

Strategic Teaching and Equilibrium Models of Repeated Trust and Entry Games

Colin F. Camerer¹

Division of Humanities and Social Sciences
California Institute of Technology
Pasadena, CA 91125
camerer@hss.caltech.edu

Teck H. Ho

The Wharton School
University of Pennsylvania
hoteck@wharton.upenn.edu

Juin-Kuan Chong

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

Keith Weigelt

The Wharton School
University of Pennsylvania
weigelt@wharton.upenn.edu

May 17, 2002

¹This research was supported by NSF grant SES-0078911. Thanks to John Kagel for rapidly supplying data and to Drew Fudenberg, Qi-Zheng Ho, David Hsia, Xin Wang, and two referees for discussions and help. Useful comments were received from seminar participants at Berkeley, Caltech, Chicago, Harvard, Hong Kong UST, New York University, Pittsburgh, Princeton, and Wharton.

Abstract

This paper tests a model of strategic teaching in repeated trust and entry games with incomplete information. The model assumes ‘short-run’ players follow a one-parameter learning model (functional experience-weighted attraction). ‘Long-run’ players either realize how others are learning and ‘teach’ by maximizing their long-run payoff, or always behave honestly or aggressively. For precision, the fraction of honest/aggressive types was first measured in an experiment with one-shot games. Using data from 28 experimental sessions of eight-period trust and entry supergames (25,000 observations), the model fits modestly better than a quantal-response equilibrium benchmark, and both models predict much more accurately than chance. Estimates show most players are sophisticated, and become more sophisticated with experience. Direct measures of subjects’ beliefs are weakly correlated with implicit model beliefs, but are extremely accurate and do not show the overconfidence found in many psychological studies.

1 Introduction

Statistical models of learning have been applied to many experimental games. But most of the learning models are adaptive, and were applied to simple games. This paper tackles a harder problem. We apply a model in which some players are “sophisticated”—that is, they believe others are adaptive. The sophisticated players also know they will play the adaptive learners repeatedly, so they have an incentive to “teach” the learners what to expect (which later benefits the teachers). Sophistication without teaching has been incorporated in Stahl (1999) and Cooper and Kagel (2001). Teaching has only been studied by Camerer, Ho, and Chong (2002).

Strategic teaching has been proposed as a boundedly rational theory of reputation formation (see Fudenberg and Levine, (1989), Watson (1993), and Watson and Battigali (1997)). It is a promising way to explain the fact that the way players are rematched changes how they behave in experiments, and the fact that experienced subjects in experiments on repeated games of incomplete information behave in ways which are reasonably approximated by very complex, counter-intuitive Bayesian sequential equilibrium predictions, except in finite games when standard theories predict unraveling and no reputation-building. Camerer and Weigelt (1988a) suggested that the remarkable approximation of sequential equilibrium to data explained by a heuristic learning process. Strategic teaching is a good candidate for the sort of heuristic model they had in mind.

Extending learning models to include teaching is important for extending their applicability to important economic settings. Many transactions in the economy—perhaps most—are conducted repeatedly by players who know the history of behavior by others and anticipate future interactions. Examples include employment relations, banking relationships, long-standing corporate rivalries, customers who are loyal to retailers, cartels, and so forth. Learning models which do not include the possibility of teaching have little to say about these important exchange relationships.

Camerer, Ho and Chong (2002) applied a parameter-rich learning model to a small sample of eight trust game sessions. However, parameter-rich learning quickly becomes unwieldy when applied to asymmetric games if players in different roles are allowed to

have different learning parameters. They used 17 parameters to model adaptive players of one type and adaptive and sophisticated teaching players of another type. As a result, parameters are not always precisely identified and vary across different experimental sessions.

This paper fits models based on strategic teaching to data from 28 experimental sessions on trust and entry-deterrence games (about 20,000 stage games). The predictive accuracy of teaching is compared to a generalized equilibrium model called agent-based quantal response equilibrium (AQRE). AQRE is more general than a conventional equilibrium model because players optimize noisily and have a common prior belief that some players always prefer to behave in a trustworthy (or spiteful) way, though it is costly for those players to do so. But players in AQRE update their beliefs using Bayes' rule and anticipate accurately what others will do, so AQRE retains key properties of Bayesian-Nash equilibrium.

This paper introduces three innovations in the statistical use of strategic teaching models. First, the parametric EWA model in the learning part of the learning-teaching model is replaced by a much leaner functional experience-weighted attraction (fEWA) approach. In fEWA, fixed parameters are replaced by functions of the data. Ho, Camerer and Chong (2002) specified functions which reproduce variation in key features of learning processes across ten games; the resulting fEWA model predicts as accurately as models with more parameters. To explain teaching, three other parameters are included, which represent the strength of cross-sequence learning, and the frequencies of players who teach and players who are always honest or aggressive.

Second, the same teaching model is applied to trust and entry games which are conceptually quite different. In these two games the unusual types of player behavior (about which others are incompletely informed) are honesty or aggression. A common interpretation of the source of the unusual types' behavior is unobserved variation in personal utilities or (for firms) corporate policies or philosophies. For example, trustworthy players feel obliged to reciprocate trust, or put a positive weight on how much others earn. Aggressive incumbent firms put a negative weight on their rivals' profits or feel spiteful. But these types of tastes are essentially opposite. An obvious fear in applying such

models is that it is too easy to invent an ad hoc player type to explain any pattern of behavior (as Ledyard (1986) proved). However, the teaching model produces these opposite behaviors *endogenously*: Both types are Stackelberg types who choose the move they would precommit to if they could. (In trust games, players would precommit to being trustworthy if they could, to win trust and earn more; and in entry games firms would precommit to fight entry if they could, to deter entry and earn more.) We also constrain the frequency of unusual types by measuring how common they are in a separate experiment with one-shot games.

Third, in eight new trust game sessions, subjects report beliefs about the chance of a default. The belief reports serve as a diagnostic check on the models (cf. Nyarko and Schotter (2002) and Manski (2002)) and incidentally, are remarkably well-calibrated to the events that actually occur, which means the subjects know something the AQRE model is not capturing. (Also, the superb calibration in reported beliefs is the first evidence that overconfidence which is routinely observed in psychology studies disappears when incentivized subjects become expert in forecasting behavior of others like themselves in an artificial domain.)

The next section describes the adaptive parametric EWA model and fEWA. Section 3 introduces teaching. Section 4 reports results from two experiments that measured the proportions of homemade prior in trust and entry-deterrence games. Section 5 reports estimates of teaching model on three data sets from repeated trust and entry-deterrence games, and discusses subtler distinctions between the teaching and AQRE approaches. Section 6 describes measurement of subjects' beliefs and assesses their accuracy. Section 7 concludes.

2 Adaptive EWA Learning

We start with notation. Each of n players are indexed by i ($i = 1, \dots, n$). The strategy space of player i , S_i consists of m_i discrete choices, that is, $S_i = \{s_i^1, s_i^2, \dots, s_i^{m_i-1}, s_i^{m_i}\}$. $S = S_1 \times \dots \times S_n$ is the Cartesian product of the individual strategy spaces and is the strategy space of the game. $s_i \in S_i$ denotes a strategy of player i , and is therefore

an element of S_i . $s = (s_1, \dots, s_n) \in S$ is a strategy combination, and it consists of n strategies, one for each player. $s_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$ is a strategy combination of all players except i . S_{-i} has cardinality $m_{-i} = \prod_{j=1, j \neq i}^n m_j$. The scalar-valued payoff function of player i is $\pi_i(s_i, s_{-i})$. Denote the actual strategy chosen by player i in period t by $s_i(t)$, and the strategy (vector) chosen by all other players by $s_{-i}(t)$. Denote player i 's payoff in a period t by $\pi_i(s_i(t), s_{-i}(t))$.¹

Like most learning models, the EWA model assumes each strategy has a numerical attraction which determines the probability of choosing that strategy (see Camerer and Ho (1999)). Learning models require a specification of initial attractions of strategies, how attractions are updated by experience, and how choice probabilities depend on attractions. The key variable is $A_i^j(a, t)$, an adaptive player i 's attraction of strategy j after period t has taken place.

The variables $N(t)$ and $A_i^j(a, t)$ begin with prior values, $A_i^j(a, 0)$ and $N(0)$ ($A_i^j(a, 0)$ could be derived from a model of first-period thinking, as in Camerer, Ho and Chong (2001); and $N(0) = 1$ below for simplicity). The experience weight is updated according to $N(t) = (1 - \kappa) \cdot \phi \cdot N(t - 1) + 1$. Attractions are updated by weighting hypothetical payoffs that unchosen strategies *would* have earned by a parameter δ , and weighting payoffs actually received by an additional $1 - \delta$ (for a total weight of 1). Define an indicator function $I(x, y)$ which equals 1 if $x = y$ and 0 if $x \neq y$. Then the weighted payoff is $[\delta + (1 - \delta) \cdot I(s_i^j, s_i(t))] \cdot \pi_i(s_i^j, s_{-i}(t))$.

Attractions are updated according to

$$A_i^j(a, t) = \frac{\phi \cdot N(t - 1) \cdot A_i^j(a, t - 1) + [\delta + (1 - \delta) \cdot I(s_i^j, s_i(t))] \cdot \pi_i(s_i^j, s_{-i}(t))}{N(t)}. \quad (2.1)$$

Attractions determine choice probabilities using a logit stochastic response rule, where λ is response sensitivity (infinite λ is best-response):

$$P_i^j(a, t + 1) = \frac{e^{\lambda \cdot A_i^j(a, t)}}{\sum_{k=1}^{m_i} e^{\lambda \cdot A_i^k(a, t)}}. \quad (2.2)$$

¹In estimation we rescale payoffs by subtracting the minimum payoff so all rescaled payoffs are positive. And we convert experimental currency units to dollars so that, in principle, values of the response sensitivity λ can be compared across games and countries.

The parameter δ can be interpreted as a kind of ‘imagination’ of, or responsiveness to, foregone payoffs. The decay rate ϕ reflects a combination of forgetting and the degree to which players realize other players are adapting, so that old observations of what other players did become less and less useful. The parameter κ affects the growth rate of attractions, which determines how sharply players lock in to a strategy.

In 1988, Francis Crick, who co-discovered DNA, wrote that “In nature hybrid species are often sterile, but in science the reverse is true”. The EWA rule was developed to hybridize the key features of simpler models, hoping Crick’s dictum might prove helpful in the domain of human learning. When $\delta = 0$ the EWA rule reduces to reinforcement rules in which only the received payoff drives learning.² When $\delta = 1$ and $\kappa = 0$ the EWA rule is, perhaps surprisingly, equivalent to belief learning according to weighted fictitious play. (In weighted fictitious play, players update beliefs by computing geometrically weighted averages of what others have done in the past, then choosing best responses, e.g., Fudenberg and Levine (1998).)

The EWA structure shows that weighted fictitious play belief learning is simply generalized reinforcement in which unchosen strategies are reinforced by foregone payoffs. The equivalence arises because when expected payoffs based on updated beliefs are written as a function of lagged expected payoffs, the belief terms disappear; the payoff implications of updating beliefs are fully encoded by directly reinforcing unchosen strategies.³

Another way to think of EWA is as a kind of reinforcement learning in which strategies only increase in probability if their reinforcement is above a “reference point” (the δ -weighted average of all foregone payoffs). This reference point changes endogeneously and requires no extra parameters.

The EWA hybrid is consistent with recent theories of “dual processes” in the brain (e.g. Kahneman and Frederick (2002), and Camerer, Loewenstein, and Prelec, in press).

²E.g., Harley (1981), Arthur (1991), McAllister (1991), Roth and Erev (1995), Roth et al. (2001).

³Note that this equivalence is only true for weighted fictitious play, not for beliefs which change in other ways. The connection was discussed by Cheung and Friedman (1997), and Fudenberg and Levine (1998). And cf. Hopkins (in press), who demonstrates another type of equivalence—when payoffs in reinforcement models are divided by the frequency of strategy choice, the resulting dynamics are the same as in belief learning.

These theories posit two separate, interacting processes. One process uses “quick and dirty” affective information, habit, analogy, and pattern matching. The other process is slower, more thoughtful, and can override faulty judgments from the first process with logical checking. The fact that received payoffs and foregone payoffs get different weight might reflect two processes. The automatic, habitual process instinctively puts plenty of weight on the payoff from the chosen strategy (in experiments, attention is usually rapidly focused on this number). Then the deliberative process kicks in and asks how well other strategies would have done. The weight δ can then be interpreted as the strength of deliberation versus automatic habit.

While it is interesting to have potential cognitive underpinnings for a behavioral model, our primary concern is that the model fits and predicts well. In previous research, best-fitting EWA parameters were estimated in thirty experimental games with a wide variety of strategic structures (see Camerer, Ho and Chong (2002)).⁴ In most games (those with mixed equilibria are an exception) the parameter values implied by the reinforcement and belief restrictions are strongly rejected.

2.1 Functional EWA

EWA parameter estimates are systematically different across games (as is usually the case for learning models; see Cheung and Friedman (1997); Crawford (1995)). Cross-game parameter variation creates a need for a more flexible specification which can explain why parameters differ across games, and can predict what parameter values will fit best in new games. To address this question, Ho, Camerer and Chong (2002) developed a *functional* model (fEWA) in which the fixed parameters ϕ, δ and κ were replaced by functions of data which self-adjust across games and over time (fixing $N(0) = 1$ for simplicity leaves only one free parameter to must be estimated in fEWA, the response sensitivity λ). In fact, if the goal is to maximize the model’s hit rate, *no* free parameters are needed once

⁴Models are typically fit on part of the data and penalized for degrees of freedom using Bayesian information criteria, then parameter values are frozen and used to predict a hold-out or validation sample. Both procedures ensure that more complex models do not fit better simply because they have more degrees of freedom.

initial conditions are specified.⁵

The functional approach is particularly useful in games with asymmetric players because the number of parameters grows very large if different player types are permitted to have different parameter values. For example, when Camerer, Ho, and Chong (2002) applied the parametric EWA model to repeated games with incomplete information and asymmetric parameters they used 17 parameters. The fEWA version has only three substantive parameters (τ , α and θ), and five nuisance parameters (an initial attraction $A_i^{IN}(0)$ and values of λ for each player type, which can be set equal if necessary). The model is also particularly good at fitting behavior from one sample of games and then predicting behavior in an entirely new game.

fEWA replaces the parameters ϕ, δ, κ with deterministic functions $\phi_i(t), \delta_i(t), \kappa_i(t)$ of player i 's experience up to period t . These functions determine parameter values for each player and period, which are then plugged into the EWA updating equation to determine attractions.

The function $\phi_i(t)$ is designed to detect change in the learning environment. Ideally, the value of ϕ should fall when change is rapid so that previous experience should be largely ignored in figuring out what to do next. The heart of the function is a “surprise index”, the difference between the other players’ strategies in the window of the last W periods and their average strategy in all previous periods (where W is the fewest number of strategies played in any Nash equilibrium). The ϕ function is defined in terms of relative frequencies of strategies, without using information about how strategies are ordered, so that it can be applied to non-ordered strategies (e.g., rows in a normal-form game). It is easy to construct similar functions which exploit order information, for application to games where strategies are prices, quantities, locations, etc.

To measure change, $\phi_i(t)$ takes the differences in corresponding elements of the two frequency vectors, square them, and sum over strategies. Formally, the change-detection

⁵The hit rate is the fraction of choices which are predicted to be most likely that are actually chosen. Since the strategy with the highest attraction will be predicted to be most likely for any λ , the hit rate does not depend on λ .

function $\phi_i(t)$ is

$$\phi_i(t) = 1 - .5 \left(\sum_{j=1}^{m-i} \left[\frac{\sum_{\tau=t-W+1}^t I(s_{-i}^j, s_{-i}(\tau))}{W} - \frac{\sum_{\tau=1}^t I(s_{-i}^j, s_{-i})}{t} \right]^2 \right) \quad (2.3)$$

The term $\frac{\sum_{\tau=t-W+1}^t I(s_{-i}^j, s_{-i}(\tau))}{W}$ is the j -th element of a vector that simply counts how often strategy j was played by the others in the W periods from $t - W + 1$ to t , and divides by W . The term $\frac{\sum_{\tau=1}^t I(s_{-i}^j, s_{-i})}{t}$ is the relative frequency count of the j -th strategy over all t periods.

The function is not derived from axioms but it does reflect several intuitions. Multiplying the sum of squared differences by .5 and subtracting half the sum from one just normalizes so that ϕ is between zero and one. Squaring deviations means that ϕ is closer to one when previous choices by other players are more dispersed. The model effectively senses that when other players have always been playing variably, experience should not be decayed too rapidly (or put differently, it does not overreact to recent fluctuations when fluctuation is normal). Furthermore, ϕ dips lowest when recent behavior by others is most surprising—i.e., when previous experience is constant but suddenly switches (then $\phi_i(t) = \frac{2t-1}{t^2}$). Note that in this case of maximal surprise, the drop in ϕ is greater when the stretch of previous experience is longer. The averaging window of W periods is used to calculate recent experience because in a mixed equilibrium (i.e., $W > 1$) choices by others will naturally fluctuate even if players are using stable mixtures. Averaging over W recent observations smooths fluctuations and ensures that ϕ does not mistake equilibrium mixing for genuine change and dip too low.

The parameter $\delta_i(t)$ is set to $\phi_i(t)/W$. Tying δ to ϕ encodes a kind of freezing or status quo bias—when the environment is changing rapidly, players tend to do what they did in the past. Dividing by W is useful empirically because in mixed games, best-fitting estimates of δ tend to be lower than in games with pure equilibria (and close to zero in many games).

The parameter κ controls the growth rate of attractions. Psychologically, κ can be interpreted as the extent to which players “explore” by choosing different strategies, relative to how quickly they “exploit” what they have learned by switching to a constant choice of the strategy which has performed the best in the past (see Sutton and Barto

(1998)). Players with low κ are constantly exploring— they just keep track of average (δ -weighted) payoffs. When players “exploit” they quit exploring and commit to a strategy by choosing it frequently, even if its average previous payoff is not much larger than the average previous payoffs of other strategies. If payoffs are positive (and $\delta < 1$), a higher κ means players are basically rewarding a strategy they choose a lot, simply for being chosen. This is one way of characterizing lock-in or exploitation empirically (cf. Polya urns).

This line of reasoning suggests using variation in how frequently a player uses different strategies to track when they explore and when they exploit. Therefore, $\kappa_i(t)$ is set equal to the Gini coefficient measuring the dispersion or “inequality” of previous choices.⁶ When a player’s strategy choices are dispersed, the Gini coefficient is low and so κ is low (reflecting her exploration). When she starts to lock in to one strategy, the Gini rises and κ moves toward one, reflecting exploitation.

Ho, Camerer and Chong (2002) apply the fEWA model to experimental data from 10 games and compare its accuracy to three other learning models with more parameters (parametric EWA, belief learning, reinforcement with payoff variability; see Roth et al (2001)) and a non-learning quantal response (noisy) equilibrium model.⁷ fEWA fits in-sample and predicts out-of-sample about as well as the other learning models in most games, and often better, even though it has fewer parameters. Most importantly, the values of ϕ and δ generated by the fEWA functions, averaged across players and time, are close to the best-fitting parametric EWA estimates across games. The fact that these

⁶To calculate the Gini coefficient for subject i , first rank strategies from most-probable to least-probable (using observed choice frequencies). Denote the rank-ordered choice proportions of these strategies by $f_i^{(1)}(t)$ to $f_i^{(m_i)}(t)$. Then plot a cumulative probability distribution which measures the total probability of the strategies used as frequently as j or less frequently, $C_i(j, t) = \sum_{k=1}^j f_i^{(k)}(t)$. This calculation gives j points; use linear interpolation to create a piecewise-linear function connecting the points. The Gini coefficient is then the area between the identity line and the interpolated function passing through the $C_i(j, t)$ points, normalized so that Gini coefficients range from zero (when all strategies are played equally often) to 1 (when one strategy is played all the time). The normalized Gini coefficient is $\kappa_i(t) = 1 - 2 \cdot \left\{ \sum_{k=1}^{m_i} f_i^{(k)}(t) \cdot \frac{m_i - k}{m_i - 1} \right\}$ where $f_i^k(t)$ are ranked from the lowest to the highest.

⁷The initial values $\phi_i(0)$ and $\kappa_i(0)$ are set to .5. A weighted average with weight $(1/t)$ on these initial values and $(t-1)/t$ on the values computed using the text formulae were used to specify the functional values, to smooth out fluctuations in values in early periods.

functions economize on parameter estimation, and reproduce cross-game variation in parameters, justifies applying them in this paper. Parametric EWA is also estimated for comparison.

3 Strategic teaching

Most of the statistical models of learning (like those described in the previous section) are adaptive. In adaptive models, players ignore information about payoffs of others and do not realize others are learning. Players who are aware that others are learning are called “sophisticated” (see Selten (1986); Milgrom and Roberts (1991); Fudenberg and Kreps (1990)). There is clear empirical evidence that some players use information about payoffs of others⁸, are sophisticated about learning of others, and become more sophisticated with experience (Stahl (1999); Cooper and Kagel (2001); Camerer, Ho and Chong (2002)).

In our specification, sophistication is modeled by assuming some fraction of players believe that others are learning adaptively. This approach is stylized to the asymmetric games studied below, in which ‘long-run’ players may be sophisticated and understand how ‘short-run’ players learn.⁹

If a player is sophisticated and believes she will be rematched with another player in the future, she might take into account the effect of her period t action on the adaptive players’ period $t + 1$ actions, because those actions will change the sophisticated player’s period $t + 1$ payoffs. Maximizing discounted expected payoffs, taking into account the effect of one’s own current behavior on future behavior of others, is called “strategic

⁸See Partow and Schotter (1993); Mookerjee and Sopher (1994); Cachon and Camerer (1996).

⁹Our earlier paper assumes sophisticated players know the mixture of adaptive learners (who use EWA rules) *and* know how many others are sophisticated like themselves, in a symmetric game (cf. Stahl (1999)). That specification embeds equilibrium concepts as a limiting case when everyone is sophisticated and is simple. Camerer et al. (2002) also estimate a more complex model in which the sophisticated player’s beliefs about how many others were sophisticated can be wrong. They found that sophisticated players tended to underestimate how many others were sophisticated (perhaps reflecting overconfidence about their own relative skill).

teaching”. There are also myopic player types who forecast sophisticatedly, but do not teach.

The idea of strategic teaching was introduced by Fudenberg and Levine (1989). They showed that a very patient strategic teacher can get almost as much utility as from the Stackelberg equilibrium, by playing an optimal precommitment strategy forever (and waiting for the adaptive player to learn to best-respond (see also Watson (1993) and Watson and Battigali (1997), and Ellison (1997)).

One implication of strategic teaching is that the way players are rematched will affect behavior. Of course, experimenters have been concerned about such possible effects for decades, and often deliberately choose a design with random rematching (typically, without repetition) to disable the possibility of teaching. In fact, when subjects play repeatedly with a fixed partner, more mutual cooperation and efficiency are achieved in prisoners’ dilemmas (Andreoni and Miller (1993)) and stag hunt games (Van Huyck, Battalio and Beil (1990); Clark and Sefton (1999)) than when players are rematched randomly with strangers. Another implication of strategic teaching is that group size matters when the incentive to teach depends on the ‘slowest learner’ in the group. For example, players reach efficiency less often in 3-player stag hunt games than in 2-player games (Knez and Camerer (2000)), experimental collusion is harder to sustain with more firms (e.g., Holt (1995)), and two buyers can force monopoly prices down by boycotting (withholding demand) but four buyers cannot (Ruffle (2001)).

To illustrate the details of how teaching works, consider a repeated trust game studied experimentally by Camerer and Weigelt (1988a). (The ideas transfer naturally to entry games). In the trust game, a single borrower B wants to borrow money from a number of lenders denoted L_i ($i = 1, \dots, N$) (cf. Kreps (1990)). A lender makes only a single lending decision (either *Loan* or *No Loan*) and the borrower makes a string of decisions, (either *Repay* or *Default*), if the lender chooses *Loan*.

In a typical experimental session, subjects are randomly assigned fixed roles of borrower, or lender (e.g., 11 subjects are divided into 3 borrowers and 8 lenders). In a single sequence, a borrower B is randomly chosen to play in every period of an eight-period

supergame.¹⁰ Each lender L_i plays in exactly one of the 8 stage games in a random order (which is unknown to the borrower). To study cross-sequence learning, the entire supergame is repeated in a series of *sequences* (typically 50 to 100).

Index each sequence of game rounds by k and each game round by t . The goal of the model is to specify the probabilities of lending and repaying in *each* round of *each* sequence.

Recall that lenders play only once in a sequence, and borrowers play in only a third of the sequences. Yet players clearly respond to experiences they observe, so we assume “observational learning”. Players can also learn from previous rounds in a sequence *and* from previous sequences. Consider round 7 in sequence 14. The round 7 lender who is deciding what to do saw what happened in the previous 6 rounds of sequence 14, and learned about the attractiveness of lending from what happened in those rounds. But the lender also knows what happened in the upcoming (7th) round of the previous sequences 1-13– a “peripheral vision” glance at the past– and learned about whether she should loan in round 7 from those previous round 7 experiences. We call the latter effect cross-sequence learning.

Within-sequence and cross-sequence learning could be integrated in various ways. Returning to our example, the strategy “loan” for a lender before period 7 of sequence 14 is influenced by two sources of experience– the attraction of “loan” after period 6 of sequence 14, and the experience of period 7 of sequences 1-13. Reflecting a prior belief that within-sequence learning is more important than cross-sequence learning, we elected to make updating attractions within a sequence the basic operation, then include an extra step of partial updating using the average payoff from the upcoming round in previous sequences.

The strength of cross-sequence learning is parameterized by a parameter τ . If $\tau = 0$ there is no cross-sequence learning; if $\tau = 1$ experience in upcoming periods of previous sequences is just as important as experience in the previous period of the current sequence. The data will tell us how strong cross- sequence learning is through the value of τ .

¹⁰After each supergame, one of the two idle borrower subjects is chosen to play the next supergame, so no borrower plays two supergames in a row, to minimize overarching reputational effects.

Define the attraction of an adaptive lender at the end of sequence k and round t by $A(a, k, t)$. The updating occurs in 2 steps. The idea is to create an “interim” attraction for round t , $B(a, k, t)$, based on the attraction $A(a, k, t - 1)$ and payoff from the round t , then incorporate experience in round $t + 1$ from previous sequences, transforming $B(a, k, t)$ into a final attraction $A(a, k, t)$.

- Step 1 (adaptive learning across rounds within a sequence):

$$B_L^j(a, k, t) = \frac{\phi \cdot A_L^j(a, k, t - 1) \cdot N(k, t - 1) + (\delta + (1 - \delta) \cdot I(j, s_L(k, t)) \cdot \pi_L(j, s_B(k, t)))}{M(k, t)}$$

$$M(k, t) = \phi(1 - \kappa) \cdot N(k, t - 1) + 1$$

- Step 2 (simulated learning in a coming round from previous sequences):

$$A_L^j(a, k, t) = \frac{\phi^\tau \cdot B_L^j(a, k, t) \cdot M(k, t) + \tau \cdot \delta \cdot \hat{\pi}_L^j(k, t + 1)}{N(k, t)}$$

$$N(k, t) = [\phi(1 - \kappa)]^\tau \cdot M(k, t) + \tau$$

Assume that the learning about an upcoming round from previous sequences is driven by the average payoff in that round in previous sequences. Formally, $\hat{\pi}_L^j(k, t + 1) = \sum_{m=1}^{k-1} \pi_L(j, s_B(m, t + 1)) / (k - 1)$.

Borrowers with normal payoffs come in two types: Myopic, and teachers. Since myopic borrowers move after the lenders do, the attractions of repay and default are simply the stage-game payoffs and they choose using a logit rule. A *teacher* borrower guesses how the lender learns, adapts those guesses to experience, and also plans actions for the remaining periods within a game sequence. A teacher borrower’s attractions are given by:

$$A_B^j(s, k, t) = \sum_{j'=Loan}^{NoLoan} P_L^{j'}(a, k, t + 1) \cdot \pi_B(j, j')$$

$$+ \max_{J_{t+1}} \left\{ \sum_{v=t+2}^T \sum_{j'=Loan}^{NoLoan} \hat{P}_L^{j'}(a, k, v | j_{v-1} \in J_{t+1}) \cdot \pi_B(j_v \in J_{t+1}, j') \right\}$$

where $\hat{P}_L^{j'}(a, k, v|j_{v-1}) = \hat{P}_L^{Loan}(a, k, v-1|j_{v-1}) \cdot P_L^{j'}(a, k, v|(Loan, j_{v-1})) + \hat{P}_L^{NoLoan}(a, k, v-1|j_{v-1}) \cdot P_L^{j'}(a, k, v|(NoLoan, j_{v-1}))$. J_{t+1} specifies a possible path of future actions by the sophisticated borrower from round $t + 1$ until end of the game sequence. That is $J_{t+1} = \{j_{t+1}, j_{t+2}, \dots, j_{T-1}, j_T\}$ and $j_{t+1} = j$. To simplify computation, we search only paths of future actions that always have default following repayment because the reverse behavior (repayment following default) is rare. $P_B^j(s, k, t + 1)$ is derived from $A_B^j(s, k, t)$ using a logit rule.¹¹

The homemade prior is the fraction θ of the types with normal payoffs who act as if they have honest-type payoffs, and choose according to a logit rule. A fraction p of the players are also induced to have honest payoffs by the experimenter, so the total fraction who behave honestly is $p + (1 - p)\theta$. Of the $(1 - p)(1 - \theta)$ who do not behave honestly, a fraction α of the borrowers are teachers and $1 - \alpha$ are myopic. The likelihood of observing the data is then given by $\Pi_k[(p + (1 - p)\theta)\Pi_t P_B^B(t)(h, k, t) + (1 - p)(1 - \theta)[(1 - \alpha) \cdot \Pi_t P_B^{S_B(t)}(m, k, t) + \alpha \cdot \Pi_t P_B^{S_B(t)}(s, k, t)]]$ where h and m denote honest and myopic types.

The model is estimated on using repeated game trust data from Camerer and Weigelt (1988a, 1988b) and entry deterrence games from Jung, Kagel and Levin (1994). Maximum likelihood estimation (MLE) was used to calibrate the model on 70% of the sequences in each experimental session, then forecast behavior in the remaining 30% of the sequences in that session. If the model fits in-sample purely by overfitting, it will perform surprisingly *poorly* out-of-sample.¹²

The delicate logic of the repeated-game equilibrium can be illustrated with the trust game. Table 1 shows payoffs in the Camerer-Weigelt repeated trust game. Recall that a single borrower is drawn to play an 8-period sequence. Her type (either honest or dishonest) is drawn randomly using a commonly-known prior and communicated only to the borrower. The borrower then plays a sequence of stage games with eight lenders who

¹¹Our earlier work included a discount factor ϵ which weighted future expected payoffs. Those estimates were usually close to one so now we restrict $\epsilon = 1$ for parsimony.

¹²We used GAUSS. To avoid getting stuck in local maxima, we typically posited two or three starting values for each parameter, and used 64 combinations of possible parameter values as different initial conditions. After 50 iterations from each initial condition, we chose the best-fitting estimates and continued iterating to convergence.

play once each in random order.

In each stage game, the lender can choose not to lend (then both earn 10 currency units) or can choose to lend. Lenders prefer to lend if the borrower will choose to repay, yielding 40 for the lender. But if the borrower defaults the lender earns -100.¹³ A dishonest borrower earns 10 if the lender does not default, and 60 if she repays. Honest-type borrower have the same payoffs except a default pays 0. Note that in the subgame after receiving a loan, the myopic dishonest borrower prefers to default while the honest borrower prefers to repay. The probability that an borrower had honest-type payoffs in a particular sequence was varied from 0.33 to 0.

The sequential equilibrium is computed from the last period forward. We sketch the equilibrium (assuming best response rather than quantal response) briefly for sessions 4-6 (see Camerer and Weigelt (1988a) for details). In the last period, risk-neutral lenders lend if their perceived $P(\text{Honest})$ is a threshold $\gamma = .79$. Anticipating this, normal borrowers mix in period 7 by repaying with enough probability to make the lender's updated $P(\text{Honest}) = .79$ in period 8 and makes lenders indifferent (they lend with probability .64). Guessing accurately what borrowers will do, the lender's $P(\text{Honest})$ threshold in period 7 is γ^2 . The same argument works by induction back to period 1. In each period the lender has a threshold of perceived $P(\text{Honest})$ which makes her indifferent between lending and not lending. The path of these threshold $P(\text{Honest})$ values is simply γ^{9-n} in period n . When $P(\text{Honest})$ in period t is above the threshold in the period $t+1$, the lender always lend and normal borrowers always repay in period t . After that phase, lenders mix and borrowers default with increasing probability if they get a loan. Bayesian updating and optimization also impose two strong restrictions: Since only normal borrowers default, after a default the borrower's type is revealed and players never lend or repay after that. And after a later period in which there is no loan the borrower missed an opportunity to improve her reputation so players should never lend or repay after that.

Jung, Kagel and Levin (1994) ran experiments with a 'chain-store' entry deterrence game, which is a workhorse in industrial organization modeling. Their design closely followed Camerer and Weigelt's and the game-theoretic models. Table 1 shows payoffs.

¹³In sessions 4-6, the lender's default payoff was -50. In sessions 7-8, it was -75.

The entrants prefer to play “in” (entering) if met by sharing behavior, earning 150, but entrants earn the least if the incumbent fights (80) and earn an intermediate payoff of 95 from staying out. A “normal” incumbent earns 300 if the entrant is out, 160 if the entrant plays in and she plays share, and earns only 90 if she fights. Fighter-type incumbents have share and fight payoffs reversed. Payoffs were increased for entrants in two sessions.¹⁴ With these parameters, the sequential equilibrium is very much like the one in the trust game: Fighting for a couple of periods (and entrants wisely staying out) followed by mixing, with an increasing tendency to share toward later periods.

The teaching model is compared below with an agent-quantal response equilibrium (AQRE) model that incorporates two behavioral features. In AQRE, players noisily best respond, and some “homemade prior” fraction of players always behave as Stackelberg types—honestly in trust games, or fighting in entry games. The homemade prior is necessary to explain why reputation-building occurs (even in the last period) in finitely-repeated games where the experimental design did not explicitly induce incomplete information, so unraveling should occur in theory. Earlier experiments imputed a value for the homemade prior (Camerer and Weigelt (1988a); Neral and Ochs (1992)) or estimated it from a structural model (Palfrey and Rosenthal (1988); McKelvey and Palfrey (1992)). We *measure* the frequency of honest or fighter types in two experimental sessions with one-shot games with random rematching. In these games, there is no reputational incentive for behaving honestly or fighting. The measured rate of those behaviors is then used to constrain their frequency in the repeated game estimation.

The agent-based version of quantal response equilibrium was introduced to fit suitable for extensive-form games (see McKelvey and Palfrey (1998)). In the agent-based form players choose a distribution of strategies at each node, rather than using a distribution over all history-dependent strategies. This model is more suitable than sequential equilibrium because the (intuitive) sequential equilibrium predicts that many events have zero probability, so some notion of error or trembling is needed to fit the data. Agent-QRE is a plausible benchmark and fits many other data sets well (e.g., McKelvey and Palfrey (1998); Goeree and Holt (2001); Ho, Camerer and Weigelt (1998)). The AQRE model is implemented with four parameters—three different response sensitivities (λ 's)

¹⁴In sessions 5-6, entrant payoffs were 150 for in-fight, 475 for in-share, and 300 for out.

for lenders, honest borrowers, and dishonest borrowers, and a perceived prior belief of lenders about $P(\text{honest})$ or $P(\text{fighter})$. To allow the model flexibility, especially when the induced $P(\text{honest})=0$ and no lending or repayment is predicted, we allow a homemade prior θ . It is noteworthy that AQRE estimation is more computationally challenging than in any previous applications because there are $3^8 = 6561$ paths through the game (see Camerer, Ho and Chong (2002) for details).

4 Measuring the Homemade Prior θ

Earlier trust and entry experiments showed that even when the induced fraction of honest borrowers or fighter types are zero, there is substantial repayment and fighting (even in the last period). Inspired by the “gang of four” model of Kreps and Wilson (1982), Camerer and Weigelt (1988a) suggested this was due to the presence of an endogenous fraction of subjects who, despite monetary incentives to default, simply preferred to repay— a “homemade prior”. Palfrey and Rosenthal (1988) used the same idea to explain contribution in public goods games. More recently, the same intuition has been formalized in models of social preference used to explain contribution (and punishment) in public good games, reciprocity, rejections of ultimatum offers, and so forth (e.g., Camerer (2002) chapter 2).

In Camerer, Ho and Chong (2002) we estimated the AQRE homemade prior θ from the data.¹⁵ The resulting estimates were high— from .5 to 1— compared to the values around .1-.2 suggested by early experiments.

Since the homemade prior is intimately tied to the extent of repaying or fighting, it is important to estimate it precisely and plausibly. By definition, honest or fighter types will repay or fight even in one-shot games. Therefore, we measured θ by conducting two

¹⁵An earlier version of this paper at <http://www.fba.nus.edu.sg/depart/mk/bizcjk/teach.htm> reports details of estimation in all sessions where θ can take any value in the AQRE model, and is not included in the teaching model. When θ is restricted as we discuss below and included in teaching, the AQRE model fit suffers a little and the teaching model fit improves a little, but the basic pattern of results and parameter values are not changed much.

experimental sessions of one-shot games, reproducing the original experimental conditions as closely as possible while generating enough data for a reliable estimate. One session used the most common payoff structure in trust games and the other session used the most common structure in entry games.¹⁶ Each session used 12 subjects playing two blocks of 6 rounds in a fixed-role protocol (as in the original experiments). In each block of six rounds, each borrower was matched with each other lender once in a “zipper” design. Each borrower therefore plays the same lender twice, but never knows which lender she is playing. A total of 72 single-shot games were played in each experimental session. Since the crucial behavior is repayment by borrowers, we used the “strategy method” in which borrowers chose whether to repay or default *before* knowing whether they received a loan. (Otherwise, potential repayments are only observed when lenders lend, which limits the sample.)

Dollar payments were the original payments, adjusted upward for inflation.¹⁷ In trust games there were 17 repayments (26%) and in entry games there were 11 fight choices (18%), percentages which are remarkably close to the 17% figure originally imputed by Camerer and Weigelt (1988). Another way to measure types is to ask what fraction of subjects usually repaid. Two subjects (2 of 12, or 16%) repaid about half the time or more, giving a similar measure of θ .¹⁸

Because the samples are modest, it makes sense to restrict θ in the teaching and AQRE models to be close to the value measured in these one-shot games but allow some leeway. So θ is restricted to lie in a 95% confidence interval around the estimated value, (.19,.29) for trust and (.11,.20) for entry.

¹⁶The lender’s payoff used was -50 when the borrower reneges. This payoff is identical to trust data sessions 6-8 where $p = 0.1$ and new trust data sessions 1-7 where $p = 0.1$. The entrant’s payoff used was 80 when the weak monopolist fights in market-entry games. This corresponds to market entry game sessions 1-3 (inexperienced) and 6 (experienced) where $p = 1/3$.

¹⁷We assumed the original experiments were in 1986 and 1990 and adjusted payments by the GDP deflator, increasing them by 50% and 23% respectively.

¹⁸In trust games, the numbers of “repay” in 12 rounds across borrower subjects were 0,0,1,1,1,5,9. In market entry games, they were 0,0,2,3,3,3.

5 Data and Results

The central contribution of this paper is fitting fEWA-teaching and AQRE models to experimental data from three sources. The first source is eight experimental sessions of a repeated borrower-lender trust game reported by Camerer and Weigelt (1988a). The second source is an unpublished sample of eight more sessions of the same game (with $P(\text{honest}) = .10$) in which players also report beliefs about whether the borrower will default if there is a loan (Camerer and Weigelt (1988b)). These data are called “new trust” games. The third source is 12 sessions of an entry-deterrence game from Jung, Kagel and Levin (1994). Eight of the sessions use “inexperienced” subjects (participating in that particular game for the first time) and four use experienced subjects who returned for a second session playing the same game. There are a total of 28 experimental sessions, roughly 2,000 8-period sequences and 26,000 choices.

Subjects in the trust games were either MBA students at NYU (in the original data) or undergraduates at the University of Pennsylvania (in the new trust data). They were paid an average of \$18 for a 2-1/2 hour session. Instructions are available in Camerer and Weigelt (1988). Subjects in the entry-deterrence games were University of Pittsburgh undergraduates. See Jung, Kagel and Levin (1994) for design details. Each session had 4-101 eight-period sequences. In each trust session, there were 11 subjects, three borrowers and eight lenders. In each entry-deterrence session, there were 7 subjects, three monopolists and four entrants.

5.1 Trust games

Typical patterns in the old trust data can be seen in Figures 1a-b (pooling across all sessions to reduce sampling error). The figures show relative frequencies of not lending (all data) and default (conditional on lending, for dishonest borrowers only), assuming there was no default earlier in the sequence. Sequences are combined into ten-sequence blocks (denoted “sequence” in the figures) and average frequencies are reported from those blocks. Periods 1,...,8 denote periods in each sequence.

Two patterns in the data are of primary interest. First, what is the rate of lending across periods (and how does it change across sequences)? Second, how do borrowers respond to loans in different periods (and how do these responses vary across sequences)? Figure 1a-b show that lenders start by generally making loans (i.e., low frequency of no-loan) in early periods, then learn to rarely loan (i.e., high frequency of no-loan) in periods 7-8. Borrowers rarely default in early periods, but frequently default in periods 7-8. The pattern of increasing default in later periods is particularly dramatic in later sequences.

How well the models fit is judged from three angles: Graphs that show predicted frequencies corresponding to the data in Figures 1a-b; overall statistics measuring fit (hit rate and log likelihood); and reported parameter values.

Figures 1c-d show frequencies predicted by the teaching model. Figures 1e-f show corresponding predictions for the AQRE model. The predictions come from estimation with session-specific parameters. Note well that while the model makes separate predictions for *every* choice, the predictions are averaged in figures 1c-f to correspond to the pooled relative frequencies in the data.

The teaching model captures the rise in the frequencies of not lending and default across periods within each sequence, and picks up a little of the drop in no-lending across sequences. AQRE predicts hardly any trend (slightly in the wrong direction). It basically fits best by assuming lenders lend almost randomly (approximately 40% not-loan rate) and that default is rare and, counterfactually, *falls* across periods within a sequence.

Figures 2a-f show frequencies for the eight new trust sessions. Neither model is as visually impressive as in Figure 1. The teaching model reproduces a small rise in not-lending and default across periods a little better than AQRE does.

Table 2 summarizes log likelihoods (LL) and hit rates (i.e., the fraction of times that the choice predicted to be more likely by the model is the choice subjects made) from various models.¹⁹ These statistics are reported separately for in-sample calibration and

¹⁹These fit statistics, along with the average predicted probability of the actual choices, are reported separately for the dishonest and honest borrower, and lender categories in the working paper on our website <http://www.fba.nus.edu.sg/depart/mk/bizcjk/teach.htm>. The estimated $A_i^{IN}(0)$ and λ parameters that are suppressed in the tables below are also reported in that working paper.

out-of-sample validation. Keep in mind that the teaching model has more estimated parameters than AQRE²⁰, so the accuracy of out-of-sample forecasting is the better test of which model predicts best.

Results are reported for the total fit statistics (summed across sessions) when parameters are session-specific (“individual”), and when parameters are common (“pooled”) for both functional EWA and parametric EWA. The teaching model usually predicts out-of-sample a little better than AQRE based on hit rate and LL, except in the new trust data where they are about equally accurate. Note also that the fEWA teaching model is almost as accurate when all sessions are pooled, with common parameters, compared to when fit statistics from session-specific estimation are totaled up. In fact, the difference between out-of-sample LL for the session-specific estimates added up (the column “Total”) and when common parameters are imposed (“Pooled”) is only 38 and 73 likelihood points for fEWA in the two data sets, although 56 fewer parameters are estimated when data are pooled. This is an important clue that the parameter estimates are quite stable across sessions for the teaching model. The teaching model also usually has a higher (or equal) hit-rate out-of-sample than in-sample while the out-of-sample hit rate usually falls for AQRE. This is a sign that the teaching model is not overfitting.

Table 2 also shows fit of two other benchmark models pooling across sessions. One is the parametric EWA model in which ϕ, κ, δ and $N(0)$ are free parameters rather than generated from functions as in fEWA. Being able to estimate free parameters generates a substantial predictive improvement over fEWA. The other benchmark is a restriction of AQRE with the homemade prior θ fixed to zero. This restriction hurts the likelihood a lot and creates hit rates around 50% (chance), which is a reminder that the homemade prior is crucial to the success of AQRE.

Table 3 gives the minimum and maximum parameter values from individual session-

²⁰The teaching model has an initial attraction for lending $A(0)$, four λ 's (for the lender, two types of myopic borrowers, and borrower teachers), θ, τ and α , a total of 8 parameters. AQRE has the homemade prior θ , and three λ 's (for lenders and the two types of borrowers), a total of 4 parameters. Since there are plenty of data to work with, we elected to have lots of freedom in λ 's for both models. Note that the very small loss in fit from pooling, particularly for teaching, suggests many of these parameters have similar values across sessions so we could often restrict them a priori to economize on degrees of freedom.

specific estimation (see tables in Appendix for details) and estimates from pooling.²¹ The averaged functional values of ϕ , δ and κ are extremely consistent across sessions. They are in the ballpark of the values estimated in parametric EWA, except that the functional ϕ is always too high (.76-.77 compared to unconstrained estimates of .45 and .25). The cross-sequence learning strength τ is close to one in most sessions, and is usually around .90 when sessions are pooled. The estimated percentages of teachers α generally range from 70-100% (except for a few low outliers), and are almost the same, .93 and .91, when new and old trust sessions are pooled. The fact that pooling across sessions degrades overall fit only a little (as Table 2 shows), and parameters are consistent across the new and old trust data sets, is very encouraging. These results suggest that economizing on parameters by restricting $\tau = \alpha = 1$ will yield a good approximation in games like these.

Recall that the estimate of θ , the homemade prior parameter, is allowed to vary within the 95% confidence interval around our estimate from the one-shot experiments. The pooled estimates in trust games are in the middle of the interval (.23-.24) for teaching and at the upper bound of the interval (.29) for AQRE. This result is consistent with the fact that estimated $\hat{\theta}$'s are very large (often close to one) when they are unconstrained (see our working paper); so the AQRE model chooses the largest value it can when restricted. Since most of the observations in the data are loans and repayment, the AQRE model fits best by simply assuming many players prefer to repay (and since it is an equilibrium model, lenders know that and lend). As a result, it does not generate the sharp rise in default rates at the end of the sequences which is evident in the data (Figures 1-2).

5.2 Entry-deterrence games

Now turn to the Jung, Kagel and Levin data on entry-deterrence. Since they ran eight sessions with inexperienced subjects and four with experienced subjects, we can see whether subjects grow more sophisticated when they repeat an entire experimental session.

Start with the data from inexperienced subjects (Figures 3a-f). Focus on rates of entry and sharing, which are predicted to start low and rise as the end of a sequence draws

²¹Keep in mind that in fEWA ϕ , δ and κ are *not* estimated, they are functions of the data; the numbers in parentheses below each estimate are the standard deviations of the functional values across periods.

near. Entry and sharing by inexperienced subjects are far too frequent in early periods but there is some convergence toward early entry-deterrence across the experimental session. Entrants just didn't quite figure out how much it pays to fight entry in early periods. Both models are visually unimpressive. The AQRE model fits these data best with roughly flat predictions across periods and sequences. Teaching predicts a modest rise in sharing across periods.

Figures 4a-f show data from experienced subjects. The correspondence of behavior to equilibrium is much more dramatic. In the first sequence block, players often enter in the first 3 periods, but they quickly learn early entry is rarely met with sharing, and they stay out in early periods of later sequences. Both AQRE and teaching capture the downward trends in fighting over periods (although in a muted way, compared to the data) and the teaching model also predicts a modest increase in entry across periods.

Summary statistics in Table 2 shows that the two models are about equally accurate for inexperienced subjects. For experienced subjects, fEWA teaching is much more accurate than AQRE, especially in hit rate. In fact, the hit rates for AQRE are typically *worse* than chance.²²

Table 3 shows parameter values. As in the trust data, the fEWA entrant parameters are consistent across sessions. The functional ϕ is again too high compared to the parametric EWA estimate, and some of the parametric EWA estimates are extreme in the inexperienced sample. The cross-sequence learning strength τ is usually around one but is sometimes very low; when sessions are pooled the estimates are .98 and .48 respectively. We suspect τ is not always well-identified in an econometric sense, but this also means it can be set to one without harming fit.²³ The estimated fraction of teachers α is smaller

²²Figures 4a and 4e show why. The probability of entry predicted by AQRE is just slightly above .5 in most sequences and periods. This means the model is 'always' predicting entry if it has to predict one choice or the other (as the hit rate measure forces it to). But Figure 4a shows that the overall frequency of entry is below a half. Thus, the model is not far off in matching the overall rate of entry, but its prediction is very often on the 'wrong side' of 50%.

²³The parameter τ will not be well-identified when experience in earlier periods in a sequence is the same as experience in the current period in previous sequences. For example, in the trust games the modal behavior is lending and repayment, until the last two periods. Since both sorts of experience are the same, the model cannot tell whether a player is looking mostly at previous *periods* (low τ) or

for inexperienced subjects (.54) than for experienced subjects (.88). This increase in sophistication is also observed by Stahl (1999) in matrix games, Camerer, Ho and Chong (2002) in dominance-solvable games, and Cooper and Kagel (2001) in signaling games.

A crucial test for both teaching and equilibrium models is whether similar parameter values can be used to explain behavior in trust games and entry-deterrence games. These games are opposite in incentive structure in the sense that Stackelberg behavior is mutually-beneficial in trust games but only privately-beneficial in entry-deterrence. If the same general model structure and parameters can explain both games that is very encouraging for broader application. In fact, the trust and experienced entry data give very similar values of α (.88-.93), ϕ (.76-.77), δ (.52-.53), κ (.60-.65), though τ is much lower in entry than in trust.

Many results are consistent across both games. Hit rates are around 70%. (Keep in mind that since equilibrium and teaching models both predict mixing, the hit rates of those models will always be below 100% even if the models actually generated the data.²⁴) The equilibrium models predict rather well, but they are helped substantially by allowing the constrained homemade prior. Restricting $\theta = 0$ degrades fit a lot, yielding hit rates which are often worse than chance guessing.

In terms of overall fit, teaching is generally a little more accurate except on the sample of inexperienced entry games, where both models fit equally well. However, a key point is that the models are closely related because the teaching borrowers and incumbent firms are exhibiting sophistication and strategic foresight, just as in equilibrium models. The main difference between the two models is that lenders and entrants learn in the teaching model but they anticipate what borrowers and incumbents will do in AQRE. at previous sequences (high τ) except in the last period or two. But identification also means that restricting $\tau = 1$ will lead to approximately the same predictive accuracy as estimating it. Indeed, restricting $\tau = 1$ only degrades in-sample LL for the four fEWA pooled estimates (left to right in Table 2) by 2, 20, 3, and 3 points.

²⁴For example, suppose the sequential equilibrium model with $\theta = 0$ actually generated the data, and that model's conditional predictions are used to make guesses. Then when $p=1/3$ the expected hit rate is 92.3%. The reason the hit rate is less than 100% is that in most of the possible sequences there is one misprediction ("a miss") because either the borrower defaults early or the lender stops lending (which the model does not predict).

The fact that the teaching model generally fits and predicts better than AQRE means that weakening sophistication of ‘short-run’ players is empirically useful.

5.3 Distinguishing teaching and AQRE models: Missed opportunity, forgiveness, and rebuilding

The equilibrium type-based view and the strategic teaching model make similar predictions and have similar comparative static properties, but can be distinguished in a couple of subtle ways.

Equilibrium models with best response predict a discontinuous sensitivity to the prior $P(\text{Honest})$ — if it is too low there is no reputation-building. These models also predict, curiously, that a player’s mixture probabilities depend on the *other* player’s payoffs but not on the mixing player’s payoffs. (Neral and Ochs (1992) report evidence against this assumption.) The AQRE model with a homemade prior does not make these counterintuitive predictions (and neither does teaching).

One way to validate the teaching model is to see whether recovered estimates of the “amount of teaching”— measured by α — correspond to visually observable patterns of lending and default. If there is little teaching (low α) there should be little lending and rampant default, and vice versa for lots of teaching (high α). Figure 5 shows the rates of not-lending and default for new trust sessions separated by values of α .²⁵ There are striking differences in the two cases which means the estimated α has some explanatory power in dividing sessions into those with low and high reputation-building.

Another way illustrates more subtler differences. In a type-based equilibrium model with best response, missing a chance to fight entry (or pay back a loan) in later periods leads *immediately* to entry and sharing (or not lending) in all subsequent periods. But the same effect of “missed opportunity” is *not* present in early periods. The reason for these effects is that the threshold $P(\text{honest})$ required to induce the lender to lend is always rising over time. If an incumbent misses a chance to *increase* her reputation

²⁵We separate the sessions in the New Trust data into two separate sets of plots; one for the high- α and the other for the low- α . We set $\alpha = 0.7$ as the cut-off point.

(i.e., the perceived $P(\text{honest})$) by repaying when she could have defaulted, the prior is not updated. In early periods the prior is above the threshold belief lenders require so missed opportunity makes no difference. But in later periods missing a chance to build reputation is fatal; on the equilibrium path, the posterior after in period t is below the belief required in period $t+1$, so the lender won't lend in period $t+1$ or any future periods.

Another feature of optimizing Bayesian models is “no forgiveness”. In the type-based models a player who defaults, for example, has clearly revealed her type. The perceived $P(\text{Honest})$ should plunge to zero, and nobody lends afterwards. Since it doesn't pay to try to rebuild one's shattered reputation, optimizing Bayesian models also predict that we will never see repayment after a default. In contrast, strategic teaching *does* allow lenders to smoothly adjust attractions so that the “bad reputation” of the lending strategy after a default could lead to more lending, and may justify subsequent repayment.

While the sequential equilibrium with best response predicts effects of missed opportunity, no forgiveness, and no rebuilding, the AQRE predictions about these effects are sensitive to parameter values. So we'll look at whether these patterns exist in the data, then compare how well teaching and AQRE models account for reputation rebuilding in the data.

To measure missed opportunity, Figure 6a plots the conditional relative frequencies (from data) of lending in period t conditional on not lending in period $t-1$. These conditional frequencies are shown across periods and sequences for the new trust data in 20-period sequence blocks.²⁶ The conditional frequencies clearly rise over time, from about $1/3$ in period 2 to $2/3$ in period 8.²⁷

Forgiveness is measured by calculating the probability of lending in period $t+1$, given a default in period t . Figure 6b shows the patterns in the new trust data. There is some

²⁶Only the new trust data are shown to save space. The comparisons never favor AQRE strongly over teaching, though often the two models are about equally accurate. Other plots can be seen on our website <http://www.fba.nus.edu.sg/depart/mk/bizcjk/teach.htm>.

²⁷The relative frequencies of these conditional events predicted by the teaching model also rise, but less than the data (from about 40% to 60%). However, the predicted frequencies from AQRE are even lower in magnitude and flatter over time. So the teaching model does a slightly better job of accounting for the missed opportunity effect and its increase across periods.

lending following default in the first two blocks of sequences but in later sequence blocks lending following default is rare, especially after periods 2 or 3.²⁸

An interesting feature of the data is whether borrowers who get a subsequent loan after default— i.e. who *are* forgiven by lenders— repay or not. Since the posterior $P(\text{Honest})=0$ after a default in optimizing equilibrium models, and the predicted $P(\text{repay})=0$ ²⁹, borrowers should never try to rebuild their shattered reputations. Since AQRE allows trembling, this sharp prediction may not hold so let’s see what the data tells us and whether the models match. Figure 7 shows the conditional repayment rate in t after a default in $t-1$. (Negative values on the 3-axis plot cells suggests no data and should not be confused with true zeros.) Borrowers clearly do try to rebuild trust, occasionally, but only in the first couple of periods. Both models predict a decline in rebuilding, but the AQRE model severely overestimates the rate of rebuilding while teaching gets it about right. Although both models overestimate the overall rate of conditional repayment. The corresponding LL measures for Figures 7b (teaching) and 7c (AQRE) are -3.29 and -3.82 respectively.³⁰ The differences in accounting for reputation rebuilding across models are modest although the teaching model does a measurably better fitting job than AQRE.

A better way to distinguish the equilibrium type-based approaches from teaching is worth mentioning, though executing it lies beyond the scope of this paper. Even the AQRE Bayesian approach is sensitive to how types are determined and what is known about them. In the standard approach, a player’s type in one period of a sequence is perfectly correlated with her “earlier” type. A different design breaks the correlation by drawing types independently each period with a fixed probability. When types are ‘refreshed’ like this, the Bayesian link between a player’s past behavior and her likely future behavior is severed. AQRE will allow stochastic fighting and lending, but it will make the same prediction in each period of a sequence. The strategic teaching

²⁸The AQRE model is slightly closer to the overall rate of lending and both models pick up a little of the decline in forgiveness in later periods.

²⁹If $P(\text{repay in } t \mid \text{default in } t-1) > 0$ for normal types, then by Bayes’ rule $P(\text{Honest in } t \mid \text{repay in } t)$ is still zero and borrowers won’t get more loans, so it doesn’t pay to try to rebuild by repaying.

³⁰The LL for period t sequence k is given by $ef_{tk} \cdot \ln(pf_{tk})$ where ef is the empirical frequency and pf is the predicted frequency. We sum across all periods and all sequences to give the LL measure for the plots.

model, in contrast, allows reputation-building even when types are refreshed, and still predicts within-sequence differences— viz., a breakdown of reputation-building toward the end of a sequence. A similar difference in type-based and teaching predictions arises when players have types, but the types are not commonly-known so that incumbents aren't sure whether reputation-building is worthwhile.³¹ Further experiments exploiting these sensitivities to information about types could be useful in contrasting the teaching and equilibrium approaches. But since the theories' differences depend on having no 'homemade' types, controlling or measuring such types is crucial to get a sharp test.

6 Measuring beliefs

A direct way to judge game theories which specify beliefs is simply to measure beliefs and compare them to those proposed by various refinement criteria. Banks, Camerer and Porter (1988) and Camerer and Weigelt (1988b) were among the earliest experimental economists to do so (inspired by McKelvey and Page (1990)). More recently, Nyarko and Schotter (2002) compared reported beliefs with those predicted by fictitious play learning; the two types of beliefs are not closely related. Manski (2002) also notes that in some games, econometric identification based on choices alone is poor, but can be improved by measuring beliefs or other data beyond choices (cf. Camerer et al. (1993); Costa-Gomes, Crawford and Broseta (2001); Johnson et al. (2002)).

In the new trust data, beliefs were measured using an incentive-compatible quadratic scoring rule. In every period seven lenders and two borrowers are sitting idle and watching the other subjects. Since these subjects have time on their hands, if a loan was made we asked two borrowers and two lenders to report a belief $P(\text{default})$, which was rewarded with an incentive-compatible quadratic scoring rule.³² This section analyzes

³¹Drew Fudenberg suggested this subtle point.

³²If a reporter stated a belief p , she earned an affine transformation of the payoffs $2p - p^2$ if default occurred and $1 - p^2$ otherwise. It is well-known that (i) this scheme induces risk-neutral players to report their true beliefs; (ii) if players are risk-averse their reports will be biased toward .5; and (iii) the expected payoff loss from small deviations around the true belief is extremely flat (e.g., Camerer (1995) p. 592-3). Given the bias and flat-maximum problems, it is all the more remarkable that reports are as

some properties of reported beliefs and compares them to corresponding beliefs implied by models.

To reduce sampling error, beliefs from all eight new trust sessions are pooled for the analysis. We first compare reported beliefs (averaged over the two borrowers only) with the implicit $P(\text{default})$ in the AQRE and teaching models.³³ Figure 8 shows a scatter plot of implied beliefs and borrower beliefs. The correlations are low the implied model beliefs are usually severely restricted in range (especially for the teaching model) while reported beliefs are more widely dispersed.

Since implied and reported beliefs are weakly correlated, a crucial question is which type of belief are a better guess of when defaults actually occur. It could be that the reports are noisy and inaccurate. Or the reports might be more accurate than the model beliefs, which would be evidence that the model is misspecified in a way that can potentially be improved in future work (cf. Nyarko and Schotter (2002)).

The accuracy of forecasts can be assessed with a calibration diagram. To measure calibration, implied and reported beliefs were grouped into probability categories (“bins”), rounding down to the nearest increment of .05. Denote the i -th forecast bin by $i \in \{0, .05, .10, \dots, .95, 1\}$. There are 21 such bins and the prediction in the i -th bin is just $P(\text{default}) = i$. Denote the total frequency of forecasts of i by N_i , the total sample size of all forecasts by $N \equiv \sum_{i=0}^1 N_i$, and the relative frequency of actual defaults in category i by f_i . Calibration is assessed by comparing each categorical probability with the relative frequency of defaults in all the cases where the stated probability fell in that category.

well-calibrated and discriminated as they are.

³³The belief is of course from the perspective of the lender, and the other subjects, who do not know the borrower’s type. Also, in the teaching model there is different borrower behavior by myopic, honest and homemade-honest, and teaching types. We take the perspective of an ‘outside’ borrower subject who knows the structure of the model, averages the different $P(\text{default})$ values each period with the correct relative frequencies of types, and Bayesian updates across periods in a sequence. Note that these implied beliefs do *not* usually correspond to lender’s beliefs because in fEWA learning the lender attractions do not usually correspond to *any* beliefs (just as a player who cooperates in a one-shot PD game behaves in a way that is not consistent with any beliefs.)

Figure 9 is a series of bubble charts plotting i (x-axis) against f_i (y-axis) for implied beliefs and for reported beliefs (averaged across the four reporters). The bubble size is proportional to $\sqrt{N_i}$ to show how relatively common different belief reports are. The data are pooled across borrowers and lenders³⁴, and across all periods, because there were no significant differences across time or subject types (using Chow tests for pooling).³⁵

Reported beliefs are *very* closely related to the actual frequencies of defaults. Implied beliefs are strongly correlated to actual frequencies but the slope of the regression of f_i on i is greater than one. Table 4 reports the estimated intercept and slope coefficients in weighted regressions of default rates against reported beliefs.³⁶ The intercepts are small and estimated slopes are close to one.

To parse accuracy of beliefs further, it helps to distinguish two properties of probability judgments, conventionally called "calibration" and "discrimination" (e.g., Yates (1990); Feltovich (2000)). Well-calibrated forecasts are those that match aggregate results well (i.e., i and f_i are close). For example, if subjects said there is a 20% chance of default 141 times (as they did), then there should actually be default in roughly 20% of the 141 cases, or 28 times. Visually, calibration is the distance between the points in Figure 9 and the identity line and is usually measured by the sum of the weighted squared distances between the predictions and their associated relative frequencies, $\sum_i N_i(i - f_i)^2 / N$. Calibration measures are reported in Table 4. Calibration is generally extremely good (the measure is small) and averaged reported beliefs are better calibrated than implied beliefs.

³⁴Borrower and lender beliefs are about equally accurate. This rules out a hunch we had, that borrower beliefs are more informative or accurate because borrowers are better than lenders at forecasting the behavior of other subjects like themselves (since it is other borrowers who choose whether to default). However, borrowers did use extreme categories, beliefs of zero or .05, about twice as often as lenders, and the correlation of the two borrowers' beliefs was much higher than the between-lender correlation, so there is some asymmetry which might have an interesting cognitive basis.

³⁵Averaging makes a substantial difference because there is a lot of disagreement across different reporters (the correlations between any two of the four belief reports, "interrater reliability" in psychometric language, range from .43 to .69).

³⁶The regression statistics reported in Figure 9 are unweighted. The weighted regressions use the square root of each point's sample size N_i as a weight, effectively weighting residual variance by the reciprocal of $\sqrt{N_i}$.

However, well-calibrated forecasts can contain little information. For example, simply guessing the base-rate of overall default (which is 19.6%) all the time leads to perfect calibration— all events are in the “19.6” bin and their overall frequency is 19.6% (this model is reported in Table 4 as a benchmark). But these judgments convey nothing about the variation in $P(\text{default})$ across periods. The amount of information in a sample of forecasts is measured by “discrimination”, the degree to which the actual default frequencies in the various forecasting categories vary from the overall default rate. Define the overall frequency as $\bar{f} \equiv \sum_i N_i f_i / N$. Then discrimination is measured by $\sum_i N_i (f_i - \bar{f})^2 / N$. (A high number is better.) Table 4 shows that the discrimination of implied beliefs is substantially worse (teaching) and a little worse (AQRE) than for belief reports.

An overall measure of statistical accuracy is the “probability score” of a set of beliefs. The probability score is simply the average of the squared deviation between each belief and one or zero (for events that occurred or didn’t, respectively). It is well-known that the average probability score is equal to calibration minus discrimination plus $\bar{f}(1 - \bar{f})$. Table 4 shows that the probability score for the implied beliefs is worse than for the averaged reports, although all the beliefs are only modest improvements over guessing the base rate every time.

These belief reports are scientifically useful in two ways (and should perhaps be used more often). First, the fact that implied model beliefs are much more strongly clustered in the middle of the range, and more poorly calibrated and discriminating than the reported beliefs, means the model is missing something which subjects themselves understand. The overall probability score favors the AQRE model over teaching but neither model is really impressive compared to always guessing the base rate.

Second, there is a large literature on calibration in judgment research (e.g., Yates (1990); Camerer (1995) p. 592). Most lab studies use college students or other regular folks who are asked general knowledge questions. For example, Lichtenstein, Fischhoff and Phillips (1982) summarize many studies in which subjects were asked a question with two possible answers, such as “Which magazine had the largest circulation in 1970, *Playboy* or *Time*?”. Subjects gave an answer *and* a probability that their answer was correct. Reported beliefs usually exhibit overconfidence: When subjects said they were

sure they were right (they were 100% confident) they were right only 80% of the time. Other studies showed that confidence intervals for numerical quantities are typically too narrow.³⁷

Most of these early studies did not use incentive-compatible scoring rules to reward subjects, and dramatize how big a mistake it is to wrongly state an extreme probability. However, many studies examined calibration of domain experts who are professionally trained and usually have direct or indirect incentives to forecast accurately. Since it is often easy to get field data on actual forecasts, studies have been done with weather forecasters (Murphy and Winkler (1977)), blackjack dealers (Keren (1987)), and market prices for orange juice and horse racing bets (Roll (1984); Ali (1979); Hausch, Lo and Ziemba (1994); Johnson and Bruce (2001)). Most of these studies show superb calibration. Our study is the first to show calibration can be extremely accurate—people are *not* overconfident—when an incentive-compatible reward structure is used, and subjects become experts over time about their own collective behavior in an artificial domain.

7 Conclusion

Models of individual learning have been actively used to fit data from many game theory experiments and understand their regularity. Almost all these models are *adaptive*—i.e., players do not explicitly account for the possibility that other players are learning. This paper adds “sophisticated” players who realize others are learning, in repeated games with incomplete information. Sophisticated players who know they are playing a repeated game have an incentive to take actions which are costly in the short-run, but which “teach” learners what to expect, in a way that benefits the teachers. Neglecting these teaching effects severely limits the application of learning models outside the lab, to domains in which economic relationships are long-lasting.

This paper specifies a precise model of strategic teaching and applies it to finitely-

³⁷One criticism of early studies is that subjects were essentially tricked because of the way the general knowledge questions were sampled. But sampling questions randomly produced similar results (Soll (1996)).

repeated experimental games of trust and deterrence of entry by incumbents. The games have incomplete information because the experimenter induced some fraction of borrowers and incumbents to always behave in a trustworthy or aggressive (entry-fighting) way. Previous experiments have shown that some features of behavior in these games is approximated by very complex and delicate equilibria (Camerer and Weigelt (1988a); Jung, Kagel and Levin (1994)). But it is unlikely that players approximate the equilibria by introspection, and their comparative static predictions are often wrong (Neral and Ochs, 1992). A boundedly rational model of learning which can explain how equilibration comes about would therefore be useful.

In the teaching model the lenders and entrants learn adaptively (using a parsimonious functional EWA rule which sophisticated players know). Some fraction of borrowers and incumbents are myopic and some fraction always behave honestly or aggressively. The others teach strategically by anticipating how lenders and entrants learn, and making choices which maximize long-run payoffs when the impact of their choices on learning of others is accounted for. In the equilibrium approaches the *players* build reputations (“this guys seems honest”); in the teaching model the learners’ *strategies* have reputations (“entry is dangerous”). Assuming updated types is designed to explain why lenders are willing to lend early and afraid to lend late. But learning and teaching can produce the same pattern without types: Teaching borrowers will pay back early so that lenders learn to lend; and cross-sequence learning means lenders learn that lending in late periods is hazardous.

The teaching model was fit to 28 sessions of data from both repeated trust and entry games (including eight new sessions), more than 25,000 choices. A key fact is that the same model can account for quite different behavior in these games: Borrowers in trust games behave in a trustworthy way that is mutually-beneficial, while “fighter” incumbents benefit only themselves. The same model explains both because the two opposite behaviors emerge endogenously from the interaction of teaching and payoffs. The models also include a fraction of ‘homemade’ types who always behave honestly or aggressively, but we constrain the estimated frequency of those types to be close to the frequency measured in a new experiment with one-shot games.

The teaching model correctly predicts about 75% of the choices and reproduces basic trends in the data (particularly increasing default and market-sharing in later periods of a sequence). One estimated parameter is τ , which measures how much learners in period t weigh experience in periods t from previous sequences, relative to experience in previous periods 1 through $t-1$ in the current sequence. Since τ is typically estimated to be around one, it could be safely restricted to one in future work.

The other key parameter is the fraction of strategic teachers, α , which is reliably estimated to be about .5 for inexperienced entry-game subjects and .9 for trust games and experienced entry-game data. The fact that α rises with subject experience corroborates other findings. A benchmark model of agent-based quantal response equilibrium (AQRE) is also estimated—the most complex QRE application so far—and fits reasonably well, though appreciably worse than the teaching model (except for inexperienced entry game data).

A second contribution is that we contrast subtle differences between teaching and AQRE approaches about “missed opportunity”, “no forgiveness”, and “no rebuilding”. In standard equilibrium models, missing an opportunity to build reputation later in the sequence leads immediately to unraveling. And players who have destroyed their reputations (e.g., defaulting or sharing after entry) will never get a chance to rebuild them and shouldn’t try to if they do. Data on the actual frequencies of these conditional events are slightly better matched by the teaching model than by AQRE.

A third contribution is measurement of subjects’ beliefs about the chance of default in eight new trust game sessions. Measuring beliefs (and other cognitive variables) can sharpen econometric identification. The reported beliefs are usually quite different (and more variable) than the beliefs implied by either model. The reported beliefs are more accurate, as conventionally measured by calibration and informativeness (“discrimination”), than beliefs implied by AQRE or teaching. Incidentally, reported beliefs exhibit *superb* calibration. Their terrific calibration is a reminder that the overconfidence in beliefs observed in many psychological studies can disappear when the agents reporting beliefs are incentivized to be accurate, and both create and become expert in the events they are forecasting.

Of course, the teaching approach will never fully replace equilibrium Bayesian types modelings because the latter is well-understood and insightful. But teaching shows promise for capturing much of the intuition and empirical regularity of reputation-building in repeated games in a boundedly rational way. Note also that the type-based models assume optimization and foresight by reputation-builders, and Bayesian updating of types by “learners”. The teaching model takes away *only* the last feature. Our overarching conclusion therefore echoes Cooper, Garvin and Kagel (1997):

game theory is at its foundation a hypothesis about how people behave. In particular, it posits that individuals will attempt to anticipate the behavior of others and respond accordingly. This is the soul of game theory, and the experiments indicate that it is alive and well. What may not be so healthy are the legion of assumptions which have been added to the theory in order to get tractable results. In the rush to get theories which give sensible outcomes and are easily employed by theorists, the reality of how people actually behave may have been left behind. We do not suggest that game theory be abandoned, but rather as a descriptive model that it needs to incorporate more fully how people actually behave. (P. 555)

There are many directions for new research. An obvious one is to apply the model to other repeated games. Another is to endogenize the τ and α parameters, or fix them to one and see what results theoretically. A third direction is to sharpen the distinction between teaching and equilibrium approaches with experiments in which players’ types are drawn independently in each period or players’ incomplete information about the presence of “crazy types” is weakened (strict equilibrium theories predict a collapse in reputation-building in these cases; teaching theories do not). A fourth direction is application of these models to repeated games of special economic interest. One is macroeconomic models in which policymakers set inflation rates in the presence of an expectational Phillips curve, knowing that the public forecasts inflation adaptively (Sargent (1999); Arifovic and Sargent (2001)). Another is price bubbles in finitely-lived assets (e.g., Smith, Suchanek and Williams (1988); cf. Lei, Noussair and Plott (2001)). We conjecture that such bubbles are explained by the interaction of adaptive and sophisticated traders.

References

- [1] Ali, Mukhtar M. (1979): “Some Evidence of the Efficiency of a Speculative Market,” *Econometrica*, 47, 387-392.
- [2] Andreoni, James and John Miller (1993): “Rational Cooperation in the Finitely Repeated Prisoner’s Dilemma: Experimental Evidence,” *Economic Journal*, 103, 570-585.
- [3] Arifovic and Sargent (2001): “Laboratory Experiments with an Expectational Philips Curve,” Stanford University working paper.
- [4] Arthur, Brian (1991): “Designing Economic Agents That Act Like Human Agents: A Behavioral Approach to Bounded Rationality,” *American Economic Review*, 81(2), 353-359.
- [5] Banks, Jeffrey, Colin F. Camerer, and David Porter (1989): “Experimental Tests of Nash Refinements in Signaling Games,” Presentation to Economic Science Association.
- [6] Cachon, Gerard P. and Colin F. Camerer (1996): “Loss-avoidance and Forward Induction in Experimental Coordination Games,” *The Quarterly Journal of Economics*, 111, 165-194.
- [7] 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.
- [8] Camerer, Colin F. (1997): “Progress in Behavioral Game Theory,” *Journal of Economic Perspectives*, 11, 167-188.
- [9] Camerer, Colin F. (2002): *Behavioral Game Theory: Experiments on Strategic Interaction*, Princeton: Princeton University Press.
- [10] Camerer, Colin F. and Teck-Hua Ho (1998): “EWA Learning in Normal-form Games: Probability Rules, Heterogeneity and Time Variation,” *Journal of Mathematical Psychology*, 42, 305-326.

- [11] Camerer, Colin F. and Teck-Hua Ho (1999a): "Experience-weighted attraction learning in games: Estimates from weak-link games," in *Games and Human Behavior: Essays in honor of Amnon Rapoport*, ed. by D. Budescu, I. Erev, and R. Zwick. New Jersey: Lawrence Erlbaum Assoc., Inc., 31-52.
- [12] Camerer, Colin F. and Teck-Hua Ho (1999b): "Experience-weighted Attraction Learning in Normal-form Games," *Econometrica*, 67, 827-874.
- [13] Camerer, Colin F., Teck-Hua Ho and Juin-Kuan Chong (2001): "Behavioral Game Theory: Thinking, Learning and Teaching," Nobel Symposium. at 2001.<http://www.hss.caltech.edu/camerer/camerer.html>
- [14] Camerer, Colin F., Teck-Hua Ho and Juin-Kuan Chong (2002): "Sophisticated EWA Learning and Strategic Teaching in Repeated Games," *Journal of Economic Theory*, forthcoming.
- [15] Camerer, Colin F., E. Johnson, S. Sen and T. Rymon (1993): "Cognition and Framing in Sequential Bargaining for Gains and Losses," in *Frontiers of Game Theory*, ed. by K Binmore, A. Kirman, and P. Tani. Cambridge: MIT Press, 27-48.
- [16] Camerer, Colin F. and Keith Weigelt (1988a): "Experimental Tests of A Sequential Equilibrium Reputation Model," *Econometrica*, 56, 1-36.
- [17] Camerer, Colin F. and Keith Weigelt (1988b): "Measuring Beliefs in Reputation Games," Presentation to the Economic Science Association.
- [18] Camerer, Colin F., George Loewenstein, and Drazen Prelec (2002): "Gray Matters: How Neuroscience can Inform Economics," under revision for *Journal of Economic Perspectives*.
- [19] Cheung, Yin-Wong and Daniel Friedman (1997): "Individual Learning in Normal Form Games: Some Laboratory Results," *Games and Economic Behavior*, 19, 46-76.
- [20] Clark, Kenneth and Martin Sefton (1999): "Matching Protocols in Experimental Games," University of Manchester Working Paper.

- [21] Cooper, David and John Kagel (2001): "Learning and Transfer in Signaling Games," working paper, Department of Economics, Case Western Reserve University.
- [22] Cooper, D., S. Garvin and J. Kagel (1997): "Adaptive Learning vs. Equilibrium Refinements in an Entry Limit Pricing Game," *Economic Journal*, 107, 553-575.
- [23] M. Costa-Gomes, V. Crawford and B. Broseta (2001): "Cognition and Behavior in Normal-form Games: An Experimental Study," *Econometrica*, 69, 1193-1235.
- [24] Crawford, Vincent P. (1995): "Adaptive Dynamics in Coordination Games," *Econometrica*, 63, 103-143.
- [25] Crick, Francis (1988): *What Mad Pursuit? A Personal View of Scientific Discovery*. New York: Basic Books.
- [26] Ellison, Glenn (1997): "Learning from Personal Experience: One Rational Guy and the Justification of Myopia," *Games and Economic Behavior*, 19, 180-210.
- [27] Feltovich, N. (2000): "Reinforcement-based vs Belief-based Learning Models in Experimental Asymmetric-information Games," *Econometrica*, 68, 605-641.
- [28] Fudenberg, D. and D. Kreps (1990): "A Theory of Learning, Experimentation and Equilibrium in Games," Stanford University working paper.
- [29] Fudenberg, Drew and David Levine (1989): "Reputation and Equilibrium Selection in Games with A Patient Player," *Econometrica*, 57, 759-778.
- [30] Fudenberg, Drew and David Levine (1998): *The Theory of Learning in Games*. Cambridge: MIT Press.
- [31] 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.
- [32] Harley, Calvin (1981): "Learning the Evolutionary Stable Strategies," *Journal of Theoretical Biology*, 89, 611-633.

- [33] Hausch, D., V. Lo, and W. Ziemba (1994): *Efficiency of Racetrack Betting Markets*. San Diego: Academic Press.
- [34] 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.
- [35] Ho, Teck-Hua, Colin Camerer, and Juin-Kuan Chong (2002): "Functional EWA: A One-parameter Theory of Learning in Games," University of Pennsylvania Working Paper.
- [36] Holt, Charles (1995): "Industrial Organization: A Survey of Laboratory Research," in *Handbook of Experimental Economics* ed. by J. Kagel and A. Roth. Princeton: Princeton University Press, 349-443.
- [37] Hopkins, E. (2002): "Two Competing Models of How People Learn in Games," *Econometrica*, in press.
- [38] Johnson, E., C. Camerer, S. Sen and T. Rymon (2002): "Detecting Backward Induction in Sequential Bargaining Experiments," *Journal of Economic Theory*.
- [39] Johnson, Johnnie E. V. and Alistair C. Bruce (2001): "Calibration of Subjective Probability Judgments in Naturalistic Settings," *Organizational Behavior and Human Decision Processes*, 85, 265-290.
- [40] Jung, Yun Joo, John H. Kagel, and Dan Levin (1994): "On the Existence of Predatory Pricing: An Experimental Study of Reputation and Entry Deterrence in the Chain-store Game," *RAND Journal of Economics*, 25, 72-93.
- [41] Kahneman, D. and S. Frederick (2002): "Representativeness Revisited: Attribute Substitution in Intuitive Judgment," in *Heuristics and Biases: The Psychology of Intuitive Judgment* ed. by T. Gilovich, D. Griffin and D. Kahneman. New York: Cambridge University Press.
- [42] Keren, Gideon (1987): "Facing Uncertainty in the Game of Bridge: A Calibration Study," *Organizational Behavior and Human Decision Processes*, 39, 98-114.

- [43] Knez, Marc J. and Colin F. Camerer (2000): "Increasing Cooperation in Prisoner's Dilemmas by Establishing a Precedent of Cooperation in Coordination Games," *Organizational Behavior and Human Decision Processes*, 82, 194-216.
- [44] Kreps (1990): "Corporate Culture and Economic Theory," in *Perspectives on Positive Political Economy* ed. by J. Alt and K. Shepsle. Cambridge: Cambridge University Press.
- [45] Kreps, D. and R. Wilson (1982): "Reputation and Imperfect Information," *Journal of Economic Theory*, 27, 253-279.
- [46] Ledyard, J. (1986): "The Scope of the Hypothesis of Bayesian Equilibrium," *Journal of Economic Theory*, 39, 59-82.
- [47] 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.
- [48] Lichtenstein, Sarah, Baruch Fischhoff and Lawrence D. Phillips (1982): "Calibration of the Probabilities: The State of the Art to 1980," in *Judgment under Uncertainty: Heuristics and Biases* ed. by Kahneman, D., P. Slovic and A. Tversky. Cambridge: Cambridge University Press, 306-334.
- [49] Manski, C. (2002): "Identification of Decision Rules in Experiments on Simple Games of Proposal and Response," *European Economic Review*, 46, 880-891.
- [50] McAllister, Patrick H. (1991): "Adaptive approaches to stochastic programming," *Annals of Operations Research*, 30, 45-62.
- [51] McKelvey, Richard D. and Talbot Page (1990): "Public and Private Information: An Experimental Study of Information Pooling," *Econometrica*, 58, 1321-1339.
- [52] McKelvey, Richard D. and Thomas R. Palfrey (1992): "An Experimental Study of the Centipede Game," *Econometrica*, 60, 803-836.
- [53] McKelvey, Richard D. and Thomas R. Palfrey (1998): "Quantal Response Equilibria for Extensive Form Games," *Experimental Economics*, 1, 9-41.

- [54] Milgrom, P. and J. Roberts (1991): "Adaptive and Sophisticated Learning in Repeated Normal Form Games," *Games and Economic Behavior*, , 82-100.
- [55] Mookerjee, Dilip and Barry Sopher (1994): "Learning Behavior in an Experimental Matching Pennies Game," *Games and Economic Behavior*, 7, 62-91.
- [56] Murphy, Alan H. and Robert L. Winkler (1977): "Can Weather Forecasters Formulate Reliable Forecasts of Precipitation and Temperature?" *National Weather Digest*, 2, 2-9.
- [57] Neral, John and Jack Ochs (1992): "The Sequential Equilibrium Theory of Reputation Building: A Further Test," *Econometrica*, 60, 1151-1169.
- [58] Nyarko, Yaw and Andrew Schotter (2002) "An experimental study of belief learning using elicited beliefs," *Econometrica*, 70, 971-1005.
- [59] Palfrey, Thomas R. and H. Rosenthal (1988): "Private Incentives in Social Dilemmas: The Effects of Incomplete Information and Altruism," *Journal of Public Economics*, 35, 309-332.
- [60] Partow, J. and A. Schotter (1993): "Does Game Theory Predict Well for the Wrong Reasons? An Experimental Investigation," C.V. Starr Center for Applied Economics Working Paper 93-46, New York University.
- [61] Roll, R. (1984): "Orange Juice and Weather," *American Economic Review*, 74, 861-880.
- [62] Roth, Alvin, Greg Barron, Ido Erev, and Robert Slonim (2001): "Equilibrium and Learning in Economic Environments: the Predictive Value of Approximations," Harvard University Working Paper.
- [63] Roth, Alvin and Ido Erev (1995): "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," *Games and Economic Behavior*, 8, 164-212.
- [64] Ruffle, Bradley J. (2001): "Some Factors Affecting Demand Withholding in Posted-offer Markets," in *Advances in Experimental Markets* ed. by T. Cason and C. Noursair. Berlin: Springer-Verlag, 33-48.

- [65] Sargent, Thomas (1999): *The Conquest of American Inflation*, Princeton: Princeton University Press.
- [66] Selten, R. (1991): "Anticipatory learning in 2-person Games," in *Game Equilibrium Models I*, ed. by R. Selten. Berlin: Springer. 98-154.
- [67] Smith, V. L., G. Suchanek and A. Williams (1988) "Bubbles, Crashes and Endogenous Expectations in Experimental Spot Asset Markets," *Econometrica*, 56, 1119-1151.
- [68] Soll, Jack (1996): "Determinants of Overconfidence and Miscalibration: The Roles of Random Error and Ecological Structure," *Organizational Behavior and Human Decision Processes*, 65, 117-137.
- [69] Stahl, Dale O. (1999): "Sophisticated Learning and Learning Sophistication," University of Texas at Austin Working Paper.
- [70] Sutton, Richard and Andrew Barto (1998): *Reinforcement Learning: An Introduction*, Boston: MIT Press.
- [71] Van Huyck, John, Raymond Battalio, and Richard Beil (1990): "Tacit Cooperation Games, Strategic Uncertainty, and Coordination Failure," *The American Economic Review*, 80, 234-248.
- [72] Watson, Joel (1993): "A 'Reputation' Refinement without Equilibrium," *Econometrica*, 61, 199-205.
- [73] Watson, Joel and Pierpaolo Battigali (1997): "On 'Reputation' Refinements with Heterogeneous Beliefs," *Econometrica*, 65, 363-374.
- [74] Yates, J. (1990): *Judgment and Decision Making*, Englewood Cliffs, N.J.: Prentice Hall.

Table 1: Payoffs for the Borrower-lender Trust Games and the Entry-deterrence Games

Payoffs in the borrower-lender trust game, Camerer & Weigelt (1988)


lender strategy	borrower strategy	payoffs to lender	payoffs to borrower normal (X)	payoffs to borrower honest (Y)
loan	default	-100*	150	0
	repay	40	60	60
no loan	no choice	10	10	10

Payoffs in the entry-deterrence game, Jung, Kagel and Levin (1994)


entrant strategy	incumbent strategy	payoffs to entrant	payoffs to incumbent normal (X)	payoffs to incumbent fighter (Y)
in	fight	80	70	160
	share	150	160	70
out	no choice	95	300	300

Note: * Loan-default lender payoffs were -50 in sessions 6-8 and -75 in sessions 9-10.


Figures 1a-f: Frequency Plots for the Trust Data from Camerer & Weigelt (1988a)




Figures 2a-f: Frequency Plots for the Unpublished Trust Data from Camerer & Weigelt (1988b)



Figures 3a-f: Frequency Plots on Inexperienced Subjects from the Entry Data from Jung, Kagel & Levin (1994)



Figures 4a-f: Frequency Plots on Experienced Subjects from the Entry Data from Jung, Kagel & Levin (1994)



Figures 5a-d: Frequency Plots on High and Low α Sessions
for the Unpublished Trust Data from Camerer & Weigelt (1988b)

Figure 5a: Empirical Frequency for No Loan for Low α Sessions




Figure 5b: Empirical Frequency for Default conditional on Loan for Low α Sessions (Dishonest Borrower)




Figure 5c: Empirical Frequency for No Loan for High α Sessions





Figure 5d: Empirical Frequency for Default conditional on Loan for High α Sessions (Dishonest Borrower)



Figures 6a-b: Empirical Frequencies for Lender (New Trust Data)

Figure 6a: Empirical Frequency of No Loan in t Given No Loan in $t-1$





Figure 6b: Empirical Frequency of Loan in t Given Default in $t-1$



Figures 7a-c: Frequency of Repay in t Given Default in $t-1$ for Dishonest Borrower (New Trust Data)






Figure 8a: Correlation of Teaching implied and reported beliefs





Figure 8b: Correlation of AQRE implied and reported beliefs








Table 2: In-sample and Out-of-sample Performance of the fEWA Teaching and AQRE Models

<u>Datasets</u>	Camerer and Weigelt (1988a)			Camerer and Weigelt (1988b)			Jung, Kagel and Levin (1994) (Inexperienced Subjects)			Jung, Kagel and Levin (1994) (Experienced Subjects)		
<u>Sample Size</u> ¹												
In-sample Calibration		5757			3820			5847			2232	
Out-of-sample Validation		2894			1882			2866			1072	
<u>The Teaching Models</u>												
Model Specification		<u>fEWA</u>	<u>pEWA</u>		<u>fEWA</u>	<u>pEWA</u>		<u>fEWA</u>	<u>pEWA</u>		<u>fEWA</u>	<u>pEWA</u>
Data Sessions	Individual	Pooled	Pooled	Individual	Pooled	Pooled	Individual	Pooled	Pooled	Individual	Pooled	Pooled
In-sample Calibration												
Hit Rate ²	77%	77%	80%	72%	74%	75%	76%	76%	79%	65%	61%	70%
Log-likelihood	-2830	-2930	-2515	-2072	-2128	-1891	-2259	-2329	-2178	-1368	-1414	-1286
Out-of-sample Validation												
Hit Rate	76%	78%	81%	72%	78%	79%	74%	73%	73%	73%	59%	79%
Log-likelihood	-1395	-1419	-1208	-998	-1028	-885	-1380	-1411	-1394	-578	-643	-546
<u>Agent-based Quantal Response Equilibrium (AQRE) Models</u>												
Model Specification		<u>AQRE</u>	<u>AQRE(θ=0)</u>		<u>AQRE</u>	<u>AQRE(θ=0)</u>		<u>AQRE</u>	<u>AQRE(θ=0)</u>		<u>AQRE</u>	<u>AQRE(θ=0)</u>
Data Sessions	Individual	Pooled	Pooled	Individual	Pooled	Pooled	Individual	Pooled	Pooled	Individual	Pooled	Pooled
In-sample Calibration												
Hit Rate	71%	69%	43%	76%	56%	46%	78%	77%	79%	53%	50%	47%
Log-likelihood	-3343	-3466	-3816	-2044	-2532	-2518	-2420	-2534	-2647	-1449	-1509	-1530
Out-of-sample Validation												
Hit Rate	70%	67%	35%	79%	56%	43%	73%	70%	75%	47%	44%	49%
Log-likelihood	-1630	-1805	-1923	-942	-1264	-1279	-1456	-1540	-1637	-696	-738	-739

Note 1: Calibrated on all observations for 70% of the subjects instead of 70% observations of all subjects.

Note 2: Number of hits counts the occasions when $\text{prob}(\text{chosen strategy}) = \text{maximum}(\text{predicted probabilities})$. Each count is adjusted by number of strategies sharing the maximum. Number of hits for the incumbent is computed using a weighted predicted probability which is the weighted average of the myopic sophisticates and the teacher where the weights are $(1-\alpha)$ and α respectively.

Note 3: All sessions are calibrated simultaneously with common parameter estimates except for scale sensitivity λ .

Table 3: Parameter Estimates for the Teaching and AQRE Models

<u>Datasets</u>	Camerer and Weigelt (1988a)				Camerer and Weigelt (1988b)				Jung, Kagel and Levin (1994) (Inexperienced Subjects)				Jung, Kagel and Levin (1994) (Experienced Subjects)			
<u>The Teaching Model</u>																
Model Specification	<u>fEWA</u>		<u>pEWA</u>		<u>fEWA</u>		<u>pEWA</u>		<u>fEWA</u>		<u>pEWA</u>		<u>fEWA</u>		<u>pEWA</u>	
Data Sessions	Individual (Min,Max)		Pooled	Pooled	Individual (Min,Max)		Pooled	Pooled	Individual (Min,Max)		Pooled	Pooled	Individual (Min,Max)		Pooled	Pooled
<u>Parameters for Adaptive Lender ¹</u>																
ϕ	0.74	0.78	0.76 (0.02)	0.41 (0.01)	0.75	0.78	0.77 (0.02)	0.24 (0.03)	0.76	0.82	0.78 (0.02)	0.57 (0.01)	0.76	0.77	0.76 (0.02)	0.53 (0.00)
δ	0.50	0.55	0.53 (0.02)	0.19 (0.01)	0.51	0.53	0.52 (0.01)	0.32 (0.02)	0.49	0.55	0.53 (0.01)	0.92 (0.00)	0.48	0.56	0.52 (0.02)	0.63 (0.00)
κ	0.58	0.75	0.65 (0.06)	1.00 (0.00)	0.51	0.75	0.62 (0.06)	1.00 (0.00)	0.56	0.76	0.65 (0.06)	0.41 (0.00)	0.53	0.67	0.60 (0.06)	1.00 (0.01)
τ	0.64	1.00	0.98 (0.00)	0.22 (0.01)	0.00	1.00	0.98 (0.00)	0.17 (0.03)	0.01	1.00	0.97 (0.00)	0.08 (0.01)	0.00	1.00	0.72 (0.00)	0.25 (0.00)
<u>Parameters for Teacher Borrower (Dishonest)</u>																
α	0.60	1.00	1.00 (0.00)	1.00 (0.00)	0.05	1.00	0.73 (0.00)	0.85 (0.03)	0.02	1.00	0.42 (0.00)	0.34 (0.00)	0.87	1.00	0.88 (0.00)	0.86 (0.01)
θ^2	0.19	0.28	0.19 (0.00)	0.24 (0.00)	0.19	0.29	0.28 (0.00)	0.19 (0.00)	0.11	0.19	0.12 (0.00)	0.11 (0.00)	0.11	0.19	0.17 (0.00)	0.12 (0.00)
<u>Agent-based Quantal Response Equilibrium (AQRE)</u>																
Model Specification	<u>AQRE</u>		<u>AQRE($\theta=0$)</u>		<u>AQRE</u>		<u>AQRE($\theta=0$)</u>		<u>AQRE</u>		<u>AQRE($\theta=0$)</u>		<u>AQRE</u>		<u>AQRE($\theta=0$)</u>	
Data Sessions	Individual (Min,Max)		Pooled	Pooled	Individual (Min,Max)		Pooled	Pooled	Individual (Min,Max)		Pooled	Pooled	Individual (Min,Max)		Pooled	Pooled
<u>Parameters for Lender</u>																
θ^3	0.24	0.29	0.29 (0.00)	<u>0.00</u> -	0.23	0.29	0.28 (0.00)	<u>0.00</u> -	0.11	0.19	0.16 (0.00)	<u>0.00</u> -	0.12	0.19	0.19 (0.00)	<u>0.00</u> -

Note 1: Each parameter is presented with its standard error (in parenthesis) directly below. For functional parameters like ϕ , δ and κ , standard deviation across subjects is reported.

Note 2: θ is constrained between 0.185 and 0.286 which is the 95% confidence interval for home-made prior derived from experiments we ran.

Note 3: The proportion of strong/honest type, p , is incorporated into the home-made prior where $q = p + \theta^*(1-p)$ where θ is estimated.

Table 4: Calibration of reported and implied beliefs in trust games

source of beliefs	<u>weighted regression</u>			<u>accuracy measures</u>		
	adjusted R ²	intercept (std. Error)	slope (std. Error)	calibration	discrimination	probability score
lender (reported)	0.74	0.664 (0.290)	0.614 (0.083)	0.011	0.030	0.152
borrower (reported)	0.72	0.778 (0.318)	0.708 (0.102)	0.006	0.040	0.141
averaged (reported)	0.86	0.115 (0.268)	0.902 (0.084)	0.004	0.042	0.135
AQRE (implied)	0.20	2.003 (0.783)	0.358 (0.197)	0.029	0.046	0.148
Teaching (implied)	0.32	1.202 (1.267)	0.545 (0.280)	0.014	0.025	0.156
overall P(default)	0.23	0.000	0.000	0.000	0.000	0.174

Appendix Table 1: In-sample and Out-of-sample Performance of the fEWA Teaching and AQRE Models in Individual Sessions for the Trust Data in Camerer & Weigelt (1988)

Session No.	1	2	3	4	5	6	7	8
No of Sequences	90	90	81	70	77	69	90	101
P(Honest)	0.33	0.33	0.33	0.10	0.10	0.10	0.00	0.00

Sample Size ¹

In-sample Calibration	744	784	742	661	673	626	703	824
Out-of-sample Validation	401	386	368	356	315	288	419	361

The Teaching Model

In-sample Calibration

Hit Rate ²	79%	81%	81%	80%	79%	81%	65%	73%
Log-likelihood	-347	-361	-365	-287	-336	-273	-427	-434

Out-of-sample Validation

Hit Rate	82%	79%	83%	79%	75%	77%	68%	69%
Log-likelihood	-181	-172	-171	-145	-175	-135	-227	-187

Agent-based Quantal Response Equilibrium (AQRE)

In-sample Calibration

Hit Rate	69%	73%	68%	79%	82%	86%	56%	60%
Log-likelihood	-405	-482	-450	-354	-357	-284	-469	-540

Out-of-sample Validation

Hit Rate	65%	71%	69%	81%	81%	88%	58%	53%
Log-likelihood	-236	-217	-213	-188	-150	-99	-282	-245

Note 1: Calibrated on all observations for 70% of the subjects instead of 70% observations of all subjects.

Note 2: Number of hits counts the occasions when $\text{prob}(\text{chosen strategy}) = \text{maximum}(\text{predicted probabilities})$. Each count is adjusted by number of strategies sharing the maximum. Number of hits for the incumbent is computed using a weighted predicted probability which is the weighted average of the myopic sophisticates and the teacher where the weights are $(1-\alpha)$ and α respectively.

Appendix Table 2: Parameter Estimates for the Teaching and AQRE Models for the Trust Data in Camerer & Weigelt (1988)

Session No.:	1	2	3	4	5	6	7	8
<u>The Teaching Model</u>								
<u>Parameters for Adaptive Lender ¹</u>								
ϕ	0.75 (0.02)	0.75 (0.02)	0.76 (0.02)	0.78 (0.02)	0.75 (0.02)	0.75 (0.02)	0.77 (0.02)	0.74 (0.02)
δ	0.55 (0.03)	0.55 (0.03)	0.55 (0.03)	0.53 (0.02)	0.51 (0.02)	0.51 (0.02)	0.52 (0.01)	0.50 (0.02)
κ	0.68 (0.05)	0.65 (0.06)	0.68 (0.05)	0.75 (0.03)	0.60 (0.06)	0.67 (0.06)	0.58 (0.07)	0.59 (0.07)
τ	0.64 (0.06)	0.96 (0.04)	0.99 (0.00)	1.00 (0.00)	0.99 (0.25)	1.00 (0.01)	0.72 (0.25)	1.00 (0.02)
<u>Parameters for Teacher Borrower (Dishonest)</u>								
α	0.60 (0.03)	1.00 (0.00)	0.99 (0.00)	1.00 (0.04)	0.99 (0.21)	0.87 (0.07)	0.82 (0.05)	1.00 (0.00)
θ^2	0.20 (0.00)	0.19 (0.00)	0.20 (0.00)	0.28 (0.02)	0.20 (0.01)	0.28 (0.00)	0.25 (0.00)	0.28 (0.00)
<u>Agent-based Quantal Response Equilibrium (AQRE)</u>								
<u>Parameters for Lender</u>								
θ^3	0.27 (0.03)	0.29 (0.03)	0.29 (0.03)	0.29 (0.00)	0.24 (0.00)	0.24 (0.00)	0.24 (0.02)	0.29 (0.03)

Note 1: Each parameter is presented with its standard error (in parenthesis) directly below. For functional parameters like ϕ , δ and κ , standard deviation across subjects is reported.

Note 2: θ is constrained between 0.185 and 0.286 which is the 95% confidence interval for home-made prior derived from experiments we ran.

Note 3: The proportion of strong/honest type, p , is incorporated into the home-made prior where $q = p + \theta^*(1-p)$ where θ is estimated.

Appendix Table 3: In-sample and Out-of-sample Performance of the Teaching and AQRE Models in Individual Sessions for the Unpublished Trust Data in Camerer & Weigelt (1988)

Session No.:	1	2	3	4	5	6	7	8
No of Sequences	40	45	61	48	54	81	75	50
P(Honest):	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10

Sample Size¹

In-sample Calibration	332	416	529	467	479	606	617	374
Out-of-sample Validation	159	203	248	230	265	315	291	171

The Teaching Model

In-sample Calibration

Hit Rate ²	56%	78%	67%	88%	80%	69%	72%	63%
Log-likelihood	-212	-197	-335	-205	-257	-298	-324	-243

Out-of-sample Validation

Hit Rate	59%	77%	61%	88%	88%	58%	69%	79%
Log-likelihood	-99	-100	-160	-111	-122	-160	-153	-92

Agent-based Quantal Response Equilibrium (AQRE)

In-sample Calibration

Hit Rate	65%	87%	61%	90%	82%	77%	82%	59%
Log-likelihood	-221	-188	-349	-183	-240	-328	-287	-248

Out-of-sample Validation

Hit Rate	77%	87%	57%	92%	90%	78%	78%	75%
Log-likelihood	-91	-89	-167	-75	-119	-158	-145	-98

Note 1: Calibrated on all observations for 70% of the subjects instead of 70% observations of all subjects.

Note 2: Number of hits counts the occasions when $\text{prob}(\text{chosen strategy}) = \text{maximum}(\text{predicted probabilities})$. Each count is adjusted by number of strategies sharing the maximum. Number of hits for the incumbent is computed using a weighted predicted probability which is the weighted average of the myopic sophisticates and the teacher where the weights are $(1-\alpha)$ and α respectively.

Note 3: All sessions are calibrated simultaneously with common parameter estimates except for scale sensitivity λ .

Appendix Table 4: Parameter Estimates for the Teaching and AQRE Models for the Unpublished Trust Data in Camerer & Weigelt (1988)

Session No.:	1	2	3	4	5	6	7	8
<u>The Teaching Model</u>								
<u>Parameters for Adaptive Lender ¹</u>								
ϕ	0.76 (0.02)	0.77 (0.02)	0.76 (0.02)	0.78 (0.02)	0.77 (0.02)	0.77 (0.02)	0.75 (0.02)	0.77 (0.02)
δ	0.51 (0.01)	0.52 (0.02)	0.51 (0.01)	0.53 (0.02)	0.52 (0.02)	0.52 (0.01)	0.51 (0.02)	0.51 (0.01)
κ	0.58 (0.07)	0.68 (0.05)	0.53 (0.06)	0.75 (0.04)	0.69 (0.05)	0.63 (0.06)	0.62 (0.07)	0.51 (0.07)
τ	0.00 (0.00)	0.97 (0.17)	0.00 (0.00)	1.00 (0.00)	1.00 (0.00)	0.62 (0.06)	1.00 (0.04)	0.98 (0.03)
<u>Parameters for Teacher Borrower (Dishonest)</u>								
α	0.48 (0.12)	0.77 (0.26)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	0.14 (0.02)	0.79 (0.01)	0.05 (0.02)
θ^2	0.29 (0.00)	0.29 (0.03)	0.29 (0.00)	0.29 (0.00)	0.29 (0.00)	0.27 (0.00)	0.28 (0.00)	0.19 (0.00)
<u>Agent-based Quantal Response Equilibrium (AQRE)</u>								
<u>Parameters for Lender</u>								
θ^2	0.27 (0.00)	0.27 (0.00)	0.29 (0.00)	0.27 (0.00)	0.25 (0.00)	0.24 (0.00)	0.24 (0.00)	0.23 (0.01)

Note 1: Each parameter is presented with its standard error (in parenthesis) directly below. For functional parameters like ϕ , δ and κ , standard deviation across subjects is reported.

Note 2: θ is constrained between 0.185 and 0.286 which is the 95% confidence interval for home-made prior derived from experiments we ran.

Note 3: The proportion of strong/honest type, p , is incorporated into the home-made prior where $q = p + \theta^*(1-p)$ where θ is estimated.

Appendix Table 5: In-sample and Out-of-sample Performance of the Teaching and AQRE Models in Individual Sessions for the Entry Data in Jung, Kagel & Levin (1994)

	Inexperienced Subjects								Experienced Subjects			
Session No.:	1	2	3	4	7	8	10	11	5	6	9	12
No of Sequences	44	90	90	90	81	81	81	75	72	72	72	72
P(Fighter):	0.33	0.33	0.33	0.33	0.00	0.00	0.00	0.00	0.33	0.33	0.00	0.00
<hr/>												
<u>Sample Size ¹</u>												
In-sample Calibration	392	754	771	784	825	816	715	790	509	535	622	566
Out-of-sample Validation	198	355	352	403	431	401	368	358	226	241	297	308
 <u>The Teaching Model</u>												
In-sample Calibration												
Hit Rate ²	63%	62%	60%	75%	87%	85%	73%	98%	72%	66%	64%	61%
Log-likelihood	-212	-408	-415	-362	-256	-218	-307	-81	-306	-336	-368	-357
Out-of-sample Validation												
Hit Rate	63%	57%	58%	77%	92%	78%	77%	75%	79%	71%	68%	73%
Log-likelihood	-132	-241	-242	-171	-94	-167	-163	-170	-114	-131	-177	-157
 <u>Agent-based Quantal Response Equilibrium (AQRE)</u>												
In-sample Calibration												
Hit Rate	73%	70%	71%	65%	81%	87%	75%	97%	46%	47%	69%	49%
Log-likelihood	-202	-391	-413	-499	-306	-219	-314	-76	-335	-362	-361	-390
Out-of-sample Validation												
Hit Rate	62%	66%	61%	71%	84%	81%	74%	82%	21%	32%	63%	63%
Log-likelihood	-125	-207	-230	-242	-130	-142	-172	-208	-154	-164	-173	-205

Note 1: Calibrated on all observations for 70% of the subjects instead of 70% observations of all subjects.

Note 2: Number of hits counts the occasions when $\text{prob}(\text{chosen strategy}) = \text{maximum}$ (predicted probabilities). Each count is adjusted by number of strategies sharing the maximum. Number of hits for the incumbent is computed using a weighted predicted probability which is the weighted average of the myopic sophisticates and the teacher where the weights are $(1-\alpha)$ and α respectively.

Note 3: All sessions are calibrated simultaneously with common parameter estimates except for scale sensitivity λ . Note that sessions with experienced subjects are pooled separately from ses

Appendix Table 6: Parameter Estimates for the Teaching and AQRE Models for the Entry Data in Jung, Kagel & Levin (1994)

	Inexperienced Subjects								Experience	
Session No.:	1	2	3	4	7	8	10	11	5	6
<u>The Teaching Model</u>										
<u>Parameters for Adaptive Entrant ¹</u>										
ϕ	0.77 (0.02)	0.76 (0.02)	0.77 (0.02)	0.77 (0.02)	0.81 (0.02)	0.81 (0.02)	0.79 (0.02)	0.82 (0.02)	0.77 (0.02)	0.76 (0.02)
δ	0.55 (0.02)	0.55 (0.02)	0.55 (0.02)	0.55 (0.02)	0.50 (0.00)	0.50 (0.00)	0.49 (0.01)	0.51 (0.00)	0.56 (0.03)	0.55 (0.03)
κ	0.56 (0.06)	0.56 (0.07)	0.58 (0.07)	0.62 (0.07)	0.71 (0.05)	0.72 (0.05)	0.65 (0.07)	0.76 (0.04)	0.67 (0.05)	0.65 (0.06)
τ	0.90 (0.10)	0.81 (0.06)	0.38 (0.24)	1.00 (0.00)	1.00 (0.00)	0.04 (0.00)	0.98 (0.01)	0.01 (0.03)	0.00 (0.02)	0.00 (0.00)
<u>Parameters for Teacher Incumbent (Weak)</u>										
α	1.00 (0.00)	0.38 (0.10)	0.69 (0.02)	0.65 (0.05)	0.45 (0.02)	0.92 (0.03)	0.52 (0.06)	0.02 (0.02)	1.00 (0.00)	0.95 (0.02)
θ^2	0.11 (0.00)	0.18 (0.01)	0.11 (0.00)	0.11 (0.00)	0.19 (0.00)	0.11 (0.00)	0.11 (0.00)	0.19 (0.00)	0.18 (0.00)	0.19 (0.00)
<u>Agent-based Quantal Response Equilibrium (AQRE)</u>										
<u>Parameters for Entrant</u>										
θ^3	0.15 (0.00)	0.17 (0.00)	0.15 (0.00)	0.11 (0.00)	0.19 (0.00)	0.19 (0.00)	0.19 (0.00)	0.15 (0.00)	0.15 (0.00)	0.15 (0.00)

Note 1: Each parameter is presented with its standard error (in parenthesis) directly below. For functional parameters like ϕ , δ and κ , standard deviation across subjects is reported.

Note 2: θ is constrained between 0.111 and 0.195 which is the 95% confidence interval for home-made prior derived from experiments we ran.

Note 3: The proportion of strong/honest type, p , is incorporated into the home-made prior where $q = p + \theta^*(1-p)$ where θ is estimated.