i and the actions by i have expected payoffs $U_i = (U_{i1}, \dots, U_{iJ_i})$, then the probability of i choosing action j as a function of λ_i is:

$$p_{ij}(\lambda_i; U_i) = \frac{e^{\lambda_i U_{ij}}}{\sum_{k=1}^{J_i} e^{\lambda_i U_{ik}}} \quad (4)$$

We call any measurable function $p_i : [0, \infty) \rightarrow \Delta A_i$ a strategy for player i . Hence given some fixed vector of expected payoffs to i , $U_i = (U_{i1}, \dots, U_{iJ_i})$, equation (4) implies an induced mixed strategy for i that depends on λ_i : $p_i(\lambda_i) = [p_{i1}(\lambda_i), \dots, p_{iJ_i}(\lambda_i)]$.

We next turn to the induced mixed strategies. Because of the different subjective beliefs about the distribution of λ , players k and k' can have

different beliefs about the induced mixed strategy of player i . However, we assume that any differences in their beliefs about i 's mixed strategy are due to differences in beliefs about the distribution of λ_i . That is, the strategy profile, p , is assumed to be common knowledge (hence we refer to this as an equilibrium model). We denote type λ_k of player k 's belief about player i 's induced mixed strategy by $\sigma_i^k(p_i)$. Therefore, given i 's strategy, $p_i(\cdot)$, the belief of player k that player i will choose action j (i.e., before λ_i is drawn) is:

$$\sigma_{ij}^k(p_i|\lambda_k) = \int_0^\infty p_{ij}(\lambda_i) f_i^k(\lambda_i|\lambda_k) d\lambda_i \quad (5)$$

Given $\sigma_{-i}^i(p_{-i}|\lambda_i)$, the beliefs of type λ_i of player i about the induced mixed strategy profile of all players other than i , type λ_i of player i 's expected payoffs, $U_i^{\lambda_i}(\sigma_{-i}^i) = (U_{i1}^{\lambda_i}, \dots, U_{iJ_i}^{\lambda_i})$, are simply:

$$U_{ij}^{\lambda_i}(\sigma_{-i}^i) = \sum_{a_{-i} \in A_{-i}} \left(\prod_{k \neq i}^n \sigma_k^i(a_k|\lambda_i) \right) u_i(a_{ij}, a_{-i}). \quad (6)$$

In a subjective HQRE with logit response functions, equations (4), (5), and (6) must all be satisfied simultaneously. This leads to the following

Definition 2 p^* is a *Subjective Heterogeneous Logit Equilibrium* if:

$$p_{ij}^*(\lambda_i) = \frac{e^{\lambda_i U_{ij}^{\lambda_i}(\sigma_k^i(p^*|\lambda_i))}}{\sum_{k=1}^{J_i} e^{\lambda_i U_{ik}^{\lambda_i}(\sigma_k^i(p^*|\lambda_i))}} \text{ for all } i = 1, \dots, n, j = 1, \dots, J_i \text{ and } \lambda_i \in [0, \infty).$$

This definition reflects the idea that in subjective HQRE players have rational expectations about strategies (that is, a player's behavior conditional

on his type λ), but may have different beliefs about the distribution of mixed strategies, which are induced by different beliefs about the distribution of types λ .

3.1 Truncated expectations and bounded imagination

Since subjective HQRE is quite general, precision in applying it must come from additional restrictions on heterogeneity and subjective beliefs (preferably empirically-plausible ones).⁷ We do this by introducing “truncated expectations:” Players act as if they are not aware of the existence of types who are more rational than some maximum upper bound, and this upper bound may depend on their own type. Given their truncated beliefs, they form expectations by integrating over their perceived type distribution, just as in HQRE.

One way to operationalize subjective HQRE is to assume that there is an upper bound on player i ’s imagined types of $\theta_i(\lambda_i)$, where $\theta_i(\lambda_i)$ is commonly known. We assume that $\theta_i(\lambda_i)$ is uniformly continuous in λ_i and for each i there exists $\bar{\theta}_i$ such that $\theta_i(\lambda_i) \leq \bar{\theta}_i \lambda_i$ for all λ_i .⁸ We assume a modified form of rational expectations, which we call *truncated rational expectations*. The beliefs of type λ_i of player i about λ_{-i} are rooted in the

⁷Earlier papers have considered variations of subjective HQRE. McKelvey, Palfrey, and Weber (2000) consider an HQRE model of self-centered subjective beliefs where players have different λ_i ’s and believe every one else is exactly like them. Weizsacker (2003) considers a more general model, where the players still have point beliefs, but these beliefs are not necessarily self-centered.

⁸This can be generalized. For example, player i could have private information about $\bar{\theta}_i$, or there could be lower as well as upper bounds.

“true” distribution, but normalized to reflect the missing density: That is, for $\lambda_i > 0$, the subjective beliefs of i about the type of player k is given by $F_k^i(\lambda_k|\lambda_i) = F_k(\lambda_k)/F_k(\theta_i(\lambda_i))$ for $\lambda_k \in [0, \theta_i(\lambda_i)]$ and $F_k^i(\lambda_k|\lambda_i) = 1$ for $\lambda_k \geq \theta_i(\lambda_i)$. This is truncated HQRE, or *TQRE*. Note that as $\theta_i \rightarrow \infty$ for all i , the upper bound on λ is lifted and the model converges to the standard HQRE model.

The truncation, $\theta_i(\lambda_i)$ can be interpreted as player i ’s *imagination*. Since $\theta_i(\lambda_i)$ is finite, this is a model of bounded imagination, in the sense that for any type λ_i of player i , all λ_{-i} – *types* beyond a certain threshold, $\theta_i(\lambda_i)$, are unimaginable in the sense that i assigns zero probability to all those higher types. Notice that if $\theta_i(\lambda_i)$ is increasing, then players who are “better” in the sense of payoff responsiveness (i.e. higher λ_i) necessarily also have more accurate expectations, in the sense that their beliefs are closer to the true distribution F . Types for which $\theta_i(\lambda_i) \approx 0$ are almost completely unimaginative in the sense that they believe all other players are nearly random. Hence these very low types will act approximately as if they are applying the principle of insufficient reason to form expectations about the other players’ strategy choices (as do the level-1 types in the cognitive hierarchy model), and then quantal respond to these beliefs. If $\theta_i(\lambda_i) \leq \lambda_i$, then we say that players are *self-limited*, because they cannot imagine types with higher λ than their own. Proving existence of TQRE requires a slightly different proof than HQRE because different λ -types have different beliefs about the other players.

Theorem 2. *In finite games, a Truncated Heterogeneous Logit Equilibrium exists.*

Proof: See appendix.

There are a number of reasons why truncated beliefs represent a reasonable manner of introducing belief heterogeneity. One rationale is that players with a low value of λ who can imagine players with higher λ , and compute what those other players will do, will generally want to switch to the higher-type behavior. Another rationale is the large body of evidence showing that people are often overconfident about their relative skill and prospects, compared to other people.⁹ A third rationale is computational complexity: If there are cognitive costs to computing expected payoffs, those costs increase as players have more other types to consider. The benefits from more imagination—the expected payoff differential from imagining what a wider range of types will do—are likely to fall as λ rises, so the truncated expectations assumption can be seen as a reduced-form model of cost-benefit calculations which lead players to ignore information that is hard to process and not too costly to ignore.

⁹Kahneman and Tversky (1972) first studied overconfidence, and much work has followed, e.g. Camerer and Lovo (1999) and Santos-Pinto and Sobel (2005).

3.2 Discretized TQRE: The connection between QRE and CH

In this section we establish a formal equivalence between a version of TQRE and CH.

3.2.1 Truncation and Heterogeneity in CH

CH introduces heterogeneity of player types of a much different kind than HQRE. In CH there is a discrete distribution $f(k)$ of players who do k steps of thinking, so k indexes *strategic sophistication*. The choice probabilities for a k -step player i choosing strategy j are $p_{ij}(k)$. A 0-step player randomizes over her (finite) number of strategies J_i , so $p_{ij}(k) = 1/J_i \forall j$. Note that these players do not form beliefs or even attend to their payoffs; their presence is just assumed to get a hierarchical process started in a simple way.

Truncation of beliefs in a similar way to TQRE (albeit relative to beliefs about the distribution of a much different parameter) is the central feature of the cognitive hierarchy (CH) model of Camerer, Ho and Chong (2004). Players who do $k \geq 1$ steps of thinking form truncated beliefs about the fraction of h -step types according to $g_k(h) = f(h) / \sum_{n=0}^h f(n) \forall h < k$ and $g_k(h) = 0 \forall h \geq k$. In this specification, players do not imagine that any others are at their level (or higher), so, in the notation of the TQRE, they effectively have $\theta < 1$. All positive-step thinkers best respond given their beliefs, so in a two-

player game,¹⁰ $p_{ij}(k) = 1$ iff $a_{ij} = \operatorname{argmax}_a \sum_{h=0}^{k-1} g_k(h) \sum_{m=1}^{J-i} p_{im}(h) u_i(a, a_{-im})$.¹¹

The expected choice probabilities for player i implied by the CH model are given by $p_{ij} = \sum_{k=0}^{\infty} p_{ij}(k) f(k)$.

For precision, Camerer, Ho and Chong (2004) assume $f(k)$ is Poisson and estimate the mean of the distribution using data from more than 100 normal-form games. Other types of hierarchical models have been explored as well. Nagel (1995) and Stahl and Wilson (1994) were the first to use strategic hierarchies to study dominance-solvable “beauty contest” games and matrix games, respectively. In Nagel’s approach k –step players think all others do $k - 1$ steps of reasoning (i.e., $g_k(h) = I(h, k - 1)$ where $I(x, y)$ is an identity function equalling one if $x = y$ and zero otherwise). Stahl and Wilson’s limited-step types have the same one-step-below beliefs as in Nagel, but they also permit equilibrium types and “worldly” types who maximize against the empirical distribution of play. Players in these models are typically modelled as using quantal responses instead of best responses.¹²

¹⁰The expressions are more cumbersome to write out with n -player games because the probabilities of other players’ types have a multinomial distribution with many terms. Roughly speaking, CH models become hard to compute as the number of players increases, while QRE models, which require finding a fixed point, become more difficult to compute as the number of strategies increases.

¹¹If more than one action is a best response they are assumed to randomize equally across all best responses.

¹²Recent applications of this approach include Costa-Gomes and Crawford (2005) and Crawford and Iribiri (2004, 2005).

3.2.2 Differences and Similarities between CH and TQRE

The general form of TQRE is different from CH in three distinct ways. First, the maximum “imagined” type of other players could be equal to, greater than, or less than a player’s actual type (depending on θ_i), and this could be a second source of heterogeneity, whereas in all the CH and related approaches the imagination parameter for all players is strictly less than 1.¹³ Second, levels of rationality are indexed by λ in TQRE, rather than k , so that types correspond to increasing payoff responsiveness rather than strategic sophistication. Third, in TQRE, all types exhibit some degree of randomness in response, reflecting the stochastic choice modelling. In CH all players with $k \geq 1$ best-respond, so the only source of stochastic choice behavior is buried in the 0-level types.

In spite of these major differences between the two approaches, there are a number of important similarities between the TQRE and CH approaches. First, both models have heterogeneity of types. Second both models incorporate stochastic behavior. Third, they share an important type in common: the bottom of the food chain ($k = 0$ or $\lambda = 0$); and these lowest types are in the support of the beliefs of *all* types. Fourth, both models assume there is a limit to the rationality of the other players, and this limit is monotonically increasing in type. Fifth, in both approaches, there is heterogeneity of be-

¹³The Stahl-Wilson (1995) and Costa-Gomes and Crawford (2005) specifications include other types that do not correspond to levels in the thinking hierarchy. If the maximum imagined type is always less than one’s own type, then the model’s solution can be computed recursively, as in CH. However, if $\theta = 1$, so that players are aware that others share their level of thinking, the model must be solved using fixed point methods.

liefs as well as heterogeneity of types, and these are correlated: higher types have more accurate beliefs, and these converge to rational expectations about $f(\lambda)$ (or $f(k)$) as λ (or k) increases. Finally, *all* players are overconfident in the sense that they underestimate the gamesmanship (be it sophistication or responsiveness) of the other players.

3.2.3 The formal connection between TQRE and CH

In this section we show that by placing two parametric restrictions on TQRE, for any CH model there exist distributions of types in TQRE that lead to behavioral predictions that are essentially equivalent to CH. By *essentially* equivalent, we mean two things. First we mean that the equivalence is in terms of approximations that can be made arbitrarily close; second, the approximating equilibria in TQRE are unique.

To make this approximation, we first consider distributions such that the set of λ values is discrete, $L_\gamma = \{0, \gamma, 2\gamma, \dots, k\gamma, \dots\}$, with grid size γ . Discretizing the distribution is a small step in practice, because it is usually done in applications to make numerical computations. A player of type k , is called a *level k* player, and has response parameter $\lambda = k\gamma$. We fix the distribution over k , so that the probabilities of types are $f = \{f(0), f(1), \dots, f(k), \dots\}$. This is simply a discretized HQRE. The HQRE is defined exactly as before, except for the discrete, distribution of valuations, and existence is easily established.

The first parametric restriction we place on TQRE is that $\theta_i(\lambda) = \frac{k}{k+1}\lambda$ for all $i, \lambda \in L_\gamma$. That is, as in CH, players only recognize lower types, but

otherwise have correct beliefs about the distribution, that is, they correctly estimate the relative proportions of lower type players. In this version of TQRE, level 0 players randomize uniformly, for any value of γ . Level 1 players quantal respond using $\lambda = \gamma \cdot 1$, assuming all other players are type 0. Level 2 players quantal respond (using $\lambda = \gamma \cdot 2$), assuming all other players are type 0 or type 1, with perceived probabilities $\frac{f(0)}{f(0)+f(1)}$ and $\frac{f(1)}{f(0)+f(1)}$, respectively. Higher-level types are defined iteratively in the obvious way. This specification is also only slightly different from a discrete heterogeneous version of QRE – the only difference being truncated vs. untruncated rational expectations.

In contrast, however, as the grid size γ grows ($\gamma \rightarrow \infty$), then players doing one or more steps of thinking have unboundedly large values of λ , so their choices approach best responses, even for low-level (other than zero) players. This special form of discretized TQRE converges to a generalized form of CH in which the type probabilities have the probability distribution $\{f(0), f(1), \dots, f(k), \dots\}$. A second parametric assumption makes this form of TQRE identical to CH as it is generally implemented with a Poisson distribution of types. That is, assume $f(k)$ follows a Poisson distribution, that is $f(k) = \frac{\tau^k}{k!} e^{-\tau}$. The important link is that TQRE with $\theta \lesssim 1$ retains the continuous types, stochastic choice, and equilibrium elements of *HQRE*, but it borrows the downward-looking and Poisson elements of *CH*.

The formal connection between TQRE and CH is asymptotic in γ . In particular, for almost all games and almost all values of τ , the aggregate

choice probabilities implied by the γ -TQRE model converge to the aggregate choice probabilities of CH. This can be stated more formally, as follows.

Fix τ . Denote the CH choice probability that level k of player i chooses action j by p_{ijk}^τ , and denote the γ -TQRE choice probability (and f distributed Poisson with parameter τ) that type $\lambda = \gamma k$ of player i chooses action j by p_{ijk}^γ . Denote the expected CH choice probability of player i choosing action j by $\bar{p}_{ij}^\tau = \sum_{k=0}^{\infty} p_{ijk}^\tau f(k)$ and the expected γ -TQRE choice probability of player i choosing action j by $\bar{p}_{ij}^\gamma = \sum_{k=0}^{\infty} p_{ijk}^\gamma f(k)$. Denote $\Delta^{\tau,\gamma} = \sum_{i=1}^n \sum_{j=1}^{J^i} (\bar{p}_{ij}^\tau - \bar{p}_{ij}^\gamma)^2$.

Theorem 3: Fix τ . For almost all finite games Γ and for any $\varepsilon > 0$, there exists $\bar{\gamma}$ such that $\Delta^{\tau,\gamma} < \varepsilon$ for all $\gamma > \bar{\gamma}$.

Proof: See appendix.

Hence the *limiting* discretized TQRE shares the features of best response, and hierarchical beliefs that characterize CH.

Thus, we have created a “family tree” surrounding HQRE, illustrated in Figure 1. The most general model is subjective HQRE. When all subjectivity takes the form of truncation at a player’s $\theta\lambda_i$, we have TQRE. From TQRE there are two branches to follow. If we send $\theta \rightarrow \infty$, then the subjectivity vanishes, and we have HQRE. From there, a limiting distribution that places all mass at one value of λ corresponds to standard QRE. Following the other branch from TQRE corresponds to discretizing TQRE so that λ takes on a countable set of values, $L_\gamma = \{0, \gamma, 2\gamma, \dots, k\gamma, \dots\}$. Sending $\gamma \rightarrow \infty$, and




Figure 1. All of the models considered are special or limiting cases of subjective HQRE. The relationships among the models are depicted.

assuming a Poisson distribution on $F(k)$ then yields the standard CH model.

Another interesting special case of TQRE is when $\theta = 1$ and $\gamma \rightarrow \infty$. The restriction $\theta = 1$ means that 1-step ($k = 1$) players are best-responding to a mixture of choices by their own types and some random (0-step, $\lambda = 0$) types. Under these restrictions, in games with strict Nash equilibria, if $F(0)$ is small enough (there are too few random types to induce the QRE types away from the Nash strategies), and $\gamma \rightarrow \infty$ (the QRE types best respond), the model is a “noisy Nash” model which has been used in previous applications as a benchmark that illuminates the empirical importance of quantal response.¹⁴

4 Experimental evidence

4.1 Games and design

We explored the fit of different HQRE and CH models in 17 complete-information normal form games, and one game with information asymmetry (discussed in Section 4.4 below). Table A1 in Appendix 2 presents the payoff matrices of all 17 games and the relative choice frequencies from our data. The data from the row and column roles are combined in the symmetric games.

One game is an “unprofitable” game (Morgan and Sefton, 2002) in which maximin strategies do not form an equilibrium, yet guarantee the same pay-

¹⁴See, for example, McKelvey and Palfrey (1992), El-Gamal, McKelvey and Palfrey (1993), and Fey, McKelvey and Palfrey (1996)

offs as equilibrium strategies. Twelve games are affine transformations of games created by Stahl and Wilson (1995) (SW) to fit models of iterated strategic thinking (which are precursors to CH that include more types). We changed some design details about how the games were presented, to see how robust the patterns of play were to such details, and to avoid focal points.¹⁵

These games were chosen because there is a high proportion of Nash play in some of the games in the original SW sample, but the CH model cannot fit those data because the Nash strategy is not reached by iterations of thinking steps with best response (see Camerer, Ho and Chong, 2004). These games are interesting to study since one of our goals is to identify strategic aspects in which some models make better predictions than others. Game 8 from SW is a good example.

Table 1 shows the payoff matrix of our game based on SW 8. The three columns following the payoff matrix list the empirical choice frequencies, and the predictions of QRE and CH, based on their fitted parameter values. We find an optimal $\lambda^* = 1.05$, which generates predictions that are close to the observed play. In contrast, the best fitting parameter for CH is $\tau^* = 0.0$, with corresponding uniformly random behavior. That is, no other parameteriza-

¹⁵The main difference is the payoff transformation. This was done to eliminate possible focal payoffs, such as 0 and 100, that appear in the original SW games. Instead, our payoffs are scaled so that all entries are two-digit numbers. We also included 4 games that were neither symmetric nor 3x3. Another difference is the matching protocol. We implemented a standard random matching procedure, whereas SW match each choice against the empirical distribution of others' choices. Also, we paid subjects exactly according to the payoff tables instead of using the lottery procedure of SW. Finally, our games were presented sequentially, without the possibility of changing choices in previous games, whereas SW allowed subjects to revise all decisions before submitting their choices.

tion of CH fits better than random choice, and the model is not consistent with the relative choice frequencies in our data in this sense. As stated, the reason for this relates to the fact that equilibrium strategies can not be reached through a process of iterated thinking, whereas the empirical choice probabilities show a strong tendency towards equilibrium, as evidenced by the relatively large value of λ that we estimate in QRE. Recall that QRE converges to a Nash equilibrium as λ increases and players become completely payoff responsive.

Four games involve “cloning”—presenting the same pure strategy more than once. These games are included because QRE and CH models can respond differently to the addition of cloned strategies. It is well-known that in stochastic choice models, splitting a single strategy into two equivalent strategies increases the predicted probability of play (the two split strategy frequencies are generally higher than the single strategy frequency) unless some hierarchical structure is imposed.¹⁶ This property can lead to different predictions in QRE and CH approaches, since a cloned strategy does not necessarily receive more weight in CH (except for 0-level players) as players best respond, rather than quantal respond, in CH.

¹⁶In multinomial logit modelling this property is called the “red bus, blue bus” problem. This term comes from early transportation applications predicting whether commuters would drive or take a bus to work. The choice between {drive,bus} and {drive,red bus, blue bus} can be different if choice is stochastic. For example, if people choose randomly then there is a $\frac{1}{2}$ probability they will take the bus in the first choice set and a $\frac{2}{3}$ of taking the bus in the second choice set. A large literature on hierarchical models with nesting has emerged to take care of this problem, by treating the choice between {drive,bus} as a top-level choice and the choice between {red bus, blue bus}, conditional on choosing *bus*, as a second-level choice (where $P(bus) = P(redbus) + P(bluebus)$).

S-W 8						
	A	B	C	Data	QRE	CH
A	11, 11	59, 91	51, 51	0.17	0.11	0.33
B	91, 59	27, 27	51, 43	0.20	0.25	0.33
C	51, 51	43, 51	53, 53	0.63	0.64	0.33
					$\lambda = 1.05$	$\tau = 0.0$

Table 1. Game 8 from SW, along with empirical choice frequencies and the optimal predictions of QRE and CH.

Cloned Matching Pennies						
	L	R	R	Data	QRE	CH
U	50, 10	10, 20	10, 20	0.78	0.51	0.73
D	10, 20	20, 10	20, 10	0.10	0.24	0.14
D	10, 20	20, 10	20, 10	0.12	0.24	0.14
Data	0.71	0.22	0.07			
QRE	0.33	0.34	0.34	$\lambda = 0.12$		
CH	0.50	0.25	0.25	$\tau = 0.90$		

Table 2. The matching pennies game where “bottom” is cloned for the row player and “right” is cloned for the column player, along with empirical choice frequencies and the optimal predictions of QRE and CH.

One of our games with cloned strategies is asymmetric matching pennies, where “down” is cloned for the row player and “right” is cloned for the column player, creating a 3×3 game. The payoff matrix is given in Table 2, along with observed choice frequencies, and the predictions from QRE and CH calculated at the best fitting parameter values. Notice the reversal in prediction quality of the two models relative to SW 8. QRE consistently overestimates the frequency of cloned strategy play.¹⁷ In addition, the data show too much “up” and “left” play relative to Nash equilibrium, a phenomenon that the CH model does a better job of accounting for, due partly to the fact that these strategies are best responses to the uniform play of level 0’s.

Another game with cloned strategies is the “Joker game” of O’Neill (1987), which was originally designed to allow a clean test of minimax play. The payoff matrix is depicted in the lower right of Table A1, where the first strategy (the joker card) has been cloned for the row player. Notice that the row player’s frequency of the joker strategy is 14%, and the column player’s choice frequency of the joker strategy is 38%, both below the Nash equilibrium probability of 40%, which is predicted by QRE. The predicted change from Nash to CH depends on the value of τ . At the pooled maximum likelihood estimate, these probabilities are 23% and 32%, respectively. The empirical frequencies are also lower than what was observed in the original

¹⁷In this game, CH also overestimates the amount of cloned play, but it is not significant, and it is a much smaller magnitude than the error of QRE.

un-cloned O'Neill experiment (36% and 43%, respectively), where both QRE and CH correctly predict the column player's choice frequency to be *above* the Nash equilibrium level of 40%.

Before proceeding to a more comprehensive analysis of our data, we briefly summarize the design features. The experimental sessions took place at Caltech and UCLA in April, 2004. There were four sessions, each consisting of 25 rounds of the betting game and one shot each of the 17 matrix games, with each ordering of the two parts done twice. The sessions had 16, 18, 20, and 20 subjects each, resulting in a total of 1210 observations in the matrix games¹⁸ and 1850 observations of the betting game. Subjects consisted of undergraduate students in the two institutions, and were randomly selected to participate in the experiments. Upon arrival, students were seated at random locations in the lab. Once in the lab, instructions were read aloud for everyone to hear, and all subsequent interactions took place only through the computers. During the phase of 17 matrix games, subjects were randomly and anonymously rematched after each decision, and the same procedure was used during the 25 repetitions of the betting game in order to minimize possible repeated game effects. Average payoffs were \$7.50 for the matrix games and \$8.95 for the betting game, resulting in an average total payoff of \$21.45 after including a \$5.00 showup fee. Sessions lasted approximately 2 hours.

¹⁸Due to technical problems, one session is missing data from games 3, 4 and 17, resulting in a reduction of $3 \cdot 16 = 48$ observations.

4.2 Complete information games

The first focus of our analysis is on estimation of the QRE, HQRE (assuming an exponential $f(\lambda)$), discretized TQRE (assuming a Poisson distribution of levels), and Poisson-CH models for the complete-information games. All four models are estimated separately for each normal-form game, as well as pooling data across games and estimating a single set of parameters for all games.

Table 3 summarizes the estimation results for the complete-information games. Each column lists the best-fitting parameter value and negative log likelihood for a particular model. The parameterizations are as follows: QRE has a single λ^* (i.e., this is HQRE with a single mass point at λ^*); Poisson-CH has a mean number of thinking steps τ ; HQRE has an exponential distribution of response sensitivities with cumulative distribution function $F(\lambda) = 1 - \exp(-\alpha\lambda)$, so it has a single parameter α to estimate; and TQRE is discretized with grid size γ and Poisson parameter τ (i.e., the probability that $\lambda_i = k\gamma$ is $f(k) = \frac{\tau^k e^{-\tau}}{k!}$, and a level k player has beliefs truncated at $k\gamma$).¹⁹

We use maximum likelihood to estimate the parameters of the models and assess their qualities of fit. Each model makes a unique statistical prediction in each game as a function of its parameter(s). At the aggregate level, our

¹⁹We also tried an HQRE model with a lognormal distribution of $f(k)$, parameterized by Gaussian mean μ and variance σ^2 , so that the mean and variance of the distribution are controlled by separate variables. This model fit slightly better than exponential-HQRE in SW games 5-6, 9-10, 13-14, and 16, (perhaps because it has two free parameters rather than one), but the numerical computation did not converge for the other games.

data consists of group-level choice counts for each strategy of each game. Denote the empirical choice count of strategy j for player i in game g by c_{ijg} , and denote model M 's prediction of the frequency at parameter value ρ by $f_{ijg}^M(\rho)$. We can express the log likelihood of model M as a function of ρ by

$$\ln L^M(\rho) = \sum_{g=1}^{17} \sum_{i \in N^g} \sum_{j=1}^{J_i^g} c_{ijg} \ln f_{ijg}^M(\rho). \quad (7)$$

Maximizing $L^M(\rho)$ allows us to estimate the parameter(s) for each model. The results from this exercise appear near the bottom of Table 3, in the row marked “Pooled,” to indicate that a single parameter is estimated across all games. In addition, we estimate the parameters separately for each of the 17 games, simply by taking the likelihood function to consist only of the terms corresponding to strategies from a particular game. That is,

$$\ln L^M(\rho, g) = \sum_{i \in N^g} \sum_{j=1}^{J_i^g} c_{ijg} \ln f_{ijg}^M(\rho). \quad (8)$$

These results occupy the bulk of the table.

The first two columns of Table 3 are important for assessing the fits of the models. The “random” log likelihood is the the likelihood that results from a model that assumes every player randomizes uniformly in every game, so that the choice frequency for player i in game g are simply $1/J_i^g$. This number represents a lower bound on the quality of fit (recall all models include uniformly random behavior as a special case). The “empirical” log

likelihood results from a (hindsight) model that assigns to every strategy its empirical frequency, that is, $f_{ijg}^e = \frac{c_{ijg}}{\sum_{j' \leq j_i^g} c_{ij'g}}$. This is the model that results in the best possible fit, and is therefore an upper bound on the quality of fit for the models we consider. Thus, the random and empirical likelihoods bracket the fitted likelihoods for all the models we consider.

In Camerer, Ho and Chong (2002) a more general form of $f(k)$ is estimated which allows each $f(k)$, for $0 \leq k \leq 6$, to be a free parameter. For identification purposes, they impose the condition that the density $f(k)$ must be single-peaked. This more general form fits only very slightly better than the Poisson distribution in the many games they estimate, suggesting that a single-peaked distribution $f(k)$ such as Poisson—which is parsimoniously characterized by only one parameter, the mean and variance τ —is a good reduced-form placeholder for more general single-peaked distributions, in fitting experimental data. We adopt this approach in our estimation of CH, as well as HQRE and TQRE.²⁰

There are two central features of Table 3. First, while the parameter values for each model are reasonably similar across games, all models have one or more games with outlying values. For example, two QRE λ estimates are infinite (Nash equilibrium) and four CH τ values are either zero (corresponding to random choice) or implausibly high (above 10).²¹ At the same time, the fit

²⁰For HQRE, we use an exponential distribution.

²¹Because the CH model has best response (except for 0-step types), the likelihood surface is potentially discontinuous and very jagged. Small changes in τ will change the pattern of choices for many step-k types which produces large changes in predicted aggregate frequencies, and hence in log likelihood for any particular set of data. As a

Matrix Game Estimates												
Game #	Game	Random neg Log L	Empirical neg Log L	QRE λ neg Log L	CH τ neg Log L	HQRE α neg Log L	γ	TQRE τ neg Log L				
1	Unprofitable	81.30	73.12	0.06	75.97	10.19	73.65	0.07	75.93	0.35	3.73	73.12
2	Cloned MP	81.30	52.34	0.12	71.72	0.90	58.39	0.13	71.76	∞	0.90	58.39
3	Cloned SH (Low)	51.96	46.50	0.15	46.74	0.44	46.72	0.15	46.73	0.50	0.62	46.67
4	Cloned SH (Hi)	51.96	40.18	0.18	41.88	1.00	41.91	0.20	41.88	0.89	1.40	40.65
5	SW 1	81.30	39.55	0.13	41.51	3.10	40.92	0.36	39.56	0.17	2.66	39.93
6	SW 2	81.30	46.00	0.13	46.75	5.59	46.03	0.25	47.04	∞	5.59	46.03
7	SW 3	81.30	53.26	0.07	73.94	0.72	63.26	0.10	73.03	∞	0.72	63.26
8	SW 4	81.30	55.02	∞	56.40	1.96	56.11	∞	56.40	∞	1.96	56.11
9	SW 5	81.30	79.27	0.14	80.50	0.00	81.30	0.17	80.25	0.00	0.00	81.30
10	SW 6	81.30	79.54	0.05	80.86	1.27	79.69	0.06	80.83	0.11	2.81	79.54
11	SW 7	81.30	73.69	∞	73.73	0.67	73.73	∞	73.73	∞	0.67	73.73
12	SW 8	81.30	70.76	0.98	73.25	0.00	81.30	1.38	73.04	0.02	51.30	78.30
13	SW 9	81.30	54.30	0.08	66.56	1.01	60.94	0.10	66.61	∞	1.01	60.94
14	SW 10	81.30	73.12	0.42	73.17	13.90	77.35	0.52	73.14	0.15	7.13	73.31
15	SW 11	81.30	73.81	0.17	74.55	0.83	73.82	0.22	74.47	0.17	0.93	73.81
16	SW 12	81.30	66.81	0.03	72.94	0.94	67.02	0.05	72.19	0.16	0.87	66.81
17	Cloned Joker	86.88	77.28	0.16	83.62	0.83	81.08	0.20	83.65	∞	0.82	81.07
Sum		1329.00	1054.55	—	1134.08	—	1103.21	—	1130.24	—	—	1092.96
Pooled		—	—	0.10	1192.00	0.53	1239.26	0.15	1180.25	0.11	1.57	1164.41

Table 3. *Maximum likelihood estimates for the matrix games. The “Random” and “Empirical” scores represent upper and lower bounds on the negative log likelihoods achievable by the four models.*

for all models degrades significantly when a single parameter configuration is estimated for all games pooled, compared to estimating separate parameters game-by-game. This can be seen in the table by comparing the likelihoods in the last two rows, marked “Sum” and “Pooled.” The decline in fit is least bad for HQRE, and substantially worse for CH relative to QRE (and HQRE). For HQRE, the likelihood ratio versus the random model declines from 199 to 149. This decline is greatest for CH, where it falls from 226 to 90. Thus, while CH provides the best fit when game-specific parameters are estimated, it provides the worst fit if the parameters are constrained to be constant across games.

The second conclusion is that despite their important structural differences, game-by-game and overall fits of the different models are surprisingly close. The QRE and CH fits differ by five or more likelihood points in only five of 17 games. Not surprisingly, TQRE also fits about as well as both QRE and CH, and slightly better in many cases, since it contains structural elements of both models. Our prior expectation was that the models would be widely separated in many of these games, but they are generally not. The surprise here is not that the models differ, but that *they differ by so little in most of these games*.

A final observation concerns the discretized TQRE estimates. In 7 of the 17 games we estimate $\gamma = \infty$, in which case the TQRE model collapses to Cognitive Hierarchy. This is interesting because TQRE incorporates both

result, values of τ that are quite far apart can produce similar log likelihoods.

QRE and CH features, and in some cases the estimates show that there is no positive effect (as far as likelihoods) to adding a quantal response element to the standard CH model. For 3 of these 7 games, there is also no improvement over QRE, but in the remaining 4, CH is clearly the better fitting model.

Also notice that HQRE and QRE offer nearly identical fits in many cases. Under a general distributional assumption for HQRE, it nests QRE (at least asymptotically, as with a log-normal distribution), and so the HQRE fit would necessarily be better. However, under the exponential distribution we estimate here, the models are not nested because there is necessarily heterogeneity in HQRE. As the parameter describing the mean responsiveness increases, the variance in responsiveness necessarily increases as well. Thus the two models are indeed substantially different. Our results indicate that this heterogeneity neither improves nor degrades the fit.

Table 3 makes clear that QRE and CH have surprisingly small differences in their qualities of fit for the game-by-game estimates. To see this relationship in another way, consider Figure 2. Each point corresponds to a single strategy from one of the 17 matrix games. The horizontal axis plots empirical choice frequencies for the strategies, and the vertical axis plots the predicted choice frequencies from the models at the pooled parameter estimate. QRE predictions are shown in black and CH in gray. For a perfect fitting model, like the “empirical” model shown in Table 3, all points would fall on the 45° line, shown in heavy black. Of course, both QRE and CH show substantial deviations from this line. Both models are also “biased” in the direction

of underpredicting extreme frequencies. That is, the models put too much weight on strategies that are empirically played the least often, and too little weight on strategies that are played the most often. This can be seen by looking at the solid and dashed lines, which show the best fitting lines to the scatter plots from the models. Both lines have positive intercept and slope less than unity. Perhaps most interestingly, the fitted lines are almost identical, and can barely be distinguished in the figure.

4.3 The negative frequency-payoff deviation relation

This surprising similarity in fit of models led us to think about whether the models might share some deeper structural properties. Note that in QRE, by construction, strategies with larger, more costly, deviations from optimal response (measured by expected payoffs) are played less often. To illustrate this property, start with the best-fitting QRE parameter λ^* for a game. Then calculate both the predicted frequency of play of each strategy, and the expected payoff difference from each strategy according to the model, relative to the optimal strategy. That is, calculate the expected loss assuming that other players' choice frequencies are given by the model. Figure 3 plots these predicted frequencies, on the y-axis, against expected payoff deviations (in pennies), for all strategies in all 14 symmetric games.²² The scatter plot shows a clear downward slope, which means that bigger mistakes are made

²²We restrict attention to the symmetric games because in our data they are exactly the games that are 3×3 . Thus the frequencies are more comparable across games.




Figure 2. Each point represents one strategy from one of the complete information games. Empirical choice frequencies are plotted on the horizontal axis, and predicted frequencies from pooled estimates of the models are on the vertical axis. QRE is shown in black, and CH in gray. The heavy black line is the 45° , which corresponds to a perfect fit, and the solid and dashed lines are the fits to the QRE and CH scatters. The fits are almost identical, and in both cases are flatter than the 45° .

less frequently.²³

Figure 4 shows the same frequency-deviation plot, constructed instead using the actual data from each game. That is, given a particular data set, one can compute expected payoffs from each strategy assuming that others use the empirical choice frequencies, and then compute expected costs of deviations from the strategy that is ex post empirically optimal. To make this precise, denote empirical choice frequencies of player i by $h_i = \{h_{i1}, \dots, h_{iJ_i}\}$. The collection $h = \{h_i\}$ of choice frequencies for each player then defines ex post expected payoffs $U_{ij}^e(h) = \sum_{a_{-i} \in A_{-i}} (\prod_{k \neq i} h_k(a_k)) u_i(a_{ij}, a_{-i})$. Denote i 's optimal strategy by $j^*(i) = \arg \max_{j \leq J_i} U_{ij}^e$. Then the expected cost of playing strategy j is $U_{ij^*(i)}^e - U_{ij}^e$. Of course, the empirically-optimal strategy produces a zero deviation and is played frequently. The basic pattern that is seen in the QRE plot—small mistakes (or zero mistakes) are common, and large mistakes are rare—is also evident in the empirical frequency plot.

As an example consider game 16, which is SW 12. Table 3 gives a QRE estimate of $\lambda^* = 0.03$. At the QRE estimate the predicted choice frequencies are (37%, 19%, 44%). These result in expected payoffs of (56.7, 36.9, 61.4). Thus the expected deviation costs are 4.7, 24.5, and 0, respectively. The three points that game 16 contributes to Figure 1 are therefore (0, 0.44), (4.7, 0.37), and (24.5, 0.19). On the other hand the empirical choice frequencies are (22%, 15%, 63%), which generate expected utilities of (62.3, 31.4, 63.0) for the

²³Note that if we made a separate plot for each game, the points would necessarily move monotonically downward by construction.




Figure 3. *QRE predicted choice frequencies computed at best fitting λ (vertical axis) for each strategy of each symmetric game, versus expected payoff cost in pennies, according to the model (horizontal axis).*




Figure 4. *Empirical choice frequencies for each strategy of each symmetric game (vertical axis), versus expected payoff cost in pennies, computed from the empirical distribution of play (horizontal axis).*

3 actions. The corresponding expected costs of deviation are $(0.07, 31.6, 0)$. Notice that while there are substantial differences between the empirical numbers and the QRE predictions (recall that the QRE estimate is fairly small for this game), the ordering is preserved both in terms of choice frequencies and expected payoffs. Thus the points contributed to Figure 2 are in a qualitatively similar pattern to those for Figure 1.

Now consider the CH model. A negative relationship between the frequency of strategy choices and their payoff deviations is *not* a structural component of the CH model. The reason is that mistakes in CH which lead to payoff deviations result from mistaken beliefs about the distribution of play, not from quantal response. As the number of thinking steps k increases, the accuracy of conditional beliefs improves, but there is no guarantee that the strategy choices for k -step thinkers will have expected payoffs that always increase with k . However, Figure 5 plots the frequency-deviation plot for the CH model, using best-fitting τ^* values (as in the QRE plot). The downward slope evident in the QRE (Figure 3) and empirical frequency (Figure 4) graphs is also evident in the CH graph (Figure 5), which looks remarkably like the QRE graph, even though the frequency-deviation decline is not built into CH as it is in QRE.

One way to understand the surprising frequency-deviation link in CH is to think about the extreme thinking-step types. Zero-step thinkers will randomize across strategies, so their predicted frequencies will be the same for all deviations (i.e., their frequency-deviation profile does not slope downward,

but it does not slope upward either). If the model is correct, the highest-step thinkers are playing the optimal strategy, because they have correct beliefs (their truncated beliefs about lower-type frequencies approximate the true distribution closely enough to generate the optimal strategy); those higher-step players will play only optimal strategies with zero deviation. Combining only these two extreme types together will produce a frequency-deviation plot that is weakly downward-sloping. Although there is no mathematical guarantee that adding in intermediate types will generate a strictly declining profile (as is true for QRE), Figure 5 nonetheless suggests that in practice the relation between frequency and deviation is usually negative.

Summarizing, the empirical frequencies of strategy play are typically declining in the deviation between a strategy's expected payoff (given the data) and the payoff of the empirically-optimal strategy. QRE reproduces this property by construction. CH also reproduces this property empirically, although it does not generally hold for all games. This gives one intuitive reason why CH and QRE approaches fit data about equally well, despite their structural difference. Also recall Figure 2, which shows that at an aggregate level, the two models make extremely similar predictions across our games.

We remark that the negative frequency-deviation relationship does not always hold for CH in our data. An example where it is violated is game 14, SW 10. We find an estimate of $\tau^* = 13.9$ for CH. Based on the corresponding choice predictions, level zero players, who randomize, earn an expected payoff of 38.1. This is better than level one players, who best respond to the uniform

mixture. This results in playing strategy 3, which earns an expected payoff of 36.9 against the true (according to the model) distribution of play. In fact, there are multiple instances where a level j player does better than a level $j + 1$ player in this example, because there are many cases where adjacent level players choose different actions.

4.4 The betting game and learning

We also studied a zero-sum betting game with asymmetric information over four states used by Sonsino, Erev and Gilat (2001) and replicated by Sovik (2004). The game is shown in Table 4. Player 1 has two information partitions, $\{A, B\}$ and $\{C, D\}$. Player 2 has three partitions, $\{A\}$, $\{B, C\}$, and $\{D\}$. Note that if the state is A or D, player 2 knows the state with certainty. The prior on the states is uniform. Players choose whether to “bet” or “not bet”. If both players bet, their payoffs are determined as in the top panel of Table 4. If at least one player opts out, then a paper-rock-scissors (PRS) game was played, with payoffs of 49, 23, and 36, for win, lose, and draw, respectively. This game has a unique equilibrium in which players randomize uniformly and have an expected payoff of 36. The reason for including the game is to avoid bias in presenting the betting game. The worry is that if the outside option consisted of a certain amount, then subjects may be tempted to over-bet, either because it is “more interesting,” or out of belief that the experiment would not make sense if they were meant to repeatedly not bet.

This game tests the “Groucho Marx Theorem” (Milgrom and Stokey,

1982)—the idea that privately-informed players should never agree to a zero-sum bet in equilibrium. With these payoffs, player 2 loses by betting on A, and wins by betting on D. As a result, although a CH 1-step risk-neutral player 1 will bet if her information is $\{A, B\}$ (thinking she is equally likely to win 31 or lose 29, relative to the expectation of PRS), in equilibrium she will never win since a rational player 2 will know the state if it is A, and won't bet. Hence if player 1 guesses that player 2 is rational, she won't bet if her information is $\{A, B\}$ because she deduces that she will never win the 31 and might lose 29. By similar logic, if player 2 is rational, thinks player 1 is rational, and thinks that player 2 thinks she (player 1) is rational, she can deduce that player 1 won't bet if player 1's information is $\{A, B\}$; player 2 therefore will not bet if her information is $\{B, C\}$, since she can only lose by so doing. One more step of iterated reasoning leads player 1 to not bet if her information is $\{C, D\}$. So there will be no state in which *both* players agree to bet, if players are sufficiently confident about rationality of others, and about others' perceptions of rationality.

However, Sonsino et al (2001) and Sovik (2004) find that players *do* bet, even after many periods of experience. In most of the Sonsino treatments, however, the marginal incentive is quite low; because they ran many periods, they used a low per-period conversion rate from experimental currency to Israeli Shekels (at stake was roughly 2.4 US cents per observation). In early periods a surprising fraction of player 2's bet when they are sure to lose in $\{A\}$ (around 20%) or don't bet in $\{D\}$ when they are sure to win (around

20% do not bet). This game was therefore included with some design changes to test the robustness of betting to higher incentives and other changes.

The main design change is that players who choose not to bet play a mixed-equilibrium game with expected value of 36 instead. This helped control for possible demand effects favoring betting, and also approximates the psychological value of betting with playing a mixed-equilibrium game in which the outcome is also uncertain.²⁴ One anomaly in our data is that in the betting game, one of four states is to be drawn for each match, with equal probability. In our data the counts for the four states are (183, 242, 255, 245). The number of times state A is drawn is significantly smaller than $\frac{1}{4}$.

In the betting game, CH and QRE both predict positive amounts of betting in all information partitions. Table 4 (bottom two panels) shows predicted betting rates for both models, using parameters estimated from our data across all periods ($\lambda^* = 0.64, \tau^* = 3.09$). Both models predict substantial betting when the information partitions contain two states, but the predictions are not sufficiently different that the data could point to one theory over the other. We played the game in 25 repeated rounds per session to see how quickly learning occurs (as do Sonsino et al (2001) and Sovik (2004)).

Figure 6 shows betting rates across time for both players in a four-period moving average. As with the Sonsino et al. (2001) and Sovik (2004) results,

²⁴Sonsino et al. (2001) also included one treatment in which there was a small fixed payment for not betting, which did not reduce betting rates. A fixed payment treatment is a step in the right direction but does not control for a taste for gambling or risk-preference.

our data show that betting is common and is slow to be extinguished by learning. However, our initial betting rates are significantly lower than Son-
sion et al (2001), showing more levels of sophisticated reasoning, a difference
due perhaps to our attempt to balance the design. Aggregating across peri-
ods, as shown in Table 4, the fit of the QRE and CH models as measured by
negative log likelihood scores are 1085.4 for QRE and 1074.2 for CH. As in
the complete-information normal-form games reported earlier, the two mod-
els are about equally accurate. We also fit models in which the parameters
 λ and τ drift up over time, as a reduced-form way of characterizing learning.
In QRE we estimate an initial $\lambda^0 = 0.55$ with a time trend of 0.012, which
results in a negative log likelihood of 1083.5, an improvement of less than two
points. Allowing for the time trend generates a larger improvement in CH.
We estimate $\tau^0 = 2.5$ with a time trend of 0.017, which has a correspond-
ing negative log likelihood of 1061.6, an improvement of over twelve points.
These results reinforce the central conclusion above, that despite their struc-
tural difference the QRE and CH approaches fit about equally well. Allowing
reduced-form learning in either model improves fit about equally well.

5 Conclusion

The quantal response equilibrium (QRE) approach in games assumes that
players “better-respond”—choosing strategies with higher expected payoffs
more often—but do not necessarily best-respond by choosing the highest

expected-payoff strategy all the time. (It is a fusion of Nash equilibrium with Luce’s (1959) stochastic utility model.) Cognitive hierarchy (CH) models go in a different direction; in the Poisson-CH form of Camerer, Ho and Chong (2004), players iterate reasoning in discrete steps but players doing one or more steps of iterated reasoning choose best responses given their beliefs. Both models have been shown to explain deviations from Nash equilibrium in many experimental datasets, and also explain why the Nash model fits well in some experiments (Goeree and Holt, 2001; Camerer, Ho and Chong, 2004).

We introduce a heterogeneous form of QRE, called HQRE, which creates a family of models that include QRE and CH as special cases. In HQRE, there is a distribution of response sensitivities, $f(\lambda)$; players know this distribution and optimize (given their λ values) accordingly. If $f(\lambda)$ is a degenerate distribution around one value of λ then HQRE is equivalent to QRE. If a player i with λ_i has subjective beliefs about the distribution $g(\lambda|\lambda_i)$, the resulting model is *subjective* HQRE. When the beliefs $g(\lambda|\lambda_i)$ are truncated at λ_i (i.e., $g(\lambda|\lambda_i) = 0 \forall \lambda > \lambda_i$) then TQRE results. The clear link to a stochastic form of CH comes from allowing a discrete distribution of TQRE with λ values $\{0, \gamma, 2\gamma, \dots\}$. When the increment γ goes to infinity, the hierarchical structure is retained and all players (except 0-step types) approximately best-respond. We prove that the discrete TQRE model with large γ approximates CH arbitrarily closely.

These models were estimated on new data from a variety of matrix games,

and a zero-sum incomplete information betting game. Surprisingly, the HQRE approach adds very little to QRE, and QRE and CH approaches fit about equally well. Both approaches can explain why there are substantial betting rates in the zero-sum game, contrary to the “Groucho Marx Theorem;” a learning model which allows QRE and CH parameters to change over time explains the modest amount of learning in a reduced-form way.

The essence of the QRE model is that larger deviations from best response are made less frequently. The CH model also tends to exhibit this property though it is not guaranteed and it is not built into the model specification as in QRE. The fact that both models exhibit this property in practice is an important clue about why they fit comparably well, despite their different surface structures.

Two challenges for future research are to find more games in which the models make sharply different predictions, and to explore further why their predictions are often so similar.

Another important area for research is extension of these ideas to extensive form games. QRE has been applied to extensive form games, typically in an “agent normal form” in which the response function at each information set is controlled by a separate agent (e.g., McKelvey and Palfrey, 1998). Extending hierarchical approaches to extensive form games is less straightforward, since it is natural to assume that agents low in the hierarchy are not responding as thoughtfully, and are also acting more myopically in a game with many stages (Camerer, et al. 2002b). Linking hierarchical reasoning

as in HQRE with differences in look-ahead could generate valuable insights and could provide a disciplined way to think about “chain store paradox”-type anomalies in which players do not appear to backward induct, even in relatively simple games (e.g., Johnson et al, 2002).

Finally, the preliminary understanding of these models by comparing them on experimental data is just meant to get a sense of where the models fit well and badly, and to see which restrictions are most plausible. The eventual hope is that these behavioral theories can be applied to the economic analysis of naturally-occurring games, just as conventional equilibrium concepts are now applied. The finding that these models can reliably explain behavioral departures from Nash equilibrium in one-shot games offers some promise that they can be useful in explaining anomalies in field data. A recent example is Crawford and Iriberri (2005), which applies level-k models to auctions. Models like these can also be used for economic design. Proposed institutions in which behavior predicted by models of limited rationality is far from what the designer intends (even if equilibrium behavior is ideal) might be bad designs in practice. Thus, applying these models is one approach to talk in a disciplined way about the robustness of mechanisms.

References

- [1] Camerer, C. (2004) *Behavioral Game Theory*, Princeton University Press: Princeton.

- [2] Camerer, C. and D. Lovo (1999) "Overconfidence and Excess Entry: An Experimental Approach," *American Economic Review*, 89, 306-318.
- [3] Camerer, C., T. Ho, and J. Chong (2004) "A Cognitive Hierarchy Model of Behavior in Games," *Quarterly Journal of Economics*, 119(3), 861-98.
- [4] Camerer, C., T. Ho, and J. Chong (2002a) "A Cognitive Hierarchy Model of Behavior in Games," *Quarterly Journal of Economics*, 119(3), 861.
- [5] Camerer, C., T. Ho, and J. Chong (2002b) "Sophisticated EWA Learning and Strategic Teaching in Repeated Games," *Journal of Economic Theory*, 104(1), 137-88.
- [6] Costa-Gomes, M. and V. Crawford (2005) "Cognition and Behavior in Two-Person Betting Games: An Experimental Study," *American Economic Review*, In Press.
- [7] Crawford, V. and N. Iriberri (2005) "Level-k Auctions: Can a Non-Equilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?" Working Paper, University of San Diego.
- [8] Crawford, V. and N. Iriberri (2004) "Fatal Attraction: Focality, Naivete, and Sophistication in Experimental Hide-and-Seek Games," Working Paper, University of San Diego.

- [9] El-Gamal, M., R. McKelvey, and T. Palfrey (1993) “A Bayesian Sequential Experimental Study of Learning in Games,” *Journal of the American Statistical Association*, 88, 428-35.
- [10] Fey, M., R. McKelvey, and T. Palfrey (1996) “An Experimental Study of Constant Sum Centipede Games,” *International Journal of Game Theory*, 25, 269-87.
- [11] Fudenberg, D. and E. Maskin (1986) “Folk Theorem for Repeated Games with Discounting or with Incomplete Information,” *Econometrica*, 54, 533-554.
- [12] Goeree, J. and C. Holt (2001), “Ten Little Treasures of Game Theory and Ten Intuitive Contradictions,” *American Economic Review*, 91(5), 1402-1422.
- [13] Goeree, J., C. Holt, and T. Palfrey (2005) “Regular Quantal Response Equilibrium,” *Experimental Economics*, 8, 347-67.
- [14] Johnson, E., Camerer, C., Sen, S. and Rymon, T. (2002), “Detecting Failures of Backward Induction: Monitoring Information Search in Sequential Bargaining,” *Journal of Economic Theory*, 104(1), 16-47.
- [15] Kahneman, D. and A. Tversky (1973) “On the Psychology of Prediction,” *Psychological Review*, 80, 237-51.

- [16] Kreps, D., P. Milgrom, J. Roberts and R. Wilson (1982) "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma," *Journal of Economic Theory* 27(2), 245-252.
- [17] Luce (1959) Individual Choice Behavior, Wiley, New York.
- [18] Maskin, E. (1986) "The Existence of Equilibrium with Price-Setting Firms," *American Economic Review*, 76(2), 382-86
- [19] McKelvey, R., and T. Palfrey (1992) "An Experimental Study of the Centipede Game," *Econometrica*, 60: 803-36.
- [20] McKelvey, R. and T. Palfrey (1995), "Quantal Response Equilibria in Normal Form Games," *Games and Economic Behavior*, 10(July):6-38.
- [21] McKelvey, R. and T. Palfrey (1998) "Quantal Response Equilibria in Extensive Form Games," *Experimental Economics*, 1, 9-41.
- [22] McKelvey, R., T. Palfrey, and R. 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, 523-48.
- [23] Milgrom, P. and N. Stokey (1982) "Information, Trade and Common Knowledge," *Journal of Economic Theory*, 26, 17-27.
- [24] Nagel, R. (1995) "Unraveling in Guessing Games: An Experimental Study," *American Economic Review*, 85(5), 1313-1326.

- [25] O'Neill, B. (1987) "Nonmetric Test of the Minimax Theory of Two-Person Zerosum Games," *Proceedings of the National Academy of Sciences*, 84(7), 2106-2109.
- [26] Santos-Pinto, L. and J. Sobel (2005) "A Model of Positive Self-Image in Subjective Assessments," *American Economic Review*, 95(5), 1386-1402.
- [27] Sonsino, D., I. Erev, and S. Gilat (2001) "On Rationality, Learning and Zero-Sum Betting – An Experimental Study of the No-Betting Conjecture," Working Paper, Technion.
- [28] Sovic, Y. (2004) "Strength of Dominance and Depths of Reasoning," Working Paper, University of Oslo.
- [29] Stahl, D. and P. Wilson (1995), "On Players' Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior*, 10(July), 218-54.
- [30] Weizsäcker, Georg (2003), "Ignoring the Rationality of Others: Evidence From Experimental Normal Form Games," *Games and Economic Behavior*, 44(July), 145-71.

6 Appendix

6.1 Proof of existence of HQRE

Theorem 1: *In finite games, a Heterogeneous Logit Equilibrium exists.*

Proof: To define the fixed point mapping, we take a slightly different approach from the standard one. Rather than identify a mapping, the fixed points of which are equilibria, we consider a fixed point in the induced mixed strategies, and then an equilibrium is constructed from the induced mixed strategies using (1). This simplifies the existence theorem because we are finding a fixed point in ΔA , a compact convex subset of $\Re^m$ rather than in a function space.

Let $\alpha \in \Delta A$. We construct the mapping $\Sigma : \Delta A \rightarrow \Delta A$ in the following way. Using (3), $U : \Delta A \rightarrow \Re^m$ maps α into $U(\alpha)$, where $m = \sum_{i=1}^n J^i$. Using (1), for each i , $P_i : \Re^{J^i} \times [0, \infty) \rightarrow \Delta A_i$ maps U_i into ΔA_i for each $\lambda_i \in [0, \infty)$. Finally, using (2), for each i , σ_i maps $P_i(U_i(\alpha))$ into ΔA_i by taking expectations over λ_i according to the distribution F_i . We define $\Sigma = \Sigma_1 \times \dots \times \Sigma_n$ by the composed mapping $\Sigma_i = \sigma_i \circ P_i \circ U_i \circ \alpha$. To see that this has a fixed point, observe that U_i is a single-valued, bounded continuous function on ΔA . Furthermore, P_i is single valued, continuous and uniformly bounded and hence $\int_0^\infty P_i(\lambda_i; U_i) dF_i(\lambda_i)$ exists for all U_i . Therefore, $\sigma_i(P_i)$ is well defined, and continuous by Lebesgue's dominated convergence theorem. Hence Σ is a continuous function from ΔA into itself and has a fixed point $\sigma^* \in \Sigma$. For each i and each $\lambda_i \in [0, \infty)$, let:

$$p_{ij}^*(\lambda_i) = \frac{e^{\lambda_i U_{ij}(\sigma^*)}}{\sum_{k=1}^{J_i} e^{\lambda_i U_{ik}(\sigma^*)}}.$$

so p^* is a Heterogeneous Logit Equilibrium. *QED*

Theorem 2: *In finite games, a TQRE exists.*

Proof: To define the fixed point mapping, we take a slightly different approach than above, because player i 's beliefs about other players strategies depends on λ_i . Rather than finding a fixed point in ΔA , a compact convex subset of $\Re^m$, we find a fixed point in distributional strategies, where a distributional strategy for i , σ_i , is a probability measure on the subsets of $[0, \infty) \times A_i$, the type-action product space, since in our approach i 's type is $\lambda_i \in [0, \infty)$. The proof is a straightforward adaptation of Milgrom and Weber (1982). The only two differences are: (1) players have truncated expectations rather than rational expectations; and (2) players quantal respond according to the logit rule instead of best responding.

Payoffs are equicontinuous because each A_i is finite. Because of the truncated distribution of beliefs, the (ex ante) expected payoff to player i is then defined slightly differently from Milgrom and Weber (p. 624), the difference being that the integral with respect to the distribution of *other* player types (λ_{-i}) is truncated at $\theta_i(\lambda_i)$ for each type λ_i . Since our distribution of types is independent and a density function exists for each f_i , and because $\theta_i(\lambda_i)$ varies continuously in λ_i , absolute continuity is satisfied, so we can express expected payoffs almost exactly as in Milgrom and Weber (1982, p. 625, expression 3.1), except for the well-behaved dependence of the upper bound for types λ_{-i} on λ_i . Consequently, using the topology of weak convergence for the distributional strategies, strategy sets are convex compact metric spaces and payoff functions are continuous and linear, so a fixed

point exists by Glicksberg's theorem (1952). The fact that we are considering quantal responses rather than best responses is of no consequence. It simply means that the fixed-point correspondence is single-valued and continuous rather than being multi-valued and upper hemicontinuous. *QED*

Theorem 3: *Fix τ . For almost all finite games Γ and for any $\varepsilon > 0$, there exists $\bar{\gamma}$ such that $\Delta^{\tau, \gamma} < \varepsilon$ for all $\gamma > \bar{\gamma}$.*

Proof: Fix τ and let Γ^τ denote the set of finite games with the property that in the CH model with parameter τ there is a unique best reply for all levels $k \geq 1$. It is straightforward to show that for each n and for each J , where J is the maximum size of any of the n players' strategy sets, the set of games without these properties has Lebesgue measure 0. Since the countable union of measure 0 sets has measure 0, this implies that Γ^τ consists of almost all finite games, in the generic sense. Let $g \in \Gamma^\tau$. Denote the unique maximizing action of a level k type of player i by a_{ik}^τ , and let δ_k^τ be the smallest difference in expected utility for a level k type of player i between choosing a_{ik}^τ and any other pure strategy. Fix $\varepsilon > 0$ and let L be an integer sufficiently large such that $\sum_{k=L}^{\infty} \frac{\tau^k}{k!} e^{-\tau} < \frac{\varepsilon}{3IJ}$. Denote $\bar{p}_{ijL}^\tau = \sum_{k=0}^L p_{ijk}^\tau \frac{\tau^k}{k!} e^{-\tau}$ and $\bar{p}_{ijL}^\gamma = \sum_{k=0}^L p_{ijk}^\gamma \frac{\tau^k}{k!} e^{-\tau}$. We immediately obtain that $|\bar{p}_{ijL}^\tau - \bar{p}_{ij}^\tau| < \frac{\varepsilon}{3IJ}$ for all i, j . Hence $\sum_{i=1}^n \sum_{j=1}^{J^i} (\bar{p}_{ij}^\tau - \bar{p}_{ijL}^\tau)^2 < \frac{\varepsilon}{3}$. Similarly, $\sum_{i=1}^n \sum_{j=1}^{J^i} (\bar{p}_{ij}^\gamma - \bar{p}_{ijL}^\gamma)^2 < \frac{\varepsilon}{3}$ for any γ . Note that for each i and k , $p_{ijk}^\tau = 0$ except if j is the index corresponding to action a_{ik}^τ . Next, we wish to examine $\bar{p}_{ij}^\gamma$, for large γ . First, we show that there exists a number $\bar{\gamma}(L)$ such that for all $\gamma \geq \bar{\gamma}(L)$, a_{ik}^τ is the unique maximizing action for *all* types $1 \leq k \leq L$

and $p_{ia_{ik}^\tau}^\gamma > 1 - \frac{\varepsilon}{3ILJ}$ for all $k \leq L$. That is, if $\gamma \geq \bar{\gamma}(L)$ then for all types L or lower types of player i , $|\bar{p}_{ijL}^\tau - \bar{p}_{ijL}^\gamma| < \frac{\varepsilon}{3LJ}$ for all $j \in S_i$. The proof is recursive. It is true for level 1 types because they have the same beliefs about other players that the CH-level-1 players have, and therefore have the same unique maximizing strategy a_1^τ . Therefore, by choosing a large enough γ we can make the probability a level 1 type of player i chooses a_{i1}^τ , as close to 1 as we wish. In particular, we can find some $\bar{\gamma}(1)$ so that it is greater than $1 - \frac{\varepsilon}{3ILJ}$ for all $\gamma \geq \bar{\gamma}(1)$. Level 2 (and higher) types are only slightly more complicated. For the level 2 types, their optimal strategy will be a_2^τ as long as the probability level 1's of the players other than i play a_{-i1}^τ is sufficiently close to 1. This will be true for all γ greater than some number, call it $\hat{\gamma}(1)$. Hence, we can find a $\bar{\gamma}(2)$ such that the probability a level 1 type of player i chooses a_{i2}^τ is greater than $1 - \frac{\varepsilon}{3ILJ}$ for all $\gamma \geq \bar{\gamma}(2)$. Proceeding recursively, we can do the same for level 3 and higher types, and so forth all the way to level L types. By construction, $\sum_{i=1}^n \sum_{j=1}^{J^i} (\bar{p}_{ijL}^\tau - \bar{p}_{ijL}^\gamma)^2 < \frac{\varepsilon}{3}$ for all $\gamma \geq \bar{\gamma}(L)$. Finally, by the triangle inequality:

$$\begin{aligned}
\Delta^{\tau,\gamma} &= \sum_{i=1}^n \sum_{j=1}^{J^i} (\bar{p}_{ij}^\tau - \bar{p}_{ij}^\gamma)^2 \\
&\leq \sum_{i=1}^n \sum_{j=1}^{J^i} (\bar{p}_{ij}^\tau - \bar{p}_{ijL}^\tau)^2 + \sum_{i=1}^n \sum_{j=1}^{J^i} (\bar{p}_{ijL}^\tau - \bar{p}_{ijL}^\gamma)^2 + \sum_{i=1}^n \sum_{j=1}^{J^i} (\bar{p}_{ijL}^\gamma - \bar{p}_{ij}^\gamma) \\
&< \frac{\varepsilon}{3} + \frac{\varepsilon}{3} + \frac{\varepsilon}{3} \text{ for all } \gamma \geq \bar{\gamma}(L) \\
&< \varepsilon
\end{aligned}$$

QED




Figure 5. *CH predicted choice frequencies computed at best fitting λ (vertical axis) for each strategy of each symmetric game, versus expected payoff cost in pennies, according to the model (horizontal axis).*

Betting Game					
Payoffs	Player 1	A	B	C	D
	Player 2	67	7	55	19
		3	63	15	51
Empirical Betting Frequencies	Player 1	A	B	C	D
	Player 2	28.5%	41.2%		
		5.5%	44.7%	96.7%	
QRE Betting Frequencies	Player 1	A	B	C	D
	Player 2	25.1%	37.1%		
		20.1%	46.0%	70.9%	
		$\lambda = 0.64$ neg Log L = 1085.4			
CH Betting Frequencies	Player 1	A	B	C	D
	Player 2	16.3%	60.4%		
		2.3%	38.0%	97.7%	
		$\tau = 3.09$ neg Log L = 1074.2			

Table 4. *The betting game payoffs (top), the empirical percentage of the time players chose “bet” by information set (second panel), and the QRE and CH predictions at their maximum likelihood estimates (third and fourth panels, respectively).*




Figure 6. Betting percentages for the different information sets of the two roles in the betting game. Each point represents a 4-period moving average.

Unprofitable				S-W 6				
40, 40	60, 10	10, 40	0.33	35, 35	39, 63	95, 91	0.44	
10, 60	10, 10	60, 40	0.24	63, 39	40, 40	56, 39	0.26	
40, 10	40, 60	40, 40	0.43	91, 95	39, 56	15, 15	0.30	
Cloned Matching Pennies				S-W 7				
50, 10	10, 20	10, 20	0.78	37, 37	93, 45	53, 53	0.59	
10, 20	20, 10	20, 10	0.10	45, 93	13, 13	85, 73	0.13	
10, 20	20, 10	20, 10	0.12	53, 53	73, 85	36, 36	0.28	
0.71	0.22	0.07						
Cloned Stag Hunt (Low)				S-W 8				
21, 21	10, 20	0.24		11, 11	59, 91	51, 51	0.17	
21, 21	10, 20	0.17		91, 59	27, 27	51, 43	0.20	
20, 10	20, 20	0.59		51, 51	43, 51	53, 53	0.63	
0.34	0.66							
Cloned Stag Hunt (Hi)				S-W 9				
31, 31	10, 20	0.45		50, 50	98, 44	70, 82	0.65	
31, 31	10, 20	0.34		44, 98	38, 38	70, 18	0.04	
20, 10	20, 20	0.21		82, 70	18, 70	70, 70	0.31	
0.90	0.10							
S-W 1				S-W 10				
35, 35	39, 47	95, 40	0.11	47, 47	51, 44	28, 43	0.52	
47, 39	51, 51	67, 15	0.81	44, 51	11, 11	43, 91	0.19	
40, 95	15, 67	47, 47	0.07	43, 28	91, 43	11, 11	0.30	
S-W 2				S-W 11				
79, 79	51, 59	55, 59	0.80	41, 41	97, 45	35, 58	0.50	
59, 51	31, 31	99, 67	0.15	45, 97	17, 17	53, 57	0.17	
59, 55	67, 99	19, 19	0.06	58, 35	57, 53	33, 33	0.33	
S-W 3				S-W 12				
73, 73	13, 77	49, 93	0.15	50, 50	30, 36	74, 42	0.24	
77, 13	41, 41	49, 41	0.17	36, 30	82, 82	18, 98	0.17	
93, 49	41, 49	46, 46	0.69	42, 74	98, 18	62, 62	0.59	
S-W 4				Cloned Joker				
42, 42	58, 50	98, 46	0.61	30, 10	10, 30	10, 30	10, 30	0.07
50, 58	54, 54	26, 66	0.35	30, 10	10, 30	10, 30	10, 30	0.07
46, 98	66, 26	18, 18	0.04	10, 30	10, 30	30, 10	30, 10	0.17
S-W 5				10, 30	30, 10	10, 30	30, 10	0.28
21, 21	93, 13	45, 29	0.33	10, 30	30, 10	30, 10	30, 10	0.41
13, 93	69, 69	53, 53	0.28	0.38	0.34	0.17	0.10	
29, 45	53, 53	61, 61	0.39					

57

Table A1. *Payoff matrices for the 17 normal form games, along with empirical choice frequencies (the row and column roles are combined in symmetric games).*