

Sophisticated EWA Learning and Strategic Teaching in Repeated Games

Colin F. Camerer*

*Division of Humanities and Social Sciences, California Institute of Technology,
Pasadena, CA 91125*

E-mail: camerer@hss.caltech.edu

Teck-Hua Ho

The Wharton School, University of Pennsylvania, Philadelphia, PA 19104-6366

E-mail: hoteck@wharton.upenn.edu

and

Juin-Kuan Chong

National University of Singapore, 10, Kent Ridge Crescent, Singapore 119260

E-mail: fbacjk@nus.edu.sg

Most learning models assume players are adaptive (i.e., they respond only to their own previous experience and ignore others' payoff information) and behavior is not sensitive to the way in which players are matched. Empirical evidence suggests otherwise. In this paper, we extend our adaptive experience-weighted attraction (EWA) learning model to capture sophisticated learning and strategic teaching in repeated games.

The generalized model assumes there is a mixture of adaptive learners and sophisticated players. An adaptive learner adjusts his behavior the EWA way. A sophisticated player rationally best-responds to her forecasts of all other behaviors. A sophisticated player can be either myopic or farsighted. A farsighted player develops multiple-period rather than single-period forecasts of others' behaviors and chooses to 'teach' the other players by choosing a strategy scenario that gives her the highest discounted net present value.

We estimate the model using data from p-beauty contests and repeated trust games with incomplete information. The generalized model is better than the adaptive EWA model in describing and predicting behavior. Including teaching also allows an empirical learning-based approach to reputation formation which predicts better than a quantal-response extension of the standard type-based approach.

JEL classification: C72, C91

1. INTRODUCTION

The process by which an equilibrium arises in a market or game has been a substantial mystery until recent years. Several mechanisms are possible. Models of general equilibrium assume that equilibration comes from price-change rules implemented by a fictional Walrasian auctioneer (who is presumably a stand-in for some dynamic process which is typically unspecified). An implicit model of equilibration in game theory is that players figure out an equilibrium in a game, or adhere to a recommendation by an outside arbiter (perhaps a consortium of advisors or a government agency) if it is self-enforcing (e.g., Kohlberg and Mertens [43]). Biological models ascribe equilibration to genetic reproduction as well as to mutation and natural selection. Early on, Nash spoke of a “mass action” interpretation of equilibration akin to natural selection (which is similar to modern accounts of cultural evolution).

None of these perspectives is likely to completely explain the actual time scale of equilibration in complex games played by humans. Humans learn faster than biological models predict so other learning dynamics have been studied. Most studies ask about theoretical convergence properties of dynamics, primarily to see which equilibria they converge to (if any). This paper is about the empirical fit of learning models to experimental data. Our goal is to explain as accurately as possible, for every choice in an experiment, how that choice arose from a player’s previous behavior and experience. We also strive to explain these choices using a general model which can be applied to any normal-form game with minimal customization.

Our model extends the “experience-weighted attraction” (EWA) model of (Camerer and Ho [10], [11], [12]).¹ The key property of EWA is that it hybridizes features of popular learning rules, particularly reinforcement and belief learning (of the weighted fictitious play type), which have been widely studied in game theory. Hybridizing these familiar rules is useful for two purposes, one empirical and one theoretical. The empirical purpose is to fit and predict data better. Studies have found that the hybrid EWA typically improves substantially (and significantly) on reinforcement and belief models, in 31 data sets spanning a dozen different types of games (see details below). We are not aware of any learning model that has performed as well in that many statistical comparisons.

* This research was supported by NSF grants SBR 9730364 and SBR 9730187. Many thanks to Vince Crawford, Drew Fudenberg, David Hsia, John Kagel, and Xin Wang for discussions and help. Helpful comments were also received from seminar participants at Berkeley, Caltech, Harvard, Hong Kong UST, and Wharton.

¹The model has also been applied to signaling games (Anderson and Camerer [2]), extensive-form centipede games (Camerer, Ho and Wang [13]) and bilateral call markets (Camerer, Hsia and Ho [14]).

The theoretical point of EWA is that belief learning and reinforcement learning are not different species; they are actually close cousins. When beliefs are formed according to weighted fictitious play, and used to calculate expected payoffs, those expected payoffs are exactly equal to a weighted average of previous payoffs, including “foregone payoffs” of strategies which were not chosen. Reinforcement models are averages (or cumulations) of previously received payoffs, excluding foregone payoffs. The *only* important difference between belief and reinforcement models is therefore the extent to which they assume players include foregone payoffs in evaluating strategies. In the EWA model, this difference is parameterized by a weight δ .²

This paper overcomes two limits of all adaptive models. One limit is that adaptive players do not anticipate how others are learning and do not use knowledge of other players’ payoffs (if they have it) to outguess their opponents. We add “sophistication” to the EWA model using two parameters. We assume a fraction α of players are sophisticated. Sophisticated players think that a fraction $(1 - \alpha')$ of players are adaptive and the remaining fraction α' of players are sophisticated like themselves. They use the adaptive EWA model to forecast what the adaptive players will do, and choose strategies with high expected payoffs given their forecast. This ‘self-consciousness’ assumption creates a small whirlpool of recursive thinking which means standard equilibrium concepts (Nash), and sensible generalizations like quantal response equilibrium (QRE; McKelvey and Palfrey [47], [48]), are special cases of sophisticated EWA.

The idea of sophistication has been used before, in models of “level- k ” learning (Ho, Camerer, and Weigelt [36], Stahl [71]; cf. Stahl and Wilson, [74]) and anticipatory learning (Selten [68]; Tang [75]), although our parameterization is different. It shows that equilibrium concepts combine “social calibration” (accurate guesses about the fraction of players who are sophisticated, $\alpha = \alpha'$) with full sophistication ($\alpha = 1$). But these two features can be separated in principle, and it proves to be empirically useful to do so. The model is applied to data from p-beauty contest games (Ho, Camerer and Weigelt [36], Nagel [53]) and improves fit substantially over purely adaptive models.

The second limit of adaptive models is that they do not explain why the way in which players are matched matters (e.g., Andreoni and Miller [3], Clark and Sefton [19]). Sophisticated players matched with the same players repeatedly usually have an incentive to “teach” adaptive players, by choosing strategies with poor short-run payoffs which will change what adaptive players do, in a way that benefits the sophisticated player in the

²When foregone payoffs are not known for sure, then elements of a set of possible payoffs or previously observed payoffs, can be used for ‘payoff learning’ (Anderson and Camerer [2], Camerer, Ho and Wang [13]).

long-run. This “strategic teaching” gives rise to repeated-game equilibria and reputation formation behavior through the interaction between “long-run” and “short-run” or myopic players (e.g., Fudenberg and Levine [29], Watson [79]). We allow for teaching by adding a parameter ϵ to the sophisticated EWA model which represents the weight on future payoffs (like a discount factor). If $\epsilon = 0$, a player is sophisticated but does not incorporate the effects of current actions on future payoffs, i.e. she does not teach. If $\epsilon = 1$, the player fully accounts for the likely effects of current actions on future payoffs (as in standard repeated-game models).

We estimate the teaching model on data from experiments on repeated trust games. The model fits and predicts reasonably well (better than a quantal response approach) although the data are noisy, and there is noticeable cross-session variation. It also exhibits the main patterns predicted by sequential equilibrium based on updating of entrants’ beliefs about an incumbent’s “type”. Sophisticated EWA with strategic teaching therefore provides a boundedly rational model of reputation formation without the complicated apparatus of Harsanyi “types”.

The next section describes the adaptive EWA model, motivates its structure, and briefly reviews earlier evidence. Section 3 introduces sophistication and shows empirical estimates from p-beauty contest games. Section 4 develops the teaching model and shows the empirical estimates from repeated trust games. Section 5 concludes.

2. ADAPTIVE EWA LEARNING

2.1. The Model

We start with notation. In n -person normal-form games, 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 a cardinality of $m_{-i} = \prod_{j=1, j \neq i}^n m_j$. 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 $\pi_i(s_i(t), s_{-i}(t))$.

EWA assumes each strategy has a numerical attraction. A learning model specifies initial attractions, how attractions are updated by experience, and how choice probabilities depend on attractions. The core of the EWA model is two variables which are updated after each round. The first variable is $N(t)$, which we interpret as the number of ‘observation-

equivalents' of past experience relative to one period of current experience. (A player with a low $N(t)$ puts little weight on past attractions; a player with a huge $N(t)$ is barely affected by immediate experience.) The second variable is $A_i^j(a, t)$, an adaptive player i 's attraction of strategy j after period t has taken place.³

The variables $N(t)$ and $A_i^j(a, t)$ begin with prior values, $N(0)$ and $A_i^j(a, 0)$. $N(0)$ can be interpreted as the number of periods of actual experience which is equivalent in attraction impact to the amount of pregame thinking.

Attractions are updated after each period, using the payoff that a strategy either yielded, or would have yielded, in a period. The model weights hypothetical payoffs that unchosen strategies would have earned by a parameter δ , and weights payoffs that are actually received, from chosen strategy $s_i(t)$, by an additional $1 - \delta$ (so they receive a total weight of 1). Using an indicator function $I(x, y)$ which equals 1 if $x = y$ and 0 if $x \neq y$, the weighted payoff can be written as $[\delta + (1 - \delta) \cdot I(s_i^j, s_i(t))] \cdot \pi_i(s_i^j, s_{-i}(t))$.

The parameter δ measures the relative weight given to foregone payoffs, compared to actual payoffs, in updating attractions. It can be interpreted as a kind of 'imagination' of foregone payoffs, 'simulation' of outcomes under alternative competitive scenarios ("counterfactual thinking" to psychologists), or responsiveness to foregone payoffs. A higher δ means players move more strongly, in a statistical sense, toward "ex post best responses".

The rule for updating attraction sets $A_i^j(a, t)$ to be the sum of a depreciated, experience-weighted previous attraction $A_i^j(a, t - 1)$ plus the (weighted) payoff from period t , normalized by the updated experience weight:

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

Note well that while we assume for simplicity that players are reinforced by monetary payoffs, the reinforcement function could be altered to account for loss-aversion (the aversion to losses compared to equal-sized gains; cf. Tversky and Kahneman [76]), or social preferences like fairness, reciprocity, and inequality-aversion (as in Cooper and Stockman [20]). An aspiration level or reference point (which may change over time) could also be subtracted from payoffs, which is useful for keeping reinforcement models from getting "stuck" at nonequilibrium "satisficing" responses. However, EWA

³To prepare our notation for subsequent inclusion of sophistication, we use a in $A_i^j(a, t)$ to identify the attraction of an *adaptive* player; s is associated with a *sophisticated* player.

does something like dynamic aspiration-based updating automatically, with no extra parameters.⁴

The decay rate ϕ reflects a combination of forgetting and “motion detection” – the degree to which players realize other players are adapting, so that old observations are obsolete and should be ignored. When ϕ is lower, players decay old observations more quickly and are responsive to the most recent observations.

The second rule updates the amount of experience:

$$N(t) = (1 - \kappa) \cdot \phi \cdot N(t - 1) + 1, \quad t \geq 1. \quad (2)$$

The parameter κ determines the growth rate of attractions, which reflects how quickly players lock in to a strategy.⁵ When $\kappa = 0$, attractions are weighted averages of lagged attractions and past payoffs (with weights $\frac{\phi \cdot N(t-1)}{\phi \cdot N(t-1) + 1}$ and $\frac{1}{\phi \cdot N(t-1) + 1}$); so that attractions cannot grow outside the bounds of the payoffs in the game. When $\kappa = 1$ attractions cumulate, so they can be much larger than stage-game payoffs.

We have not explicitly subscripted the key parameters δ , κ , and ϕ , but they can obviously vary across players and games (see Camerer, Ho and Wang [13]).

Attractions must determine probabilities of choosing strategies in some way. That is, $P_i^j(a, t)$ should be monotonically increasing in $A_i^j(a, t)$ and decreasing in $A_i^k(a, t)$ (where $k \neq j$). Three forms have been used in previous research: Exponential (logit), power, and normal (probit). We use the logit because it has compared favorably to the others in direct tests (Camerer and Ho [10]) and gracefully accommodates negative payoffs. The logit form is:

$$P_i^j(a, t + 1) = \frac{e^{\lambda \cdot A_i^j(a, t)}}{\sum_{k=1}^{m_i} e^{\lambda \cdot A_i^k(a, t)}}. \quad (3)$$

The parameter λ measures sensitivity of players to attractions. Sensitivity could vary due to the psychophysics of perception, whether subjects are highly motivated or not, or could reflect an unobserved component of payoffs (including variety-seeking, errors in computation, and so forth).

⁴A strategy only increases in probability (holding previous attractions constant) if its payoff is above an average of the δ -weighted foregone payoffs. Thus, EWA mimics a process in which reinforcements are payoffs minus an aspiration level which adjusts endogenously (reflecting foregone payoff).

⁵In earlier papers (Camerer and Ho [10], [11], [12]), we define $\rho = (1 - \kappa) \cdot \phi$ and call it the rate of decay for experience. The κ notation makes it clearer that the key difference is the extent to which attractions either average or cumulate.

2.2. The EWA Learning Cube

Figure 1 shows a cube with axes representing the imagination parameter δ , the change parameter ϕ , and the lock-in parameter κ . Many existing theories are special kinds of EWA learning represented by corners or edges of the cube. Cumulative reinforcement, average reinforcement, weighted fictitious play are edges and Cournot and fictitious play are vertices of this cube, as shown in the figure.

When $\delta = 0, \kappa = 1$ (and $N(0) = 1$), then $N(t) = 1$ and the attraction updating equation becomes $A_i^j(a, t) = \phi \cdot A_i^j(a, t-1) + I(s_i^j, s_i(t)) \cdot \pi_i(s_i^j, s_{-i}(t))$. This is the simplest form of cumulative choice reinforcement (Roth and Erev [63] and Erev and Roth [28]). When $\delta = 0, \kappa = 0$ (and $N(0) = 1/(1-\phi)$), the attraction updating equation becomes $A_i^j(a, t) = \phi \cdot A_i^j(a, t-1) + (1-\phi) \cdot I(s_i^j, s_i(t)) \cdot \pi_i(s_i^j, s_{-i}(t))$. This is a form of averaged choice reinforcement (attractions are averages of previous attractions and incremental reinforcement) (e.g., Mookerjee and Sopher [51]; cf. Sarin and Vahid [65]). The key property of reinforcement models is that they assume people ignore foregone payoffs. This simplifying assumption is defensible in low-information environments where players know little about foregone payoffs.⁶ However, in most experimental games that have been studied empirically, players *do* know foregone payoffs and seem to respond to them. There is even evidence that pigeons are sensitive to foregone payoffs!⁷

A more surprising restricted case is weighted fictitious play (Brown [6], Fudenberg and Levine [30]).⁸ When $\delta = 1$ and $\kappa = 0$ the attractions are updated according to $A_i^j(a, t) = \frac{\phi \cdot N(t-1) \cdot A_i^j(a, t-1) + \pi_i(s_i^j, s_{-i}(t))}{\phi \cdot N(t-1) + 1}$. That is, attractions are weighted averages of lagged attractions and either realized or foregone payoffs. This sort of belief learning is a special kind of generalized reinforcement because beliefs can be written in the form of a difference equation. When beliefs are used to calculate expected payoffs for strategies,

⁶Even in those environments, however, some reinforcement rules learn too slowly (see Van Huyck, Battalio and Rankin [78]). Rapid learning probably occurs because players learn about foregone payoffs over time; (e.g., Camerer, Ho, and Wang [13], Anderson and Camerer [2]).

⁷Gallistel [32] (chapter 11) explains that the tendency of pigeons to “probability match” in binary choice experiments is affected by information about foregone payoffs. Specifically, pigeons tended to maximize, choosing one of two levers with the highest chance of delivering a reward all the time, when the pigeons knew after an unsuccessful trial that the other lever would have delivered a reward. (How did the pigeons “know”? Because a light displayed above a lever came on afterwards only if the lever had been armed for reward. If the light came on above the lever they did not choose, they ‘knew’ the foregone payoff.) When the pigeons did not know about the foregone payoffs (no light told them which lever had been armed to deliver food), they tended to “probability match” (to choose each lever about as often as that lever delivered a reward). So even pigeons notice foregone payoffs.

⁸Weighted fictitious play is a discrete dependent variable form of the adaptive expectations equation introduced by Cagan and Friedman in macroeconomics.

then the expected payoffs can also be written in the form of a difference equation: Expected payoffs are equal to previous expected payoffs and the increment in expected payoff which results from the updated belief. In the expected payoff equation, the belief disappears. The trick is that since beliefs are only used to compute possible future payoffs, and beliefs are backward-looking, possible future payoffs can be computed directly by incrementing expected payoffs to account for the “recently possible” foregone payoff. Seen this way, the difference between simple reinforcement and belief learning is a matter of degree, rather than kind (particularly the value of δ).

The relation between belief and reinforcement models is subtle and went unnoticed for decades.⁹ Why? For one thing, behaviorist psychologists liked the idea of reinforcement precisely because it *avoided* “mentalistic” constructs like beliefs; so the last thing they were interested in was linking reinforcement and belief formation. And when weighted fictitious play was introduced in game theory, it was thought of as a heuristic way for players to reason their way to an equilibrium, not as a literal theory of how players learn from observation. It therefore emerged from a way of thinking that was (apparently) quite different from reinforcement learning.

Indeed, there is no compelling *empirical* reason to think parameter configurations which characterize human behavior will necessarily lie on the cube edges corresponding to belief and reinforcement learning, rather than on other edges or some interior regions. The kind of ‘empirical privilege’ that would justify focusing attention in those regions *might* have come from a variety of studies which continually show that measured parameters cluster in one portion of the cube. But that never happened. Instead, most studies compare models from one corner or vertex with a static benchmark (usually Nash equilibrium). These studies provide little information about which rules—i.e., points in the cube—best characterize how people learn.

2.3. Empirical Evidence

In previous empirical research, EWA has been used to fit and predict data from order-statistic coordination games (Camerer and Ho [11], [12]), p -beauty contests (Camerer and Ho [12]), mixed strategy games (Camerer and Ho [12]), n -person (Hsia [40]) and bilateral (Camerer, Hsia and Ho [14]) call-markets, cost allocation processes (Chen and Khoroshilov [18]), extensive-form centipede games (Camerer, Ho and Wang [13]), “unprofitable” games (Morgan and Sefton [52]), signaling games (Anderson and Camerer

⁹For example, Selten [69] wrote “. . . in rote (reinforcement) learning success and failure directly influence the choice probabilities. . . . Belief Learning is very different. Here experiences strengthen or weaken beliefs. Belief learning has only an indirect influence on behavior.” EWA makes clear, however, that the indirect influence of learning of beliefs (for weighted fictitious play) can be exactly mimicked by direct influence.

[2]), patent race games with iteratively dominated strategies (Rapoport and Amaldoss [60]), patent race games (Amaldoss [1]), and 5x5 matrix games (Stahl [73]).¹⁰

Table 1a-b summarize EWA parameter estimates and goodness-of-fit statistics from these 31 data sets. The goodness-of-fit statistic is -1 times log likelihood except in Chen and Khoroshilov [18]. The column “EWA” reports the $-LL$ of the EWA model. The reinforcement and belief models report the *difference* between the $-LL$'s of those models and the EWA statistic. (A positive difference means EWA fits better.)

Values of δ tend to be between .5 and 1 in most studies except those in which games have only mixed-strategy equilibria, where δ is close to zero. The value of ϕ is reliably around .9 or so, with some exceptions.

What about model comparisons? The fairest comparisons estimate parameters on part of a sample of data and forecast choices out-of-sample, so that models with more parameters will not necessarily fit better. (Indeed, if overly complex models succeed in-sample only by overfitting, they will predict *worse* out-of-sample.) In the 11 out-of-sample comparisons (denoted “OUT” in the third column from the right), EWA always outperforms reinforcement, although usually modestly. EWA outperforms belief learning in 9 of 11 cases, quite dramatically in some data sets.

Of course, EWA necessarily fits better in the other 20 in-sample comparisons than reinforcement and belief models which are special cases. One can use standard statistical techniques for penalizing more complex models – the χ^2 test, and Akaike and Bayesian criteria. These techniques are created so that if the special case restriction is true, the penalized fit of the more complex model will be worse than the fit of the restricted model. EWA generally does better even after penalizing it. For example, if the difference in LL is 4 points or more then the special-case restriction will be rejected by the χ^2 test. By this criterion, EWA is more accurate than belief learning in all in-sample comparisons, and more accurate than reinforcement in 16 out of 20 comparisons.

Figure 1 shows the locations of estimated parameter combinations from 20 games in Table 1¹¹ in the EWA cube. Each point represents a triple of estimates of ϕ, δ and κ in a single game. The first observation is that points do not particularly cluster on the edges or corners corresponding to extreme-case theories, except for a few points in the lower corner corresponding to averaged reinforcement ($\delta = \kappa = 0$, and ϕ close to one).

¹⁰Ho and Chong [37] applied a variant of the EWA model to fit and predict 130,000 consumer product purchases in supermarkets and found that the EWA model performed substantially better, in fit and prediction, than existing reinforcement models.

¹¹In some studies, the same game was played at different stakes levels. In these cases, estimates were averaged across stakes levels, which explains why the 31 estimates in Table 1 shrink to 20 points in Figure 1.

The second observation is that points are dispersed throughout the cube. Either learning rules are fundamentally different in different games— which creates the need for some theory of which parameter combinations are used in which games (see Ho, Camerer and Chong [35])— or there may be some way to add something to the model to create more stability in parameter estimates.

Interestingly, three of four vertices on the $\kappa = 0$ and $\kappa = 1$ faces of the cube have been studied previously, but one has not. The fourth vertex, in which players are fully responsive to foregone payoffs ($\delta = 1$) but attractions cumulate rather than average past payoffs ($\kappa = 1$), does not correspond to any familiar learning theory. However, the estimates from the three order-statistic coordination games are close to this segment. This vertex is also prominent in our estimates of the p-beauty contest game reported below. These results show an advantage of thinking about points in the learning cube: Parameter configurations never imagined before fit the learning path in some data sets better than models like simple belief and reinforcement learning which have been studied for fifty years.

We end this introductory section with four comments. First, others have explored the econometric properties of reinforcement, belief and EWA learning, and the news is not all good. For example, Salmon [67] finds with simulations that in 2x2 games, reinforcement, belief, and EWA models are often poorly recoverable in the sense that rule Y cannot be rejected as a good fit of data actually generated by a different learning rule X. EWA does least poorly in this sense because it does properly identify the value of δ . That is, when the data are generated by reinforcement (belief) models with $\delta=0$ ($=1$), EWA model estimates are close to the correct value of δ . Blume et al. [4] find fairly good econometric performance of EWA and some other rules when there are repeated samples and a substantial span of data, and poor performance in small samples. Cabrales and Garcia-Fontes [7] find excellent performance. These studies are harsh reminders that we should be more careful about investigating econometric properties of estimators before rushing to apply them. Econometric pretests can also guide the choice of design features that are most likely to produce good econometric recovery.

Second, it is often useful to economize on EWA parameters that must be estimated. One approach is to make free parameters function of observed data. For example, since ϕ captures something about subjects' awareness of how rapidly their opponents' moves are changing, its parameter value can be tied to a function that detects rate of change. Indeed, we have recently developed a function-based "EWA Lite" and found that it performs almost as well as adaptive EWA with free parameters (Ho, Camerer, and Chong [35]). Excluding the initial conditions, the EWA Lite theory has

only one parameter (λ) to be estimated. This one-parameter EWA Lite model should appeal to modelers who want a highly parsimonious model.

Third, it would be useful to prove something about the long-run behavior of EWA players (cf. Hopkins [39]). Heller and Sarin [34] make a much-needed start in this direction. We conjecture that if $\kappa = 0$ (so that attractions are weighted averages of previous attractions and payoffs), then EWA players will converge to something akin to ϵ -equilibrium (at least in those classes of games where fictitious play converges) and ϵ will depend on δ and the payoffs in the game. The idea is that players could converge to a non-best response, but only if their stable payoff π_{stable} is greater than δ times the highest (best response) foregone payoff π_{br} . The gap $\pi_{br} - \pi_{stable}$ is a measure of ϵ . Since mixed-strategy equilibria are knife-edge equilibria (the equilibrium mixtures are only weak best responses, by definition), the set of ϵ -equilibrium will often be much larger than the mixed equilibrium. This helps explain why convergence in games with mixed-strategy equilibria is often so noisy, and such games are often a very poor way to distinguish models (e.g. Salmon [67]). Perhaps what we see in these games is players wandering among a large set of ϵ -equilibria which are produced by EWA equilibration.

Finally, while the EWA cube spans a large family of plausible learning rules, other theories are not special cases of EWA and also deserve further exploration. Crawford [22] created a very general form of belief learning in which decay parameters ϕ can vary over time and people, and beliefs may be subject to period-specific shocks. Broseta [5] economized on some parameters in the Crawford model by specifying ARCH (autoregressive conditionally heteroskedastic) errors. They both applied their models to experimental data on order-statistic coordination games (and see Crawford and Broseta [23]). Studying belief learning models which are more general than weighted fictitious play (like Crawford's and Broseta's) is important because one study of direct measurement of beliefs found that beliefs did not correspond closely to fictitious play calculations (Nyarko and Schotter [56]). Another important approach is "rule learning", in which players shift weight among various rules (e.g., Salmon [66], Stahl [72], [73]).

3. EWA LEARNING WITH SOPHISTICATION

For game theorists steeped in a tradition of assuming players reason thoughtfully about the behavior of others, introducing sophistication into learning is a natural step; indeed, *not* assuming sophistication might seem strange. However, our standards are entirely empirical. We would like to know whether adding sophistication to an adaptive model parsimoniously (and the reverse, "dumbing down" sophisticated models by adding *unsophistication*) helps explain how people behave.

There are several empirical reasons to allow sophistication:

1. Players *do* use information about others' payoffs. Experiments that compare behavior with and without other-payoff information found a significant difference (Partow and Schotter [59], Mookerjee and Sopher [50], Cachon and Camerer [8]). The use of other-payoff information can also be tested directly, by measuring whether players open boxes on a computer screen that contain payoffs of other players. They do (Costa-Gomes, Crawford, and Broseta [21]; cf. Camerer, Johnson, Sen and Rymon [15]).

2. If players are sophisticated, the way in which they are matched when a game is played repeatedly can affect behavior. For example, compared to the random-opponent matching protocol, the fixed-opponent matching protocol should encourage players to adopt repeated game strategies.

3. Ho et al. [36] show that experienced subjects who played a second p -beauty contest converge significantly faster to Nash equilibrium than inexperienced subjects. This can be interpreted as evidence that players learned from the first p -beauty contest about how others were learning, which means they became increasingly sophisticated.

4. In some experiments, players change strategies in ways that are inconsistent with adaptation but are consistent with sophistication. For example, Rapoport, Lo and Zwick [61] studied market entry games in which players had to enter one of three markets (see Ochs [57] for an overview). If a particular market was "under-entered", relative to the Nash equilibrium entry rate, then any sensible adaptive model (such as EWA and the restricted cases) predict more entry into that market in the next trial. In fact, players tended to enter even *less* frequently on subsequent trials, which is consistent with sophistication (i.e., expecting too much entry by other adaptive players, and hence avoiding that market).

3.1. The Model

Assume a population of both adaptive learners and sophisticated players. Denote the proportion of sophisticated players by α and the proportion of adaptive players by $(1 - \alpha)$. Adaptive learners follow the EWA learning rules and sophisticated players develop forecasts of others by assuming $(1 - \alpha')$ proportion of the players are adaptive EWA learners and the rest are like themselves and best-respond to those forecasts.

Adaptive EWA learners follow the updating and probability equations (1)-(3). The sophisticated players have attractions and choice probabilities

specified as follows¹²

$$A_i^j(s, t) = \sum_{k=1}^{m-i} [\alpha' P_{-i}^k(s, t+1) + (1 - \alpha') \cdot P_{-i}^k(a, t+1)] \cdot \pi_i(s_i^j, s_{-i}^k) \quad (4)$$

$$P_i^j(s, t+1) = \frac{e^{\lambda \cdot A_i^j(s, t)}}{\sum_{k=1}^{m_i} e^{\lambda \cdot A_i^k(s, t)}}. \quad (5)$$

For a given player i , the likelihood function of observing a choice history of $\{s_i(1), s_i(2), \dots, s_i(T-1), s_i(T)\}$ is given by:

$$\alpha \cdot [\prod_{t=1}^T P_i^{s_i(t)}(s, t)] + (1 - \alpha) \cdot [\prod_{t=1}^T P_i^{s_i(t)}(a, t)] \quad (6)$$

The proposed model passes a form of the “publishability test” articulated by McKelvey and Riddihough [49]. They argue that a good social science theory should still apply even after it is “published” or widely-understood; or if behavior changes after publication, the theory should contain an explanation for why change occurs. Our model passes this test if only sophisticated players can “read”, since sophisticated players will not change their behavior as long as adaptive learners remain unsophisticated. The only theory which passes full-readability test is when $\alpha = 1$ ($\alpha < 1$ corresponds to ‘limited circulation’ or ‘illiteracy’.)

Because the model assumes that sophisticated players think others are sophisticated (and those others think others are sophisticated...), it creates a whirlpool of recursive thinking which nests equilibrium concepts. Quantal response equilibrium (McKelvey and Palfrey [47], [48], Chen, Friedman and Thisse [17]; c.f. Rosenthal [62]) is equivalent to everyone being sophisticated ($\alpha = 1$) and the sophisticated players having rational expectations or “social calibration” about the proportion of sophisticates ($\alpha = \alpha'$). (Nash equilibrium, which we prefer to call ‘hyper-responsive QRE’, is just $\alpha = \alpha'$ along with infinite responsive sensitivity λ .) Weizsacker [81] allows players to mistakenly believe other players’ sensitivities are different that this difference explains data from some one-shot games.

Our parameterization emphasizes that QRE consists of the conjunction of two separate modeling assumptions: Players are sophisticated ($\alpha=1$) and sophisticated players are socially calibrated ($\alpha = \alpha'$). The two assumptions can be evaluated separately (and our model does). Including

¹²This specification assumes that sophisticated players’ and the modeler’s forecasts of the adaptive players are identical. A more general specification can allow them to be different.

both α and α' allows for two (opposite) kinds of judgment biases in assessing relative sophistication: Sophisticated subjects could underestimate the number of subjects who are sophisticated like themselves ($\alpha' < \alpha$, “false uniqueness” or overconfidence about relative sophistication), or could overestimate the number of sophisticates ($\alpha' > \alpha$, “false consensus” or “curse of knowledge”).

Many studies document various types of optimism, overconfidence, or “false uniqueness”. For example, most people say they are above average on good traits and below average on bad traits. Most of these studies simply use self-reports and do not pay people according to their actual ranking, but a few studies have used experimental economics methods, which include financial incentive for accuracy and a clear definition of trait and rank, and replicate the basic finding. Applied to the sophisticated EWA model, overconfidence about relative sophistication would imply that sophisticates think there are fewer people as “smart” as themselves than there actually are, so $\alpha' < \alpha$. This kind of overconfidence is built into “level- k types” models like those of Stahl and Wilson [74] (see also Costa-Gomes, Crawford, and Broseta [21], Ho et al. [36]). In those models, level 0 players choose randomly and level $k+1$ players best-respond to behavior of level k players. In a sense, this structure means players at every level think that nobody is as smart as them, and that everybody else is one level below. In our model, setting $\alpha' = 0$ corresponds to a level 1 learning type.

The opposite mistake is called “false consensus”: People overestimate how much like themselves other people are.¹³ A related effect is the inability of people who have learned new information to imagine what not knowing the information is like, the “curse of knowledge”.¹⁴ In sophisticated EWA a false consensus bias would imply that sophisticated people overestimate how many others are sophisticated, so that $\alpha' > \alpha$.

3.2. Dominance-solvable p -beauty contest games

We estimate the sophisticated EWA model using data of the p -beauty contests collected by Ho et al. [36]). In a p -beauty contest game, n players simultaneously choose numbers x_i in some interval, say $[0,100]$. The average of their numbers $\bar{x} = \frac{\sum_i^n x_i}{n}$ is computed, which establishes a target

¹³The term “false consensus” is a misnomer because a ‘rational bias’ of this sort will result if people use their own behavior or tastes as a sample of tastes and update their beliefs about population tastes. The effect is correctly defined as overweighting your own tastes relative to information about the tastes of others. Engelmann and Strobel [27] show that there is no such bias when it is defined this way and information about other people is presented, along with financial incentives for accuracy.

¹⁴An example of false consensus is due to George Wu. He asked his students whether they had cell phones, and also asked them to estimate what fraction of students in the class had cell phones. Students with cell phones thought 55% had them, and those without cell phones thought only 35% had them.

number, τ , equal to $p \cdot \bar{x}$. The player whose number is closest to the target wins a fixed prize $n \cdot \pi$ (ties are broken randomly).¹⁵

P -beauty contest games were first studied experimentally by Nagel [53] and extended by Duffy and Nagel [25] and Ho et al. [36]). These games are useful for estimating the number of steps of iterated dominance players use in reasoning through games. To illustrate, suppose $p = .7$. Since the target can never be above 70, any number choice above 70 is stochastically dominated by simply picking 70. Similarly, players who obey dominance, and believe others do too, will pick numbers below 49 so choices in the interval (49, 100] violate the conjunction of dominance and one step of iterated dominance. The unique Nash equilibrium is 0.

In experiments, initial choices are widely dispersed and centered somewhere between the interval midpoint and the equilibrium (see Nagel [54]). This basic result has been replicated with students on three continents and with several samples of sophisticated adults, including economics Ph.D.'s and corporate CEOs (see Camerer [9]). When the game is repeated, numbers gradually converge toward the equilibrium.

Explaining beauty contest convergence is surprisingly difficult for adaptive learning models. Choice reinforcement converges too slowly, because only one player wins each period and the losers get no reinforcement. Belief models with low values of ϕ (close to Cournot) fit better, but also learn too slowly (Ho et al. [36]).

The sophisticated EWA model was estimated on a subsample of data collected by Ho et al. [36]. Subjects were 196 undergraduate students in computer science and engineering in Singapore. Each seven-person group of players played 10 times together twice, with different values of p in the two 10-period sequences. (One sequence used $p > 1$ and is not included below.) The prize was .5 Singapore dollars per player each period, about \$2.33USD per group for seven-person groups. They were publicly told the target number τ and privately told their own payoff (i.e., whether they were closest or not).

We analyze a subsample of their data with $p = .7$ and $.9$, from groups of size 7. This subsample combines groups in a 'low experience' condition (the game is the first of two they play) and a 'high experience' condition (the game is the second of two, following a game with $p > 1$).

Some design choices are needed to implement the model. The subjects chose integers in the interval [0,100], a total of 101 strategies. If we allow 101 possible values of $A^j(0)$ we use too many degrees of freedom estimating initial attractions. Rather than imposing too many structural requirements

¹⁵Formally, $\pi(x_i, x_{-i}) = \frac{n \cdot \pi \cdot I(x_i, \operatorname{argmin}_{x_j} |x_j - \tau|)}{\sum_i I(x_i, \operatorname{argmin}_{x_j} |x_j - \tau|)}$.

on the distribution of $A^j(0)$ to economize on parameters, we use the first-period data to initialize attractions.

Denote the empirically observed frequency of strategy j in the first period by f^j . Initial attractions are recovered from the equations¹⁶

$$\frac{e^{\lambda \cdot A^j(0)}}{\sum_k e^{\lambda \cdot A^k(0)}} = f^j, j = 1, \dots, m. \quad (7)$$

Some algebra shows that the initial attractions can be solved for, as a function of λ , by

$$A^j(0) - \frac{1}{m} \sum_j A^j(0) = \frac{1}{\lambda} \ln(\tilde{f}^j), j = 1, \dots, m \quad (8)$$

where $\tilde{f}^j = \frac{f^j}{(\prod_k f^k)^{\frac{1}{m}}}$ is a measure of relative frequency of strategy j . We fix the strategy j with the lowest frequency to have $A^j(0) = 0$ (which is necessary for identification) and solve for the other attractions as a function of λ and the frequencies \tilde{f}^j .

Since subjects who lost did not know the winning number, they could not precisely compute foregone payoffs. Therefore, we assume they have a simple uniform belief over a range of possible winning numbers.¹⁷ We assume the losers reinforce numbers in the interval $[\tau - \frac{\delta \cdot n \cdot \pi}{d}, \tau + \frac{\delta \cdot n \cdot \pi}{d}]$. The amount of reinforcement is of a triangular form with the maximum of δ times the prize at the target number and decreases linearly at a slope of d (which is a parameter to be estimated). Denote the winning number to be w and the distance between the target and the winning number by $e = |\tau - w|$. Winners reinforce numbers in the intervals $(\tau - e, \tau + e)$ by δ times the prize. Winners reinforce the boundary number they choose, either $\tau - e$ or $\tau + e$, by the prize divided by the number of winners, and reinforce the other boundary number by δ times the prize divided by the number of winners. If there is only one winner, she also reinforces numbers in the intervals $(\tau - e, \tau - e - \frac{\delta \cdot n \cdot \pi}{d})$ and $(\tau + e, \tau + e + \frac{\delta \cdot n \cdot \pi}{d})$ by a reinforcement of a triangular form with the maximum of δ times the prize at $\tau - e$ and $\tau + e$ and decreases linearly at a slope of d with smaller and larger number respectively.

¹⁶This procedure is equivalent to choosing initial attractions to maximize the likelihood of the first-period data, separately from the rest of the data, for a value of λ derived from the overall likelihood-maximization.

¹⁷In Camerer and Ho [10], we assume that subjects know the winning number. Assuming subjects having a belief over the possible winning numbers provides a significantly better fit for the observed data.

Table 2 reports results and parameter estimates.¹⁸ For inexperienced subjects, adaptive EWA generates Cournot-like estimates ($\hat{\phi} = \hat{\rho} = 0$ and $\hat{\delta} = .90$). Adding sophistication increases $\hat{\phi}$ and improves LL by 60 and 24 points in- and out-of-sample. The estimated fraction of sophisticated players is .24 and their estimated perception $\hat{\alpha}'$ is zero. The imagination parameter δ is estimated to be .78 in the sophisticated model.

Experienced subjects show a large increase in sophistication. The estimated proportion of sophisticates, and their perceptions, rise to .77 and .41 in the experienced sample. As a result, adaptive EWA fits much worse than sophisticated EWA, differences in LL of more than 200 points both in- and out-of-sample. The increase in sophistication due to experience reflects a kind of “cross-period” learning which is similar to rule learning (Stahl [72]) or “rule switching” (Salmon [66]). The difference is that in Salmon and Stahl’s approaches, players either keep track of actual or prospective performance of different rules, and switch in the direction of better-performing rules (Stahl [72]) or switch away from poorly-performing rules (Gale, Binmore and Samuelson [31], Salmon [66]). In our current specification, this change in rules can only occur between games, but it could be easily adapted to allow within-session rule changes (see Camerer and Ho [10]).

Figures 2a and 3a show actual choice frequencies for inexperienced and experienced subjects, respectively. Experienced subjects actually start by choosing somewhat higher numbers (probably due to “negative transfer” from their earlier experience with a $p > 1$ game with an equilibrium at 200), but converge more rapidly. By round four nearly half the experienced subjects choose numbers 1-10 (the large spikes on the back and left of Figure 3a). By contrast, after nine rounds only a third of the inexperienced subjects choose 11-20. The fact that experienced subjects start *farther* from equilibrium, and end up much closer, is an indication that they are learning more rapidly or behaving more sophisticatedly (i.e., anticipating learning by others).

Figures 2b-c show the frequencies predicted by the adaptive EWA (2b) and sophisticated EWA (2c) models, for inexperienced subjects. Both models fit the general patterns in the data fairly well. Keep in mind that if

¹⁸Standard errors are derived from a bootstrap procedure. Two hundred sets of bootstrapped estimates are produced using the method of maximum likelihood. Each bootstrapped estimate is derived by maximizing a weighted log-likelihood. The weight for each observation in this log-likelihood is randomly generated such that the sum of all weights add to one. For each parameter, we sort the 200 bootstrapped estimates and note the 2.5 and 97.5 percentile estimates (i.e. the 6th and the 295th estimates). We take the difference between these two estimates to be 3.92 times the pseudo-standard errors and infer pseudo-standard errors from the difference in estimates. If the bootstrap procedure generates many outliers the standard errors will be large so we think our procedure is conservative (i.e., does not understate standard errors).

the model overfits in periods 2-7, it would predict particularly badly in the out-of-sample periods 8-10. It does not. The difference between Figures 2b and 2c shows that adding sophistication adds only a small visual improvement for inexperienced subjects (consistent with the modest increase in LL).

Figures 3b-c show the frequencies predicted by the adaptive EWA (3b) and sophisticated EWA (3c) models, for experienced subjects. Here there is a substantial improvement from including sophistication (compare 3c with 3b and 3a), which appears to fit nicely.

Restrictions of the sophisticated model generally degrade fit and predictive accuracy a lot. Imposing $\alpha' = \alpha = 1$ creates quantal-response equilibrium, which is 376 and 232 points worse in LL for inexperienced and experienced subjects. One way for QRE to capture learning in a very reduced form is to allow the response sensitivity λ to vary over time. Allowing this *does* produce increasing values of λ (reported in Table 2), but the difference in in-sample LL is still very large, 365 and 165.¹⁹ One problem with QRE in these games is that the data are multi-modal, with “spikes” reflecting discrete levels of reasoning. (For example, in very large samples from newspaper games with $p = 2/3$, there are sharp spikes at 33 and 22, reflecting one or two steps of iterated reasoning from a perceived mean of 50.) QRE does not produce spikes of this type so it captures central tendencies but not multimodality. A QRE model in which different players have different values of λ (but commonly know the distribution so the ‘E’ assumption is satisfied) will accommodate spikes and might fit better (cf. Goeree and Holt [33]).

We did not estimate simple choice reinforcement models on these data because they do an awful job. Only one of seven players receives any reinforcement each period, so learning is far too slow.²⁰ Cournot-like belief models will fit better; indeed, for both groups of subjects the EWA estimate of $\hat{\delta}$ is quite close to one and ϕ is low.

We also estimated the effects of rational expectations (RE, $\alpha = \alpha'$) and egocentric bias or level-type ($\alpha' = 0$) restrictions. For inexperienced subjects, the losses in log likelihood relative to sophisticated EWA are 3.63 and 2.85 for RE, in- and out- of-sample, and 21.24 and .16 for egocentric

¹⁹Our procedure of estimating the model in-sample and fixing parameter values to forecast out-of-sample makes life difficult for the varying- λ QRE model, since we fix λ at the last (period 7) estimate to forecast periods 8-10. A better procedure would impose some increasing functional structure on the $\lambda(t)$ function so that λ would continue to increase in the out-of-sample periods. Note, however, that the in-sample LL is substantially worse than sophisticated EWA for inexperienced subjects, and slightly worse with experienced subjects.

²⁰This is a general problem for choice reinforcement models in games where $n - 1$ players earn nothing, such as auctions, winner-take-all tournaments, market games with one seller and many buyers (or vice versa), and so forth.

bias. For experienced subjects the analogous figures are 10.78 and 62.43 for RE, and 14.29 and 13.50 for egocentric bias.²¹ While these differences in LL are small, both restrictions can be rejected (particularly RE), which shows the predictive appeal of a model that separates sophistication and perceived sophistication without imposing the strict level structure. Also, the gap between α and α' grows with experience, from 24% to 34% (and consequently, the RE restriction is rejected more strongly for experienced subjects). While players get more sophisticated between sessions, it seems that they also overestimate how many others become sophisticated.

A final observation creates a segue to the next section of this paper. In the actual frequency plots Figures 2a and 3a, the attentive eye can't help but notice the small number of very large choices (typically 100), particularly in later rounds. Ho et al. [36] called these large numbers "spoilers" and tested several explanations for why people might choose them. The most likely possibility is that subjects believe others are learning according to some adaptive rule that responds to the previous mean. By choosing a large number in round t , they throw adaptive learners off the trail, causing the adaptive learners to choose artificially high numbers in round $t + 1$, which improves the spoiler's chance of winning by choosing a low number. This kind of behavior combines two ingredients: A belief that others are learning adaptively; and a willingness to sacrifice period t profits (since picking 100 guarantees a loss) for the sake of increased future profits. This is our first glimpse of strategic teaching.

4. STRATEGIC TEACHING

For a sophisticated player who anticipates that others learn, it is natural to 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. We call the behavior which maximizes discounted expected payoffs in this way "strategic teaching".

The basic idea is described by Fudenberg and Levine [30] (pp. 261-263; cf. Ellison [26]). They write:

...imagine that one player is myopic and follows the type of learning procedure we have discussed in this book, while another player is sophisticated and has a reasonably good understanding that his opponent is using a learning procedure of this type. What happens in this case?... [much as] the results on equilibrium learning carry over to the case of nonequilibrium

²¹Imposing these restrictions does not change other parameter estimates much, except to increase δ to one for the egocentric restriction in both inexperienced and experienced subject samples.

learning, so we expect that the lessons of the literature on reputation will carry over also to the case of nonequilibrium learning.

Fudenberg and Levine [29] showed that by playing an optimal precommitment strategy forever (and waiting for the adaptive player to come around to best-responding), a sufficiently patient strategic teacher can get almost as much utility as from the Stackelberg equilibrium. In their book they add (p. 262) that “the basic argument carries over in a straightforward way to the case of nonequilibrium learning” (cf. Watson [79], Watson and Battigali [80]).

Strategic teaching extends the reach of the EWA model to incorporate two phenomena which are beyond the grasp of adaptive models: (1) The influence of fixed-matching versus re-pairing protocols, and (2) emergence of repeated-game behavior including, importantly, reputation formation without cumbersome updating of “types” (a la Harsanyi).

If some players are capable of strategic teaching, then how players are matched, and feedback that they are given, should affect learning. In fact, there is evidence that fixed-pair matching and random rematching produce different behaviors, which shows indirectly the likely influence of strategic teaching. For example, Andreoni and Miller [3] show that there is more mutual cooperation in finitely-repeated prisoners’ dilemma games when subjects play repeatedly with a fixed “partner” than when they are re-paired with “strangers” in each period. Van Huyck, Battalio and Beil [77] found a similar phenomenon in two-player “weak-link games” (which are stag hunt or assurance games with seven strategies rather than two). They compared partner pairings with stranger re-pairing. The distributions of choices in the first period of the two pairing conditions were similar, but partner pairs were able to converge to the efficient equilibrium reliably (12 of 14 did so) while the stranger re-pairing behavior did not. Clark and Sefton [19] reported a similar result. It appears that subjects who make efficient choices in the partner pairings, and see their partner choose an inefficient number in the first period, are inclined to “patiently” make an efficient choice once or twice more, as if holding their behavior steady and anticipating that the other player will learn to play efficiently.²²

²²The same difference in partner and stranger matching does *not* seem to be present in three-player groups (see Knez and Camerer [42]). We conjecture that the difference in two- and three-player dynamics can be traced to strategic teaching. In both cases, the success of a strategic teacher who makes the efficient choice repeatedly depends on behavior of the “least teachable” player. Since there are two other players being taught in the three-player game, and only one in the two-player game, strategic teachers are more likely to give up and converge toward inefficiency in three-player games. The same dynamics might help explain the fact that collusion is sustainable in repeated pricing and quantity games when the number of players is small, and difficult to sustain when the number is large (e.g., Holt [38]).

These stylized facts are consistent with strategic teaching. Strategic teachers who are matched with a different partner each time cannot use their current choices to influence what will happen in the future (to their benefit) if their future partners do not know the teachers' history of choices and anticipate similar choices in the future.

Introducing teaching allows the possibility that repeated-game behavior is different than simply repeating stage-game behavior. Of course, in theory strategies which are not equilibrium strategies in a stage game can be used in repeated-game equilibria (by the threat of reversion to a Pareto-inefficient equilibrium if a defection occurs), as the "folk theorem" of repeated games would suggest. A familiar example is the repeated-PD, in which playing tit-for-tat is a repeated-game equilibrium (if the discount factor is large enough, relative to payoffs), supported by the threat of reversion to mutual defection. This kind of dynamic is precisely what teaching can explain: A strategic teacher may play a strategy which is not myopically optimal (such as cooperating in a PD) in the hope that it induces adaptive players to expect that strategy in the future, which triggers a best-response that benefits the teacher. Furthermore, reversion to the Pareto-inefficient equilibrium is credible because the teacher knows that if she defects, her adaptive opponent will learn to quit playing the repeated-game strategy.

Strategic teaching is a different way to comprehend repeated-game equilibria than standard analyses, and could prove *better* as a way of explaining actual behavior. Consider the influence of the length of the horizon of future play. In standard (pre-1980) theory, folk theorem results unravel when the horizon of the repeated game is finite. Standard theory therefore cannot explain why players cooperate until the last couple of periods of a finitely-repeated PD, as is typically observed in experiments (e.g., Selten and Stoeker [70]). Since strategic teaching assumes that the learners are adaptive, and do not use backward induction, strategic teaching will generally predict folk theorem-type results until a point near the end of the finite horizon, when it no longer pays to teach because the end is too near. Therefore, strategic teaching does *not* predict unraveling in long finitely-repeated games, which is consistent with most experimental data and everyday intuition but contrary to standard theory.

Of course, it is now well-known that repeated-game behavior can arise in finite-horizon games when there are a small number of "irrational" types (who act like the horizon is unlimited), which creates an incentive for rational players to behave as if the horizon is unlimited until near the end (e.g., Kreps and Wilson [44]). But specifying why some types are irrational, and how many they are, makes this interpretation difficult to test.²³ In the

²³Some models allow the number and nature of irrational types to be a free parameter, as in the "homemade prior" account of Camerer and Weigelt [16] and Palfrey and

teaching approach, which “crazy” type the teacher wants to pretend to be arises endogenously from the payoff structure— they are generally Stackelberg types, who play the strategy they would choose if they could commit to it. For example, in trust games, they would like to commit to behaving nicely; in entry-deterrence, they would like to commit to fighting entry.

4.1. The Model

To illustrate the details of how teaching works, consider the repeated trust game. (Below, we estimate the teaching model on a sample of experimental data for this game.) A borrower B who want to borrow money from each of a series of lenders denoted L_i ($i = 1, \dots, N$). A lender makes only a single lending decision (either *Loan* or *No Loan*) and the borrower makes a string of N decisions each of which, either (*repay* or *default*), is made after observing the lender’s decision.

In a typical experimental session, 11 subjects are divided into 3 borrowers and 8 lenders and their roles are fixed. In a single sequence, a borrower B is randomly chosen to play an 8-round supergame. Each lender L_i ($i = 1, \dots, 8$) plays in 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 many times with fresh random assignments of borrowers, and orders of lenders, in each sequence.²⁴

Denote each sequence of game rounds by k and each game round by t . Note that within the sequence of game rounds, there is a common borrower. In a typical experimental session, there are about 81 sequences. The goal is to specify the probabilities for the borrower and the lender for each of their actions, in each round of each sequence.

Recall that each lender plays only once, and each borrower plays in only a third of the sequences. Yet they watch all the other plays, and clearly respond to observed behavior of others. Therefore, we assume “observational” learning— all lenders learn from what they observe as strongly as from what they do. This assumption is plausible and well-supported by other evidence (e.g., Duffy and Feltovich [24]). It is also necessary to explain what we see (lenders in later rounds who have not played yet clearly react to what happened in earlier rounds of a sequence). We assume that the lenders are purely adaptive (because the matching scheme gives them no incentive to teach) and that the dishonest borrower may be sophisticated, and may also be a strategic teacher. In sessions where there is a possibility of an honest borrower (e.g. sessions 3 to 8 of the trust data),

Rosenthal [58], executed formally by McKelvey and Palfrey [46]. The agent-based QRE model we use below incorporates this idea.

²⁴To limit reputation-building across sequences, if a borrower plays in sequence t she cannot be chosen to play in sequence $t + 1$.

we use a simple logit probability to model the honest behavior.²⁵ Note that the best response of an honest type is always to repay and a simple logit form is used to allow for trembling. Henceforth, all borrowers in our discussion are meant to be the dishonest type.

An important twist in our model is that players are assumed to learn about the attraction of a strategy in a current round in two separate ways: They learn from previous rounds in a sequence; *and* from how the strategies performed in the current round in previous sequences. For concreteness, consider round 7 in sequence 14. A lender presumably sees what happened in the previous 6 rounds, and learns about whether to loan from what happened in those rounds. It is also plausible that the lender looks at what happened in the 7th round of previous sequences 1-13, and learns about whether she should loan in round 7 from those sequences.

We include both types of learning in the model. Learning about a specific round across sequences is like repeated-stage-game learning across similar games; where the “similar” games are identical rounds in previous sequences. This sort of transfer has been explored by Stahl [72] and resembles the similarity-based “spillover” of reinforcement from a chosen strategy to neighboring strategies explored by Sarin and Vahid [64]. And we thought that including cross-sequence learning was necessary to explain the data better. After all, the reason why experimenters conduct a (long!) series of repeated game sequences, rather than simply one, is presumably a prior belief that learning required many repetitions of the entire sequence.

The strength of cross-sequence learning is parameterized by τ . If that parameter is zero there is no cross-sequence learning. So the data can tell us whether allowing cross-sequence learning is helpful through the value of τ .

It is not clear how to integrate the two sorts of learning. Returning to our example, the strategy Loan for a lender before period 7 of sequence 14 can be said to have two different attractions - the attraction of Loan after period 6, and the attraction of Loan after period 7 of sequence 13. Simply averaging these attractions is an obvious, but hamfisted, way to include them both in a general learning process. 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 previous sequences.

Call the attraction of an adaptive lender at the end of sequence k and round t for strategy j , $A_L^j(a, k, t)$. Updating occurs in 2 steps. The idea

²⁵For example, using the payoff function from the trust data, the honest type will choose to repay with probability $P_H = \frac{e^{\lambda_H \cdot 60}}{e^{\lambda_H \cdot 60} + e^{\lambda_H \cdot 0}}$ where λ_H is a scale parameter to be estimated.

is to create an “interim” attraction for round t , $B_L^j(a, k, t)$, based on the attraction $A_L^j(a, k, t-1)$ and payoff from the round t , then incorporate experience in round $t+1$ from previous sequences, transforming $B_L^j(a, k, t)$ into a final attraction $A_L^j(a, k, t)$.

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

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

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

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

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

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

We assume that the learning about an upcoming round from previous sequences is driven by the average payoff in that round in previous sequences. Formally, $\hat{\pi}_L^j(k, t+1) = \sum_{m=1}^{k-1} \pi_L(j, s_B(m, t+1)) / (k-1)$.²⁶ As usual, we derive $P_L^j(a, k, t+1)$ from $A_L^j(a, k, t)$.

Next we specify learning by an **adaptive** borrower. The updating occurs in 2 steps.

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

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

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

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

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

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

²⁶We have also explored a specification in which only the payoff received in the previous sequence from a particular strategy is used. That is, $\hat{\pi}_L^j(k, t+1) = \pi_L(j, s_B(k-1, t+1))$. That specification is too “fickle” and fits worse than the average-payoff specification.

As above, we assume that the learning from an upcoming round, from previous sequences, is driven by the average payoff in that round in previous sequences ($\hat{\pi}_B^j(k, t+1) = \sum_{m=1}^{k-1} \pi_B(j, s_L(m, t+1))/(k-1)$). We derive $P_B^j(a, k, t+1)$ from $A_B^j(a, k, t)$.

Now we are ready to specify how a **sophisticated** borrower will behave. A sophisticated borrower guesses how the lender learns, and adapts those guesses to experience, and also plans actions for the remaining periods within a game sequence. Specifically, we assume a sophisticated borrower's attractions are specified as follows:

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

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

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 economize in computing, we search only paths of future actions that always have default following repay because the reverse behavior (repay following default) generates a lower return. It is therefore in the sophisticated borrower's interest to allocate repay to earlier rounds. Finally, $P_B^j(s, k, t+1)$ is derived from $A_B^j(s, k, t)$ using a logit rule.

Note that if $\epsilon = 0$, the player is sophisticated but myopic (she does not take into account the future learning effects of current actions). If $\epsilon > 0$, the sophisticated player is a teacher who takes into account the effects of current actions on learned behavior of others.

Assume that a proportion α of the borrowers are sophisticated. Then the likelihood of observing the dishonest borrowers is given by $\Pi_k[(1-\alpha) \cdot \Pi_t P_B^{S_B(t)}(a, k, t) + \alpha \cdot \Pi_t P_B^{S_B(t)}(s, k, t)]$.²⁷

We estimate the model using repeated game trust data from Camerer and Weigelt [16]. As in our earlier work, we use maximum likelihood estimation (MLE) to calibrate the model on about 70% of the sequences in each experimental session, then forecast behavior in the remaining 30% of

²⁷We assume rational expectations (i.e., $\alpha = \alpha'$) to conserve on parameters.

the sequences.²⁸ If the model fits better in-sample only by overfitting, it will perform poorly out-of-sample.²⁹

4.2. Repeated Trust Game

Table 3 shows payoffs in the repeated trust game. The lenders earn 10 if they do not lend; they earn 40 if a loan is repaid and lose -100 if the borrower defaults.³⁰ A normal borrower earns 10 if the lender does not lend, 150 if the lender lends and she defaults, and earns only 60 if she pays back. Honest-type borrower have default and repayment payoffs of 0 and 60 respectively (note that they earn more from repaying).

The probability that a borrower had honest-type payoffs in a particular sequence, $P(\text{honest})$, was .33 (sessions 3-5), .10 (sessions 6-8) and 0 (sessions 9-10). Subjects were MBA students at NYU or University of Pennsylvania. They were paid according to performance and earned an average of \$18 for a 2-1/2 hour session. Each session had 70-101 eight-period sequences.

We now discuss the sequential equilibrium predictions, then return to the data. With the payoffs used in Table 3, the analysis proceeds as follows: Start from period 8. In this period, lenders know that the borrower will play Default if she loans (and the honest type will, of course, repay) so the only question is the probability that the borrower is a honest type. Simple algebra shows that if lenders are risk-neutral, they should loan if $P(\text{Honest})$ is above $55/70$, about .79. Define this important threshold to be γ . In the second-to-last period, period 7, normal borrowers are torn between two forces: Conditional on loan, they would like to choose Default to earn the higher payoff; but if they do so, they would have revealed their type and would earn the No-Loan payoff in period 8. However, if their reputation (i.e., the perception $P(\text{Honest})$ lenders have) in period 7 is below γ , then lenders will not lend and Bayesian updating would lead lenders to have the same perception in period 8 which, by assumption, is too low to induce the lenders to lend in period 8. The trick is for borrowers to play a mixed strategy, repaying frequently enough that if they do default, the updated $P(\text{Honest})$ will equal γ , so that lenders will mix and sometimes lend in the last period.

²⁸The likelihood of a session consists of the likelihood for the lenders, the likelihood for the dishonest borrowers and, in some sessions, the likelihood for the honest borrowers.

²⁹We used GAUSS. To avoid settling into local maxima, we 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.

³⁰Payoffs were varied for lenders for the loan-default outcome, -50 in sessions 6-8 and -75 in sessions 9-10. These parameter variations provide a small 'stress test' for whether the same structural model can account for behavior across sessions with minimal parameter variation.

Given a particular $P(\text{Honest})$ in period 7, the normal borrower should choose a repayment probability p which keeps the lender indifferent in period 7, and allows Bayesian updating of his reputation to the threshold $P(\text{Honest})=\gamma$ in period 8. Combining these two conditions gives a threshold of perceived $P(\text{Honest})$ which happens to be γ^2 , and a mixed strategy probability of lending in period 7 of .560.

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 $P(\text{Honest})$ values is simply γ^n . Figure 4 shows this path, and the mixed-strategy probabilities of repayment by normal borrowers which keep the lender's perceptions along this path (for an initial prior $P(\text{Honest})$ of .33). The Figure can be used to illustrate all the key properties of this equilibrium.³¹ In the first three periods, the threshold $P(\text{Honest})$ is below the prior of .33, so borrowers can "afford" to always default and lenders should lend. Beginning in period 4, normal borrowers must mix in order to boost their reputation, conditional on getting a loan, to stay along the equilibrium path of $P(\text{Honest})$ which increases. If the borrower ever defaults, the lender should not lend in all subsequent periods.

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)?

Typical patterns in the data can be seen in Figures 5a-b. The figures show relative frequencies of No Loan and Default (conditional on a loan).³² Sequences are combined into ten-sequence blocks (denoted "sequence") and average frequencies are reported from those blocks. Periods 1,...,8 denote periods in each sequence.

Figure 5a-b are data from all eight sessions pooled. Lenders start by making loans in early periods (i.e., there is a low frequency of no-loan), but they rarely lend in periods 7-8. Borrowers default infrequently in early periods, but usually default in periods 7-8. The within-sequence pattern is particularly dramatic in later sequences.

As a benchmark alternative to the teaching model, we estimated an agent-based version of quantal response equilibrium suitable for extensive-form games (see McKelvey and Palfrey [48]). An Appendix explains pre-

³¹Characteristically, there are other sequential equilibria. For example, the normal borrower might never repay, if she thinks that the lender will perceive Repay as an indication of a normal type. The intuitive criterion selects the equilibrium we discuss however, so we will casually refer to it as "the" equilibrium.

³²To distinguish cases in which there are no data from true zeros (e.g., no repay after several loans), we plot cases with no data as -.1.

cisely how the the agent-QRE model is implemented and estimated.³³ (Note that while the model has fewer free parameters than the teaching model, there are many nuances in implementation which make it more difficult to work with in some respects.) We use this model, rather than sequential equilibrium, because the (intuitive) sequential equilibrium predicts many events to have zero probability, so some notion of error or trembling is needed to fit the data (otherwise the logarithm of likelihood explodes). Agent-QRE is a plausible form and fits many data sets well (see McKelvey and Palfrey [48], Goeree and Holt [33]).

We implement the model with four parameters – the prior belief of lenders about $P(\text{Honest})$ (which can differ from the prior induced by the experimental design to reflect “homemade priors”, giving the model more flexibility), and different response sensitivities λ for lenders, honest borrowers, and normal borrowers. Agent-QRE is a good benchmark because it incorporates the key features of repeated-game equilibrium– (stochastic) optimization, accurate expectations about actions of other players, and Bayesian updating. Also, while it makes the same conditional predictions in every sequence, even under AQRE players will be predicted to behave differently in later sequences if the early-period play in those sequences is different, so in a crude way it can fit cross-sequence change. AQRE therefore presents a stiff challenge to any adaptive learning model which tries to explain learning both within and across sequences.

The models are estimated separately on each of the eight sessions to gauge cross-session stability (pooling sessions yields similar results about relative fit of teaching and AQRE). Table 4 shows measures of fit. To measure fit we report both log likelihood and the average predicted probability of events that occurred. Table 4 shows that for in-sample calibration, the average predicted probabilities range from 71% to 88% for the teaching model, compared to 69% to 81% for agent-QRE. The teaching model fits better in every session. Of course, an important test for overfitting is how a model fares in out-of-sample forecasting. The average predicted probabilities for the teaching model range from 71% to 89% with no noticeable decline in performance. The agent-QRE model probabilities range from 68% to 82% with a slight decline in 2 out of the 8 sessions and always fits worse than the teaching model by this measure. The log likelihood measure yields a similar result: The teaching model fits substantially better than agent-QRE in all sessions.

Table 5 shows parameter values (and standard errors) for the teaching model (excluding the λ values).³⁴ The interesting lender parameters are δ and τ . The weights on foregone payoff, δ range from .19 to .80, and average

³³We use an agent-based form in which players choose a distribution of strategies at each node, rather than using a distribution over all history-dependent strategies.

³⁴The remaining parameter estimates are available from the authors upon request.

.44, with low standard errors (.01). Estimates of τ range from .95 to 1.00, except for one outlier at .51, which indicates a high degree of cross-sequence learning.

The interesting parameters for sophisticated borrowers are α and ϵ . The degree of sophistication α ranges from .02 to .70. The “horizon” parameter ϵ is close to one in five sessions and around .20 in three sessions.³⁵

Figures 5c-d show average predicted probabilities from the teaching model for the no-loan and default rate (conditional on no previous default). It captures most of the key regularities in the data, although it does not always fit aggressive patterns in the data. No-loan frequencies are predicted to start low and rise across periods, as they do. There are some slight downward trends over time (e.g., periods 1-4) in the data, which the model captures in periods 2-4. Since the model is forecasting out-of-sample in sequence blocks 7-9, its ability to forecast those trends in those sequences is particularly noteworthy. Notice that the no-loan rate pops up quite a bit in periods 3-5 in those later sequences, and the model predicts an increase as well. There is less change in the default rate across sequences for the model to capture. It does not pick up the drop in default rate in early periods across sequences well, but it does predict the rate of increase in default rate across periods reasonably well, except for under-predicting default in the last period.

The key parameter in AQRE is the prior $P(\text{Honest})$ — including any “homemade prior” belief that borrowers with normal payoffs will behave honestly by always repaying. The estimates are .93, .92, 1.0, 1.0, .82, .67, .51, .53. These correlate .78 with the prior induced by the experimenter. However, these figures just seem too high. Earlier studies of this type estimate numbers on the order of 10-20% with zero induced prior (see Camerer and Weigelt [16], McKelvey and Palfrey [46], Neral and Ochs [55]).

4.3. Distinguishing strategic teaching from type-based reputation formation

Strategic teaching generates behavior which is similar to reputation-building in repeated games where there is Bayesian updating of players’ unobserved “types”. In the type-based models the presence of honest types who prefer to repay creates an incentive (depending on payoffs and the prior $P(\text{Honest})$) for types with normal payoffs to repay loans.

Many of the predictions of the teaching and type-based models across periods and changes in payoffs are similar. The crucial difference is that in the type-based models a particular *player* has a reputation (i.e., a posterior $P(\text{Honest})$ and associated probabilities of repayment in each period). In the teaching model a *strategy* has a “reputation” or attraction. More

³⁵Two of the three sessions with low $\hat{\epsilon}$ are sessions with a zero-prior on the honest type, where reputation-building is less common.

precisely, there are four important differences between the teaching and types approaches: Sensitivity to priors, independence of own payoffs and own mixture probabilities, the effect of missed opportunity, and “no second chances”. We sketch these differences here and plan to explore them further in future work.

Type-based models have the following properties: If the prior $P(\text{Honest})$ is below some threshold (depending on payoffs and the horizon of the finitely-repeated game) there is no reputation-building; mixture probabilities depend only on the other players’ payoffs; if a borrower does not receive a loan in an early period that missed opportunity does not affect future behavior, but if a borrower does not receive a loan in a later period (and hence has no chance to repay and build reputation) then she never receives any more loans; and the sensible sequential equilibrium requires the assumption that if a borrower defaults in one period, then repays in a subsequent period, her reputation is not restored (there are “no second chances”). The teaching model does not make these same predictions. When the type-based models are extended to include quantal-response and a “homemade prior”, as in the AQRE model we used as a static benchmark, the sensitivity to priors and independence properties no longer hold, but the missed-opportunity and no-second-chances properties still hold (in a probabilistic sense).

There is a simple experimental way to discriminate between the teaching and type-based approaches. The type-based approach requires that a player’s type remain fixed throughout a sequence. If types are drawn randomly in each period, the link between past behavior, inferences about types, and future incentives is broken and there is no equilibrium reputation-building by normal types. In the teaching approach the presence of nice types *does* matter³⁶ but it makes little difference whether a player’s type is fixed within a sequence or drawn independently in each period. Comparing behavior in experiments with fixed types and independent types therefore provides a way to distinguish the type-based and teaching approaches. The type-based approach predicts a big difference in behavior across those two protocols, while the teaching approach predicts little difference.

We do not suggest that the teaching approach should completely replace type-based equilibrium models of reputation-formation. However, it has always seemed dubious that players are capable of the delicate balance of reasoning required to implement the type-based models, unless they learn the equilibrium through some adaptive process. The teaching model is the only available parametric model of that process in which matching protocols matter, and is therefore worth exploring further. Note also that

³⁶Presence of honest types matters because they alter the attractions of loan and no-loan strategies for lenders, and hence alter the marginal incentives to teach.

the type-based models assume optimization and foresight by reputation-builders, and Bayesian updating of types by “learners”. The teaching model only changes the last feature, replacing Bayesian updating by learners with learning about their strategies. Our adaptive EWA work showed that some Bayesian learning models which are used to compute expected payoffs (weighted fictitious play) can be perfectly operationalized by generalized reinforcement which keeps track of historical payoffs. In a similar way, assuming that entrants update strategy attractions may be a sensible empirical alternative to Bayesian updating of types, and softens the sharp predictions which result from the Bayesian approach (particularly the missed opportunity and no-second-chance features).

5. CONCLUSION

This paper extends earlier work on adaptive EWA learning to include sophistication and strategic teaching. Before proceeding to summarize our conclusions, it is helpful to think of the properties one would like an empirical model to have. (i) The model should use all the information that subjects have and use. (Reinforcement and belief models don’t have this property; see Salmon [66].) (ii) The parameters of the model should have psychological interpretations, preferably consistent with accepted ideas in neighboring social sciences. (iii) The model should be as simple as possible, in the sense that every parameter should play a distinct role that is predictively useful. (iv) The model should fit well, both in- and out-of-sample, judging by statistical criteria which permit model comparison.

EWA does well on all four criteria. A fifth property is that a model be tractable enough to explore its theoretical implications. Heller and Sarin [34] have made initial progress, using a variant of EWA, which is promising. In current work, we have endogenized EWA parameters, making them functions of experience (Ho et al. [35]). Our method opens the door to easier theorizing.³⁷

Adaptive EWA is incomplete by the information-use (i) and psychological fidelity criteria (ii), because it does not explain how players’ information about the payoffs of others is used, and it does not allow the sort of anticipatory learning which is plausible for intelligent experienced players. Therefore, we extended the model by assuming some fraction α of players are sophisticated in a specific sense: They believe others adapt according to EWA, but also believe that a fraction α' are sophisticated like themselves.

³⁷In the one-parameter “EWA Lite” of Ho et al. [35], when an opponent’s behavior stabilize the parameters ϕ and δ both converge toward one, which means the learning rule converges toward fictitious play (if $\kappa = 0$). When another player’s behavior stabilizes we can therefore apply convergence theorems which are used to show convergence of fictitious play in some settings (see Fudenberg and Levine [29]).

We estimate the sophisticated EWA model on a sample of data from dominance-solvable ‘ p -beauty contest’ games. In these games, each of the n players choose a number from the interval $[0,100]$ and the player whose number is closest to p times the average number wins a fixed prize. We chose these games because there is substantial learning evident in the data, but the adaptive EWA model (and the special cases of reinforcement and weighted fictitious play) do not fit particularly well (see Camerer and Ho [12]).

Introducing sophistication improves fit substantially. More interestingly, we find that the estimated fraction of sophisticated players, α , rises substantially between sessions with inexperienced subjects and those with experienced subjects (who play a second p -beauty contest with a different value of p). This shows that what experience creates is not just learning about the success of strategies, but also *learning about learning*—or increasing sophistication. Players seem to learn that others are adaptive, and learn to “jump ahead” by anticipating changes by others.³⁸

Once sophistication is introduced, whether players will be matched together repeatedly or not could matter. Sophisticated players who understand that others are learning will have an incentive to take actions in period t , which “teach” adaptive players how strategies perform, so the sophisticated can earn a higher payoff in period $t+1$ and beyond. (If players are rematched in each period, then this kind of teaching motivation disappears.)

Strategic teaching captures the incentive players have, in repeated games with fixed partners, to implement repeated-game strategies which can lead to results that are not stage-game equilibria. We explore this possibility in borrower-lender trust games. In the trust game, borrowers have an incentive to repay loans in early periods, in order to obtain further loans, but toward the end of the eight-period horizon they should quit repaying. We show that a model with adaptive lenders, and sophisticated borrowers who “strategically teach”, can explain the basic patterns in these data reasonably well (and consistently better than agent-based quantal response equilibrium).

The teaching approach shows promise for capturing much of the intuition and empirical regularity of reputation-building in repeated games, without using the complicated type-based equilibrium approach. The device of assuming updated types is useful for explaining why lenders are afraid to loan early, and willing to loan late. Sophisticated EWA with strategic teaching produces the same effect more directly—borrowers have an incentive to repay early because they know that lenders will be convinced, not

³⁸A similar point is made, with quite a different model, by Stahl [71]).

because they believe the borrower's reputations per se, but simply because they learn that loaning in early periods is good.

The teaching model helps resolve a mystery in the experimental literature on repeated games. Basic patterns in the data do go in the direction predicted by sequential equilibrium— viz., borrowers repay more often in early periods of a sequence, and lenders seem to anticipate or learn this, loan more frequently in earlier periods. But changes in treatment variables do not always create predicted changes in behavior (see Neral and Ochs [55] and Jung, Kagel and Levin [41]), subtle predictions of the equilibrium theory are not confirmed, and equilibrium calculations are so complex that it is hard to believe subjects are calculating rather than learning. As a result, Camerer and Weigelt [16] concluded their paper as follows:

...the long period of disequilibrium behavior early in these experiments raises the important question of how people learn to play complicated games. The data could be fit to statistical learning models, though new experiments or new models might be needed to explain learning adequately. (pp 27-28).

Strategic teaching is one possible answer to the question they raised almost 15 years ago.

**APPENDIX: COMPUTING AGENT QUANTAL
RESPONSE EQUILIBRIUM IN REPEATED TRUSTED
GAMES**

The goal is to compute the AQRE given λ_H (honest borrower), λ_D (dishonest borrower), λ_L (lender), and q ($P(\text{Honest})$). We specify $P(\text{Honest})$ to be a function of the experimenter-induced prior (of the honest type), p and the home-made prior θ . In particular, we let $q = p + \theta \cdot (1 - p)$ where θ will be estimated.

Let's get the easy thing out of the way first. The honest type will choose "repay" with a probability of $P_H = \frac{e^{\lambda_H \cdot 60}}{e^{\lambda_H \cdot 60} + e^{\lambda_H \cdot 0}}$. Note that this probability applies to every period and is independent of lender's belief on borrower's type.

A.1. THREE MATRICES

The computation involves constructing 3 matrices which have a dimension 101 x 8 (101 belief grids about borrower type and 8 time periods, with beliefs rounded to the nearest .01):

1. **Dishonest Borrower Matrix:** Each row of the matrix corresponds to a belief grid and each column a time period. For example, row 1 corresponds to a belief of 0.0 and row 65 a belief of 0.64. Similarly, column 3 corresponds to period 3, and etc. Element (r, t) of the Dishonest Borrower matrix gives the probability that the dishonest borrower will repay given that lender's belief is r at time t . Formally, we have

$$P_D(t|r) = \frac{e^{\lambda_D \cdot (60 + V(t+1|r'))}}{e^{\lambda_D \cdot (60 + V(t+1|r'))} + e^{\lambda_D \cdot (150 + V(t+1|r''))}} \quad (\text{A.1})$$

where $V(t+1|r')$ refers to the *ex ante* value of the borrower for future rounds of the game given lender's posterior belief r' at time $t+1$. r' and r'' are the posterior beliefs of choosing repay and default respectively. Formally, we have:

$$r' = \frac{P_H \cdot r}{P_H \cdot r + P_D(t|r) \cdot (1 - r)} \quad (\text{A.2})$$

$$r'' = \frac{(1 - P_H) \cdot r}{(1 - P_H) \cdot r + (1 - P_D(t|r)) \cdot (1 - r)} \quad (\text{A.3})$$

2. **Lender Matrix:** Element (r,t) of this matrix gives, $P_L(t|r)$, the probability of lending given that the lender's belief of borrower being honest is

r at period t . In general,

$$P_L(t|r) = \frac{e^{\lambda_L \cdot [r \cdot (P_H \cdot 40 + (1 - P_H) \cdot (-100)) + (1 - r) \cdot (P_D(t|r) \cdot 40 + (1 - P_D(t|r)) \cdot (-100))]} }{e^{\lambda_L \cdot [r \cdot (P_H \cdot 40 + (1 - P_H) \cdot (-100)) + (1 - r) \cdot (P_D(t|r) \cdot 40 + (1 - P_D(t|r)) \cdot (-100))]} + e^{\lambda_L \cdot 10}} \quad (\text{A.4})$$

3. Value Matrix: Element (r,t) of the Value matrix provides the *ex ante* value for the borrower (as in a dynamic program) for future rounds of the game when the lender's belief is r at time t . Formally, it is given by:

$$\begin{aligned} V(t|r) = & (1 - P_L(t|r)) \cdot (10 + V(t + 1|r)) \\ & + P_L(t|r) \cdot [P_D(t|r) \cdot (60 + V(t + 1|r')) \\ & + (1 - P_D(t|r)) \cdot (150 + V(t + 1|r''))] \end{aligned} \quad (\text{A.5})$$

4. Observations:

(i) Equations (A.1)-(A.3) define a fixed point $P_D(t|r)$. Note that we need to know what $V(t + 1|r')$ and $V(t + 1|r'')$ are in order to solve the fixed point. This suggests we need to work backward.

(ii) Equation (A.4) can easily be solved once the fixed point is found.

(iii) Equation (A.5) will only be used when we create elements in column $t - 1$ of the value matrix. Computing it is easy after finding the fixed point.

(iv) Note that in the case where No Loan is made in a period, the belief in that period will not be updated.

5. Ending Values: We have to create these matrices by working backward (i.e., starting from column 8). The ending dishonest borrower probabilities are computed as follows:

$$P_D(8|r) = \frac{e^{\lambda_D \cdot 60}}{e^{\lambda_D \cdot 60} + e^{\lambda_D \cdot 150}} \quad (\text{A.6})$$

Similarly, the ending values $V(8|r)$ as follows:

$$\begin{aligned} V(8|r) = & (1 - P_L(8|r)) \cdot 10 \\ & + P_L(8|r) \cdot [60 \cdot P_D(8|r) + 150 \cdot (1 - P_D(8|r))]. \end{aligned} \quad (\text{A.7})$$

A.2. SOLVING THE FIXED POINT

Let us see how we can compute $P_D(7|r)$. We have the following set of equations:

$$P_D(7|r) = \frac{e^{\lambda_D \cdot (60 + V(8|r'))}}{e^{\lambda_D \cdot (60 + V(8|r'))} + e^{\lambda_D \cdot (150 + V(8|r''))}} \quad (\text{A.8})$$

$$r' = \frac{P_H \cdot r}{P_H \cdot r + P_D(7|r) \cdot (1 - r)} \quad (\text{A.9})$$

$$r'' = \frac{(1 - P_H) \cdot r}{(1 - P_H) \cdot r + (1 - P_D(7|r)) \cdot (1 - r)} \quad (\text{A.10})$$

We use $P_D(8|r)$ as the initial value and plug it into (A.9) and (A.10) to find r' and r'' . Round r' and r'' to the nearest 0.01 values. Use them to look for $V(8|r')$ and $V(8|r'')$ in the value matrix and substitute these values into (A.8) to get a new $P_D(7|r)$. Altogether we have 101x7 fixed points to solve.

A.3. LIKELIHOOD CALCULATION

Assume sequence k , we have {(Lend, Repay), (Lend, Repay), (No Loan), (No Loan), (Lend, Default), (Lend, Default), (No Loan), (No Loan)} So the likelihood of lender is:

$$\begin{aligned} \text{Likelihood}_L(k) = & P_L(1|x(1)) \cdot P_L(2|x(2)) \cdot (1 - P_L(3|x(3))) \cdot (1 - P_L(4|x(4))) \\ & \cdot P_L(5|x(5)) \cdot P_L(6|x(6)) \cdot (1 - P_L(7|x(7))) \cdot (1 - P_L(8|x(8))) \end{aligned}$$

where

$$\begin{aligned} x(1) &= p + \theta \cdot (1 - p) \\ x(2) &= \frac{P_H \cdot x(1)}{P_H \cdot x(1) + P_D(1|x(1)) \cdot (1 - x(1))} \\ x(3) &= \frac{P_H \cdot x(2)}{P_H \cdot x(2) + P_D(2|x(2)) \cdot (1 - x(2))} \\ x(4) &= x(3) \\ x(5) &= x(4) \\ x(6) &= \frac{(1 - P_H) \cdot x(5)}{(1 - P_H) \cdot x(5) + (1 - P_D(5|x(5))) \cdot (1 - x(5))} \\ x(7) &= \frac{(1 - P_H) \cdot x(6)}{(1 - P_H) \cdot x(6) + (1 - P_D(6|x(6))) \cdot (1 - x(6))} \\ x(8) &= x(7) \end{aligned}$$

The conditional likelihood of the dishonest borrower is:

$$CLikelihood_D(k) = P_D(1|x(1)) \cdot P_D(2|x(2)) \\ \cdot (1 - P_D(5|x(5))) \cdot (1 - P_D(6|x(6)))$$

The conditional likelihood of the honest borrower is:

$$CLikelihood_H(k) = P_H \cdot P_H \cdot (1 - P_H) \cdot (1 - P_H)$$

A minor complication arises because while we observe the borrower's *pay-off* type, we cannot tell directly whether a borrower with dishonest-type payoffs is playing like an honest type (which is allowed by the homemade prior component $\theta > 0$). But the conditional probability we would like to use when observing dishonest-payoff types is a (weighted) mixture of the probabilities with which dishonest types and those who behave honestly repay. There are really 2 types of borrowers endowed with dishonest payoff matrix: one who plays dishonest (DD) and one who actually plays honest (DH). However, we are not able to directly distinguish these types. The trick is we assume that DH plays like an honest borrower with $P_{DH} = P_H$ where P_{DH} is the choice probability of DH type. Note that the prior probability $P(DH | \text{plays like an honest borrower}) = \frac{\theta \cdot (1-p)}{p + \theta \cdot (1-p)}$. (That is the prior chance that the borrower has dishonest payoffs conditional on playing like a honest borrower. Note that if $\theta = 0$ that prior is zero— i.e., any borrower who plays like an honest one has honest payoffs. And if $\theta = 1$ — everyone plays honestly— then the conditional probability is 1-p.) The key insight is that since honest-payoff types and those with dishonest payoffs who play honestly play in *exactly* the same way (by definition), then the posterior probabilities of DH type and true honest type will always be the same as the priors for any observed data (since the likelihood ratio is always one). In other words, we have $\frac{P(DH,t)}{P(H,t)} = \frac{\theta \cdot (1-p)}{p}$ where $P(y,t)$ is the perceived probability at period t that the player is type y (subscript everywhere for sequence k). We can use this property to create the correct weighted average of predicted behavior when we observe a borrower's pay-offs but do not directly observe whether they are of the type that always plays honestly. Note that $P(DH,t) + P(H,t) = x(t)$.

If the borrower in sequence k is given a dishonest payoff matrix, the likelihood of observing a repay in period 1 is given by

$$P_B(1|x(1)) = \frac{P_D(1|x(1)) \cdot P(DD,1) + P_{DH} \cdot P(DH,1)}{P(DD,1) + P(DH,1)}$$

$$= \frac{P_D(1|x(1)) \cdot (1 - x(1)) + P_H \cdot \left[\frac{\theta \cdot (1-p)}{p+\theta \cdot (1-p)}\right] \cdot x(1)}{1 - x(1) + \left[\frac{\theta \cdot (1-p)}{p+\theta \cdot (1-p)}\right] \cdot x(1)}$$

Hence, the likelihood of observing a borrower given a dishonest payoff function in sequence k above is:

$$\begin{aligned} \text{Likelihood}_B(k) = & P_B(1|x(1)) \cdot P_B(2|x(2)) \\ & \cdot (1 - P_B(5|x(5))) \cdot (1 - P_B(6|x(6))) \end{aligned}$$

REFERENCES

1. W. Amaldoss, "Collaborating to Compete: A Game-theoretical Model and Experimental Investigation of Competition among Alliances," Unpublished Ph.D. dissertation, University of Pennsylvania, 1998.
2. C. Anderson and C. F. Camerer, "Experience-weighted Attraction Learning in Sender-receiver Signaling Games," *Economic Theory*, 16(3), (2001), 689-718.
3. J. Andreoni and J. Miller, "Rational Cooperation in the Finitely Repeated Prisoner's Dilemma: Experimental Evidence," *Economic Journal*, 103 (1993), 570-585.
4. A. Blume, D. DeJong, G. Neumann and N. Savin, "Learning in Sender-receiver Games," University of Iowa working paper, 1999.
5. B. Broseta. "Adaptive Learning and Equilibrium Selection in Experimental Coordination Games: An ARCH(1) Approach," *Games and Economic Behavior*, 32 (2000), 25-50.
6. G. W. Brown, "Iterative Solution of Games by Fictitious Play," in *Activity Analysis of Production and Allocation*, John Wiley & Sons, New York, 1951.
7. A. Cabrales and W. Garcia-Fontes, "Estimating Learning Models with Experimental Data," University of Pompeu Febrer working paper 501, 2000.
8. G. P. Cachon and C. F. Camerer, "Loss-avoidance and Forward Induction in Experimental Coordination Games," *The Quarterly Journal of Economics*, 111 (1996), 165-194.
9. C. F. Camerer, "Progress in Behavioral Game Theory," *Journal of Economic Perspectives*, 11 (1997), 167-188.
10. C.F. Camerer and T-H Ho, "EWA Learning in Normal-form Games: Probability Rules, Heterogeneity and Time Variation," *Journal of Mathematical Psychology*, 42 (1998), 305-326.
11. C.F. Camerer and T-H Ho, "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., Lawrence Erlbaum Assoc. Inc., New Jersey, 1999, 31-52.
12. C. F. Camerer and T-H Ho, "Experience-weighted Attraction Learning in Normal-form Games," *Econometrica*, 67 (1999), 827-874.
13. C. F. Camerer, T-H Ho, and X. Wang, "Individual Differences and Payoff Learning in Games," University of Pennsylvania working paper, 1999.

14. C. F. Camerer, D. Hsia, and T-H Ho, "EWA Learning in Bilateral Call Markets," in *Experimental Business Research*, ed. by A. Rapoport and R. Zwick, in press.
15. C. F. Camerer, E. Johnson, S. Sen and T. Rymon, "Cognition and Framing in Sequential Bargaining for Gains and Losses," in *Frontiers of Game Theory*, ed. by K. Binmore, A. Kirman, and P. Tani., MIT Press, Cambridge, 27-48, 1993.
16. C. F. Camerer and K. Weigelt, "Experimental Tests of A Sequential Equilibrium Reputation Model," *Econometrica*, 56 (1988), 1-36.
17. H. Chen, J. W. Friedman, and J. F. Thisse, "Boundedly Rational Nash Equilibrium: A Probabilistic Choice Approach," *Games and Economic Behavior*, (in press).
18. Y. Chen and Y. Khoroshilov, "Learning Under Limited Information," University of Michigan working paper, 2000.
19. K. Clark and M. Sefton, "Matching Protocols in Experimental Games," University of Manchester working paper, 1999.
20. D. J. Cooper and C. Stockman, "Fairness, Learning and Constructive Preferences: An Experimental Examination," Case Western Reserve University working paper, 1999.
21. M. Costa-Gomes, V. Crawford, and B. Broseta, "Cognition and Behavior in Normal-form Games: An Experimental Study," *Econometrica*, Sept (2001).
22. V. Crawford, "Adaptive Dynamics in Coordination Games," *Econometrica*, 63 (1995), 103-143.
23. V. Crawford and B. Broseta. "What Price Coordination? The Efficiency-enhancing Effect of Auctioning the Right to Play," *The American Economic Review*, 88 (1998), 198-225.
24. J. Duffy and N. Feltovich, "Does Observation of Others Affect Learning in Strategic Environments? An Experimental Study," *International Journal of Game Theory*, 28 (1999), 131-152.
25. J. Duffy, and R. Nagel, "On the Robustness of Behavior in Experimental 'Beauty Contest' Games," *Economic Journal*, 107 (1997), 1684-1700.
26. G. Ellison, "Learning from Personal Experience: One Rational Guy and the Justification of Myopia," *Games and Economic Behavior*, 19 (1997), 180-210.
27. D. Engelmann and M. Strobel, "The False Consensus Effect Disappears if Representative Information and Monetary Incentives are Given," *Experimental Economics*, 3 (2001), 241-260.
28. I. Erev and A. Roth, "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed-strategy Equilibria," *The American Economic Review*, 88 (1998), 848-881.
29. D. Fudenberg and D. Levine, "Reputation and Equilibrium Selection in Games with A Patient Player," *Econometrica*, 57 (1989), 759-778.
30. D. Fudenberg and D. Levine, *The Theory of Learning in Games*, MIT Press, Cambridge, 1998.
31. J. Gale, K. Binmore, and L. Samuelson, "Learning to be Imperfect: The Ultimatum Game," *Games and Economic Behavior*, 8 (1995), 56-90.
32. R. Gallistel, *The Organization of Learning*, MIT Press, Cambridge, 1990.
33. J.K. Goeree and C. A. Holt, "Stochastic Game Theory: For Playing Games, Not Just for Doing Theory," *Proceedings of the National Academy of Sciences*, 96 (1999), 10564-10567.

34. D. Heller and R. Sarin, "Parametric Adaptive Learning," *Econometric Society World Congress 2000 Contributed Papers*, No. 1496, 2000.
35. T-H Ho, C. F. Camerer, and J. K. Chong, "Economic Value of EWA Lite: A Functional Theory of Learning in Games," University of Pennsylvania working paper, 2001.
36. T-H Ho, C. F. Camerer, and K. Weigelt, "Iterated Dominance and Iterated Best-response in p -Beauty Contests," *The American Economic Review*, 88 (1998), 947-969.
37. T-H Ho and J. K. Chong, "A Parsimonious Model of SKU Choice," University of Pennsylvania working paper, 1999.
38. C. A. Holt, "Industrial Organization: A Survey of Laboratory Research," in *Handbook of Experimental Economics* ed. by J. Kagel and A. Roth, Princeton: Princeton University Press, 1995, 349-443.
39. E. Hopkins, "Two Competing Models of How People Learn in Games," University of Edinburgh and Pittsburgh working paper, 1999.
40. D. Hsia, "Learning in Single Call Markets," Unpublished Ph.D. Dissertation, University of Southern California, 1999.
41. Y. J. Jung, J. H. Kagel, and D. Levin, "On the Existence of Predatory Pricing: An Experimental Study of Reputation and Entry Deterrence in the Chain-store Game," *RAND Journal of Economics*, 25 (1994), 72-93.
42. M.J. Knez and C. F. Camerer, "Increasing Cooperation in Prisoner's Dilemma by Establishing a Precedent of Efficiency in Coordination Games," *Organizational Behavior and Human Decision Processes*, 82 (2000), 194-216.
43. E. Kohlberg and J-F Mertens, "On the Strategic Stability of Equilibria," *Econometrica*, 54 (1986), 1003-1038.
44. D. Kreps and R. Wilson, "Reputation and Imperfect Information," *Journal of Economic Theory*, 27 (1982), 253-279.
45. P. H. McAllister, "Adaptive approaches to stochastic programming," *Annals of Operations Research*, 30 (1991), 45-62.
46. R. D. McKelvey and T. R. Palfrey, "An Experimental Study of the Centipede Game," *Econometrica*, 60 (1992), 803-836.
47. R. D. McKelvey and T. R. Palfrey, "Quantal Response Equilibria for Normal-form Games," *Games and Economic Behavior*, 7 (1995), 6-38.
48. R. D. McKelvey and T. R. Palfrey, "Quantal Response Equilibria for Extensive-form Games," *Experimental Economics*, 1 (1998), 9-41.
49. R. D. McKelvey and G. Riddihough, "Editorial: The *hard sciences*" *Proceedings of the National Academy of Sciences*, 96 (1999), 10549-10549.
50. D. Mookerjee and B. Sopher, "Learning Behavior in an Experimental Matching Pennies Game," *Games and Economic Behavior*, 7 (1994), 62-91.
51. D. Mookerjee and B. Sopher, "Learning and Decision Costs in Experimental Constant-sum Games," *Games and Economic Behavior*, 19 (1997), 97-132.
52. J. Morgan and M. Sefton, "An Empirical Investigation of Unprofitable Games," Princeton University working paper, 1999.
53. R. Nagel, "Experimental Results on Interactive Competitive Guessing," *The American Economic Review*, 85 (1995), 1313-1326.

54. R. Nagel, "A Review of Beauty Contest Games," in *Games and Human Behavior: Essays in honor of Amnon Rapoport*, ed. by D. Budescu, I. Erev, and R. Zwick., Lawrence Erlbaum Assoc. Inc., New Jersey, 1999, 105-142.
55. J. Neral and J. Ochs, "The Sequential Equilibrium Theory of Reputation Building: A Further Test," *Econometrica*, 60 (1992), 1151-1169.
56. Y. Nyarko and A. Schotter, "An experimental study of belief learning using elicited beliefs," *Econometrica*, in press.
57. J. Ochs, "Entry in Experimental Market Games," in *Games and Human Behavior: Essays in honor of Amnon Rapoport*, ed. by D. Budescu, I. Erev, and R. Zwick., Lawrence Erlbaum Assoc. Inc., New Jersey, 1999.
58. T. R. Palfrey and H. Rosenthal, "Private Incentives in Social Dilemmas: The Effects of Incomplete Information and Altruism," *Journal of Public Economics*, 35 (1988), 309-332.
59. J. Partow and A. Schotter, "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, 1993.
60. A. Rapoport and W. Amaldoss, "Mixed Strategies and Iterative Elimination of Strongly Dominated Strategies: An Experimental Investigation of States of Knowledge," *Journal of Economic Behavior and Organization*, 42 (2000), 483-521.
61. A. Rapoport, A. K-C Lo, and R. Zwick, "Choice of Prizes Allocated by Multiple Lotteries with Endogenously Determined Probabilities," University of Arizona, Department of Management and Policy working paper, 1999.
62. R. W. Rosenthal, "Games of Perfect Information, Predatory Pricing and Chain-Store Paradox," *Journal of Economic Theory*, 25 (1981), 92-100.
63. A. Roth and I. Erev, "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," *Games and Economic Behavior*, 8 (1995), 164-212.
64. R. Sarin and F. Vahid, "Strategy Similarity and Coordination," Texas A&M University working paper, 2000.
65. R. Sarin and F. Vahid, "Predicting How People Play Games: A Simple Dynamic Model of Choice," *Games and Economic Behavior*, 34 (2001), 104-122.
66. T. Salmon, "Evidence for 'Learning to Learn' Behavior in Normal-form Games," Caltech working paper, 1999.
67. T. Salmon, "An Evaluation of Econometric Models of Adaptive Learning," *Econometrica*, (in press).
68. R. Selten, "Anticipatory learning in 2-person Games," ed. by R. Selten, *Game Equilibrium Models I*. Springer, Berlin. pp. 98-154.
69. R. Selten, "Evolution, Learning, and Economic Behavior," *Games and Economic Behavior*, 3 (1991), 3-24.
70. R. Selten and R. Stoecker, "End Behavior in Sequences of Finite Prisoner's Dilemma Supergames: A Learning Theory Approach," *Journal of Economic Behavior and Organization*, 7 (1986), 47-70.
71. D. O. Stahl, "Boundedly Rational Rule Learning in a Guessing Game," *Games and Economic Behavior*, 16 (1996), 303-330.
72. D. O. Stahl, "Rule Learning in Symmetric Normal-form Games: Theory and Evidence," *Games and Economic Behavior*, 32 (2000), 105-138.

73. D. O. Stahl, "Population Rule Learning in Symmetric Normal-form Games: Theory and Evidence," *Journal of Economic Behavior and Organization*, 45 (2001), 19-35.
74. D. O. Stahl, and P. Wilson, "On Players Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior*, 10 (1995), 213-54.
75. F.-F. Tang, "Anticipatory learning in two-person games: some experimental results," *Journal of Economic Behavior and Organization*, 44 (2001), 221-232.
76. A. Tversky and D. Kahneman, "Loss Aversion in Riskless Choice: A Reference-Dependent Model," *The Quarterly Journal of Economics*, 106 (1991), 1039-1061.
77. J. Van Huyck, R. Battalio, and R. Beil, "Tacit Cooperation Games, Strategic Uncertainty, and Coordination Failure," *The American Economic Review*, 80 (1990), 234-248.
78. J. Van Huyck, R. Battalio, and F. W. Rankin. "Selection Dynamics and Adaptive Behavior Without Much Information," Texas A & M Department of Economics working paper, 1996.
79. J. Watson, "A 'Reputation' Refinement without Equilibrium," *Econometrica*, 61 (1993), 199-205.
80. J. Watson and P. Battigali, "On 'Reputation' Refinements with Heterogeneous Beliefs," *Econometrica*, 65 (1997), 363-374.
81. G. Weizsacker, "Ignoring the Rationality of Others: Evidence from Experimental Normal-form Games," Harvard Business School working paper, 2000.

Figure 1: The EWA Learning Cube




Figure 2a: Actual Choice Frequencies for Inexperienced Subjects




Figure 2b: Adaptive EWA Model Frequencies for Inexperienced Subjects




Figure 2c: Sophisticated EWA Model Frequencies for Inexperienced Subjects




Figure 3a: Actual Choice Frequencies for Experienced Subjects




Figure 3b: Adaptive EWA Model Frequencies for Experienced Subjects




Figure 3c: Sophisticated EWA Model Frequencies for Experienced Subjects




Figure 4: Borrower's optimal mixed-strategy of pay back as predicted by sequential equilibrium




Figure 5a: Empirical Frequency for No Loan




Figure 5c: Predicted Frequency for No Loan




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




Figure 5d: Predicted Frequency for Default conditional on Loan (Dishonest Borrower)




Table 1a: A summary of EWA parameter estimates and forecast accuracy (games estimated by us)

CITATION	GAME	EWA estimates (standard error)			Model accuracy					Comments
		δ	ϕ	$\rho=(1-\kappa)\phi$	EWA	Choice reinforcement – EWA	Belief – EWA	In / Out of sample	Fit technique	
Camerer, Ho and Hsia (2000)	Sealed bid mechanism+	n.a.	1.00	0.91	1102.0	30.8	65.5	IN	-LL	ψ, ω & κ replace δ & ρ
Camerer, Ho and Wang (1999)	“Continental divide” coordination	0.75	0.61	0.00	346.9	86.1	235.8	OUT	-LL	
Camerer and Ho (1998)	Weak-link coordination	0.65	0.58	0.20	358.1	29.1	438.6	IN	-LL	
Anderson and Camerer (in press)	Signaling games (game 3) 95% Confidence Interval	0.69 (0.47,1.00)	1.020 (0.99,1.04)	1.00 (0.98,1.00)	72.2	6.5	10.1	OUT	-LL	
	Signaling games (game 5) 95% Confidence Interval	0.54 (0.45,0.63)	0.65 (0.59,0.71)	0.46 (0.39, 0.54)	139.5	14.1	23.7	OUT	-LL	
Camerer and Ho (1999b)	Median-action coordination	0.85 (0.01)	0.80 (0.02)	0.00 (0.00)	41.1	39.2	72.8	OUT	-LL	
	4x4 Mixed-strategy games	0.00 (0.04)	1.04 (0.01)	0.96 (0.01)	326.4	9.1	-40.8	OUT	-LL	Payoff = 5 rupees
		0.73 (0.10)	1.01 (0.01)	0.95 (0.01)	341.7	18.0	8.4	OUT	-LL	Payoff =10 rupees
	6x6 Mixed-strategy games	0.41 (0.08)	0.99 (0.01)	0.94 (0.01)	301.7	6.8	-5.4	OUT	-LL	Payoff =5 rupees
		0.55 (0.05)	0.99 (0.01)	0.93 (0.02)	362.3	13.7	8.9	OUT	-LL	Payoff =10 rupees
	p-beauty contests*	0.95 (0.01)	0.11 (0.00)	0.00 (0.00)	1917.0	647.0	35.0	OUT	-LL	Experienced and Inexperienced Combined
Camerer, Ho and Wang (1999)	Normal form centipede (odd player)	0.32 (0.32)	0.91 (0.14)	0.00 (0.00)	1016.8	57.6	536.3	OUT	-LL	Clairvoyance full update, κ
	Normal form centipede (even player)	0.24 (0.32)	0.90 (0.14)	0.95 (0.03)	951.3	46.4	604.7	OUT	-LL	Clairvoyance full update, κ

+ In Figure 1, we did not include this study.

* Unlike the previous estimates, these new estimates assume that subjects do not know the winning numbers.

Table 1b: A summary of EWA parameter estimates and forecast accuracy (games estimated by others)

CITATION	GAME	EWA estimates (standard error)			Model accuracy					Comments
		δ	ϕ	$\rho=(1-\kappa)\phi$	EWA	Choice reinforce – EWA	Belief – EWA	In / Out of sample	Fit technique	
Chen and Khoroshilov (1999)	Cost allocation+	.80~1.0	1 (fixed)	.1~.3	.73~.88	-.01~.07	n.a.	IN	MSD	
Morgan and Sefton (1999)	“Unprofitable” games (baseline games)	0.08 (0.07)	0.93 (0.01)	0.92 (0.03)	1729.5	0.6	n.a.	IN	-LL	
	“Unprofitable” games (upside games)	0.14 (0.06)	0.89 (0.01)	0.00 (0.00)	1906.5	16.2	n.a.	IN	-LL	
Stahl (1999)	5x5 matrix games	.66 (0.02)	.34 (0.04)	.086 (0.08)	4803.7	64.7	n.a.	OUT	-LL	
Hsia (1999)	Call markets	0.47 (0.32)	0.97 (0.01)	0.74 (0.06)	1915.0	0.0	403.0	IN	-LL	
Amaldoss (1998)	Same function alliance – equal profit sharing	0.00	0.95	0.89	886.3	1.6	529.6	IN	-LL	High reward
		0.00	0.99	0.99	767.5	30.1	390.4	IN	-LL	Med. Reward
		0.07	0.89	0.87	1399.7	9.4	541.5	IN	-LL	Low reward
	Same function alliance – proportional sharing	0.00	0.99	1.00	910.8	36.4	813.0	IN	-LL	High reward
		0.00	0.92	0.96	1055.0	18.3	615.9	IN	-LL	Med. Reward
		0.00	0.97	0.96	1013.7	13.3	1095.6	IN	-LL	Low reward
	Parallel development of product – equal sharing	0.00	0.91	0.59	1194.2	0.1	566.3	IN	-LL	High reward
		0.17	0.90	0.93	1321.5	9.5	497.2	IN	-LL	Med. Reward
		0.21	0.88	0.66	1297.7	4.5	484.1	IN	-LL	Low reward
Rapoport and Amaldoss (2000)	Patent race game – symmetric players	0.00	0.94	0.93	3551.7	12.1	1097.7	IN	-LL	Low reward
		0.00	0.97	0.98	2908.1	20.2	725.9	IN	-LL	High reward
	Patent race game – asymmetric players	0.48	0.90	0.86	3031.5	89.1	706.8	IN	-LL	Strong player
		0.14	0.96	0.97	2835.5	15.7	611.0	IN	-LL	Weak player

+ In Figure 1, we did not include this study.

Table 2. Model Parameter Estimates for p -beauty Contest Game

	INEXPERIENCED SUBJECTS			EXPERIENCED SUBJECTS		
	Sophisticated EWA	Adaptive EWA	QRE ¹	Sophisticated EWA	Adaptive EWA	QRE
ϕ	0.44 (0.05) ²	0.00 (0.00)	- -	0.29 (0.03)	0.22 (0.03)	- -
δ	0.78 (0.08)	0.90 (0.05)	- -	0.67 (0.05)	0.99 (0.02)	- -
ρ	0.00 (0.00)	0.00 (0.00)	- -	0.01 (0.00)	0.00 (0.00)	- -
α	0.24 (0.04)	0.00 (0.00)	1.00 (0.00)	0.77 (0.02)	0.00 (0.00)	1.00 (0.00)
α'	0.00 (0.00)	0.00 (0.00)	- -	0.41 (0.03)	0.00 (0.00)	- -
d	0.16 (0.02)	0.13 (0.01)	0.04 (0.01)	0.15 (0.01)	0.11 (0.01)	0.04 (0.00)
LL						
(in sample)	-2095.32	-2155.09	-2471.50	-1908.48	-2128.88	- 2141.45
(out of sample)	-968.24	-992.47	-1129.25	-710.28	-925.09	- 851.31
Avg. Prob.						
(in sample)	6%	5%	3%	7%	5%	5%
(out of sample)	7%	7%	4%	13%	9%	9%

¹We also estimated the QRE model with different λ s in each period. For inexperienced players, the λ s in period 2 to 6 are: 0.590; 0.663; 0.941; 1.220; 1.221; 1.381. In sample $LL=-2460.51$; out of sample $LL=-1100.09$. For experienced players, the λ 's are 1.382; 2.627; 3.970; 5.249; 5.363; 8.399. In sample $LL=-2074.25$; out of sample $LL=-769.19$.

²Standard errors in parenthesis.

Table 3: Payoffs in the borrower-lender 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

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

Table 4: A Comparison of In-Sample and Out-of-Sample Performance Between the Teaching and AQRE Models

Experiment No.:		3	4	5	6	7	8	9	10
No of Sequence (Total):		90	90	81	70	77	69	90	101
<u>In-Sample Calibration</u>									
Sample size	Total	744	784	742	661	673	626	703	824
	Borrower (Dishonest)	177	233	238	267	238	240	223	288
	Borrower (Honest)	87	71	72	26	27	18	0	0
	Lender	480	480	432	368	408	368	480	536
<u>The Teaching Model</u>									
Average Probability	Total	80%	82%	83%	88%	81%	83%	71%	84%
	Borrower (Dishonest)	60%	72%	80%	86%	73%	73%	53%	81%
	Borrower (Honest)	86%	86%	81%	86%	83%	86%	-	-
	Lender	87%	86%	84%	90%	86%	90%	79%	85%
Log-likelihood	Total	-282.15	-321.84	-294.50	-182.38	-285.73	-204.52	-433.73	-330.13
	Borrower (Dishonest)	-62.75	-88.41	-67.74	-57.50	-97.25	-86.53	-145.53	-88.03
	Borrower (Honest)	-13.31	-10.86	-20.01	-3.98	-5.95	-2.75	-	-
	Lender	-206.09	-222.56	-206.75	-120.90	-182.53	-115.24	-288.19	-242.11
<u>Agent-based Quantal Response Equilibrium (AQRE)</u>									
Average Probability	Total	79%	78%	78%	81%	76%	77%	69%	77%
	Borrower (Dishonest)	73%	69%	75%	76%	67%	66%	52%	72%
	Borrower (Honest)	86%	86%	81%	86%	83%	86%	-	-
	Lender	80%	81%	80%	84%	80%	85%	76%	80%
Log-likelihood	Total	-370.37	-433.48	-374.90	-293.89	-371.63	-300.69	-475.81	-455.30
	Borrower (Dishonest)	-79.28	-122.32	-99.21	-103.85	-130.09	-125.40	-153.36	-112.64
	Borrower (Honest)	-13.32	-10.88	-20.01	-3.98	-5.93	-2.75	-	-
	Lender	-277.77	-300.29	-255.68	-186.07	-235.61	-172.54	-322.45	-342.66
<u>Out-Of-Sample Validation</u>									
Sample size	Total	401	386	368	356	315	288	419	361
	Borrower (Dishonest)	108	99	121	157	99	99	179	89
	Borrower (Honest)	53	47	31	7	8	5	0	0
	Lender	240	240	216	192	208	184	240	272
<u>The Teaching Model</u>									
Average Probability	Total	80%	82%	84%	89%	81%	85%	71%	86%
	Borrower (Dishonest)	65%	75%	81%	86%	76%	77%	57%	78%
	Borrower (Honest)	86%	86%	86%	86%	77%	86%	-	-
	Lender	86%	85%	85%	92%	84%	89%	81%	88%
Log-likelihood	Total	-147.47	-148.45	-145.66	-70.62	-140.99	-99.44	-215.17	-139.13
	Borrower (Dishonest)	-24.90	-26.89	-42.00	-26.40	-32.73	-30.16	-91.14	-52.59
	Borrower (Honest)	-8.11	-7.19	-4.74	-1.07	-3.02	-0.77	0.00	0.00
	Lender	-114.46	-114.37	-98.92	-43.15	-105.23	-68.52	-124.03	-86.54
<u>Agent-based Quantal Response Equilibrium (AQRE)</u>									
Average Probability	Total	78%	80%	80%	82%	76%	77%	68%	79%
	Borrower (Dishonest)	77%	78%	79%	78%	73%	71%	60%	71%
	Borrower (Honest)	86%	86%	86%	86%	77%	86%	-	-
	Lender	77%	80%	80%	85%	77%	81%	75%	81%
Log-likelihood	Total	-199.77	-187.43	-176.23	-137.60	-189.84	-163.08	-270.71	-195.80
	Borrower (Dishonest)	-36.87	-30.99	-40.13	-54.62	-41.03	-40.72	-101.02	-45.86
	Borrower (Honest)	-8.12	-7.20	-4.74	-1.07	-3.02	-0.77	-	-
	Lender	-154.78	-149.24	-131.35	-81.91	-145.79	-121.59	-169.70	-149.95

Note 1: Data source from Camerer and Weigelt (1988)

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: Average probability is computed by taking arithmetic mean of the predicted probability of chosen strategies.

Table 5: Parameter Estimates of the Teaching Model

Experiment No.:	3	4	5	6	7	8	9	10
Proportion of Induced Honest Type, p :	0.33	0.33	0.33	0.10	0.10	0.10	0.00	0.00
<u>Parameters for Adaptive Lender</u>								
ϕ	0.29 (0.01)	0.65 (0.01)	0.67 (0.01)	0.73 (0.01)	0.57 (0.01)	0.69 (0.02)	0.74 (0.01)	0.09 (0.00)
δ	0.34 (0.01)	0.21 (0.01)	0.44 (0.01)	0.80 (0.01)	0.19 (0.01)	0.65 (0.01)	0.42 (0.01)	0.47 (0.01)
κ	0.99 (0.00)	0.95 (0.00)	0.35 (0.01)	0.40 (0.02)	1.00 (0.00)	0.12 (0.01)	0.00 (0.00)	0.88 (0.00)
τ	0.51 (0.01)	0.99 (0.00)	1.00 (0.00)	0.95 (0.00)	0.99 (0.00)	0.97 (0.00)	1.00 (0.00)	0.98 (0.00)
<u>Parameters for Adaptive Borrower (Dishonest)</u>								
ϕ	1.00 (0.00)	0.60 (0.01)	0.81 (0.01)	0.48 (0.01)	1.00 (0.00)	0.69 (0.01)	0.09 (0.00)	0.87 (0.00)
δ	0.17 (0.00)	0.48 (0.01)	1.00 (0.00)	0.49 (0.01)	0.53 (0.02)	0.55 (0.01)	0.05 (0.00)	0.47 (0.01)
κ	0.90 (0.00)	1.00 (0.00)	1.00 (0.00)	0.76 (0.01)	0.73 (0.01)	0.83 (0.01)	0.92 (0.00)	0.59 (0.01)
τ	0.10 (0.00)	0.24 (0.01)	1.00 (0.00)	0.01 (0.00)	0.76 (0.01)	0.35 (0.01)	0.21 (0.01)	0.67 (0.01)
<u>Parameters for Teaching Borrower (Dishonest)</u>								
ε	1.00 (0.00)	1.00 (0.00)	0.95 (0.00)	0.93 (0.00)	0.12 (0.00)	0.93 (0.00)	0.15 (0.00)	0.24 (0.01)
α	0.70 (0.01)	0.34 (0.01)	0.18 (0.01)	0.02 (0.00)	0.13 (0.01)	0.28 (0.01)	0.68 (0.01)	0.02 (0.00)

Note: Each parameter is presented with its standard error (in parenthesis) directly below.