



Experienced-Weighted Attraction Learning in Normal Form Games

Colin Camerer; Teck-Hua Ho

Econometrica, Vol. 67, No. 4. (Jul., 1999), pp. 827-874.

Stable URL:

<http://links.jstor.org/sici?sici=0012-9682%28199907%2967%3A4%3C827%3AEALINF%3E2.0.CO%3B2-B>

Econometrica is currently published by The Econometric Society.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/econosoc.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

EXPERIENCE-WEIGHTED ATTRACTION LEARNING IN NORMAL FORM GAMES

BY COLIN CAMERER AND TECK-HUA HO¹

In ‘experience-weighted attraction’ (EWA) learning, strategies have attractions that reflect initial predispositions, are updated based on payoff experience, and determine choice probabilities according to some rule (e.g., logit). A key feature is a parameter δ that weights the strength of hypothetical reinforcement of strategies that were *not* chosen according to the payoff they would have yielded, relative to reinforcement of chosen strategies according to received payoffs. The other key features are two discount rates, ϕ and ρ , which separately discount previous attractions, and an experience weight. EWA includes reinforcement learning and weighted fictitious play (belief learning) as special cases, and hybridizes their key elements. When $\delta = 0$ and $\rho = 0$, cumulative choice reinforcement results. When $\delta = 1$ and $\rho = \phi$, levels of reinforcement of strategies are exactly the same as expected payoffs given weighted fictitious play beliefs. Using three sets of experimental data, parameter estimates of the model were calibrated on part of the data and used to predict a holdout sample. Estimates of δ are generally around .50, ϕ around .8–1, and ρ varies from 0 to ϕ . Reinforcement and belief-learning special cases are generally rejected in favor of EWA, though belief models do better in some constant-sum games. EWA is able to combine the best features of previous approaches, allowing attractions to begin and grow flexibly as choice reinforcement does, but reinforcing unchosen strategies substantially as belief-based models implicitly do.

KEYWORDS: Learning, behavioral game theory, reinforcement learning, fictitious play.

1. INTRODUCTION

HOW DOES AN EQUILIBRIUM arise in a noncooperative game? While it is conceivable that players reason their way to an equilibrium, a more psychologi-

¹This research was supported by NSF Grants SBR-9511001, 9511137, 9601236, and 9730187, and the hospitality of the Center for Advanced Study in Behavioral Sciences. We have had helpful discussions with Bruno Broseta, Vince Crawford, Ido Erev, Dan Friedman, Dave Grether, Elef Gkioulekas, John Kagel, Tom Palfrey, Matthew Rabin, Al Roth, Yuval Rottenstreich, Rajiv Sarin, Rakesh Sarin, Dale Stahl, John Van Huyck, Roberto Weber, and an editor, and research assistance from Hongjai Rhee, Chris Anderson, and Juin Kuan Chong. Barry Sopher generously provided data. Many helpful comments were received from anonymous referees and seminar participants at the Society for Mathematical Psychology conference in honor of Amnon Rapoport (July 1996), the Russell Sage Foundation Summer Institute in Behavioral Economics (July 1996), the Economic Science Association (October 1996), the Marketing Science Conference (March 1997), Bonn Conference on Theories of Bounded Rationality (May 1997), FUR VIII in Mons, Belgium (July 1997), Gerzensee ESSET Economic Theory Conference (July 1997), BDRM (June 1998), and seminars at Caltech, Harvard, and Washington Universities, the Stockholm School of Economics, and the Universities of Alicante, Autonoma, California (Berkeley, Los Angeles), Chicago, Pennsylvania, Pittsburgh, Pompeu Fabra, Texas (Austin), and Texas A&M.

cally plausible view is that players adapt or evolve toward it.² The flurry of recent research on adaptation and evolution mostly explores theoretical questions, such as to which types of equilibria specific evolutionary or adaptive rules converge. We are interested in a fundamentally empirical question: Which models describe human behavior best? In this paper we propose a general ‘experience-weighted attraction’ (EWA) model and estimate the model parametrically, using three sets of experimental data.

The EWA model combines elements of two seemingly different approaches, and includes them as special cases. One approach, belief-based models, starts with the premise that players keep track of the history of previous play by other players and form some belief about what others will do in the future based on past observation. Then they tend to choose a best-response, a strategy that maximizes their expected payoffs given the beliefs they formed.

A different approach, choice reinforcement, assumes that strategies are ‘reinforced’ by their previous payoffs, and the propensity to choose a strategy depends in some way on its stock of reinforcement. Players who learn by reinforcement do not generally have beliefs about what other players will do. They care only about the payoffs strategies yielded in the past, not about the history of play that created those payoffs.

The belief and reinforcement approaches have been treated as fundamentally different since the 1950s. Until recently, nobody asked whether the two might be related, or how. But like two rivers with a surprising common source, or children raised apart who turn out to be siblings, belief and reinforcement are special kinds of one learning model. The common heritage of these approaches was probably not discovered earlier because reinforcement models were used primarily by psychologists, and belief models primarily by decision and game theorists. In addition, the information used by each approach is quite different. Belief-based models do not specially reflect past successes (reinforcements) of chosen strategies. Reinforcement models do not reflect the history of how others played. The EWA approach includes both as special cases by incorporating both kinds of information, using three modelling features.

The crucial feature is how strategies are reinforced. In the choice reinforcement approach, when player 1 picks strategy s_1^j , and player 2 picks s_2^k , player 1’s strategy s_1^j is reinforced according to the payoff $\pi_1(s_1^j, s_2^k)$. Unchosen strategies $s_1^h (h \neq j)$ are not reinforced at all. In EWA, the unchosen strategies *are* reinforced based on a multiple δ of the payoffs $\pi_1(s_1^h, s_2^k)$ they would have earned. This makes psychological sense because research on human and animal learning shows that people learn from many kinds of experiences other than those that are directly reinforcing. An expanded notion of reinforcement therefore liberates learning from the straitjacket of behaviorist psychology, toward something more cognitive and descriptive of humans.

²Like most good ideas in economics, the adaptive and evolutionary interpretations of equilibration have a long pedigree. Weibull (1997) pointed out that in Adam Smith’s famous passage where he said that the division of labor emerged as a consequence of the “propensity to truck, barter, and exchange,” Smith also noted that the division of labor emerged in a “very slow and gradual” way (1981).

The second EWA feature controls the growth rates of attractions. Attractions are numbers that are monotonically related to the probability of choosing a strategy. In cumulative reinforcement models attractions can grow and grow, which implies that convergence can be sharper (in the sense that choice probabilities diverge toward one and zero). In belief learning, attractions are expected payoffs, which are always bounded by the range of matrix payoffs. The EWA model allows growth rates to vary between these two bounds by using separate decay rates, ϕ for past attractions, and ρ for the amount of experience (which normalizes attractions).

The third modelling feature is initial attraction and experience weight. In belief models initial attractions must be expected payoffs given prior beliefs. In reinforcement models initial attractions are usually unrestricted. Therefore, initial attractions are unrestricted in EWA too. The initial experience weight $N(0)$ reflects a strength of prior in belief models, or the relative weight given to lagged attractions versus payoffs when attractions are updated.

When $\delta = 0$, $\rho = 0$, and $N(0) = 1$, the EWA attractions of strategies are equal to reinforcements, as used in many models. When $\delta = 1$, and $\phi = \rho$ (and initial attractions are determined by prior beliefs), the attractions of strategies are equal to their expected payoffs given beliefs in a general class. That is, reinforcing each strategy according to what it would have earned (or did earn) is behaviorally equivalent to forming beliefs, based on observed history, and calculating expected payoffs. The equivalence holds because looking back at what strategies earned (or would have) in the past is the same as forming beliefs based on what others did in the past, then computing forward-looking expected payoffs based on those backward-looking beliefs.

EWA tries to mix appropriate elements of reinforcement and belief learning approaches in a way that makes sense. We think this can be judged by whether the parameters have clear psychological interpretations, and whether adding them improves statistical fit (adjusting, of course, for added degrees of freedom) and predictive accuracy. To test the empirical usefulness of EWA, we derived maximum-likelihood parameter estimates from three data sets. The data sets span a wide range of games: constant-sum games with unique mixed-strategy equilibria; coordination games with multiple Pareto-ranked equilibria; and ‘*p*-beauty contests’ with unique dominance-solvable equilibria. Some empirical studies have evaluated belief and reinforcement models, but most have not compared them directly with statistical tests. Because EWA is a generalization which reduces to belief and reinforcement learning when parameters have certain values, it is easy to compare them to EWA and to each other.

In the next section, the EWA approach is defined and we show how a general class of choice reinforcement and adaptive belief-based approaches are special cases. The third section provides interpretations of the model parameters and discusses how they relate to principles of human learning. The fourth section describes previous findings and shows how our empirical implementation goes further than earlier work. The fifth section reports parameter estimates from several data sets. The last section concludes and mentions some future research directions.

2. THE EXPERIENCE-WEIGHTED ATTRACTION (EWA) MODEL

We start with notation. We study n -person normal-form games. Players are indexed by i ($i = 1, \dots, n$), and the strategy space of player i , S_i consists of m_i discrete choices, that is, $S_i = \{s_i^1, s_i^2, \dots, s_i^j, \dots, s_i^{m_i-1}, s_i^{m_i}\}$. $S = S_1 \times \dots \times S_n$ 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 $m_{-i} = \prod_{k=1, k \neq i}^n m_k$. 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, which determines the probability of choosing that strategy (in a precise way made clear below). Learning models require a specification of initial attractions, how attractions are updated by experience, and how choice probabilities depend on attractions.

2.1. The EWA Updating Rules

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. The second variable is $A_i^j(t)$, player i 's attraction of strategy s_i^j after period t has taken place.

The variables $N(t)$ and $A_i^j(t)$ begin with some prior values, $N(0)$ and $A_i^j(0)$. These prior values can be thought of as reflecting pregame experience, either due to learning transferred from different games or due to introspection. (Then $N(0)$ can be interpreted as the number of periods of actual experience, which is equivalent in attraction impact to the pregame thinking.)

Updating is governed by two rules. First,

$$(2.1) \quad N(t) = \rho \cdot N(t-1) + 1, \quad t \geq 1.$$

The parameter ρ is a depreciation rate or retrospective discount factor that measures the fractional impact of previous experience, compared to one new period.

The second rule updates the level of attraction. A key component of the updating is 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 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 $I(x, y)$ that 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 rule for updating attraction sets $A_i^j(t)$ to be the sum of a depreciated, experience-weighted previous attraction $A_i^j(t-1)$ plus the (weighted) payoff

from period t , normalized by the updated experience weight:

$$(2.2) \quad A_i^j(t) = \frac{\phi \cdot N(t-1) \cdot A_i^j(t-1) + [\delta + (1-\delta) \cdot I(s_i^j, s_i(t))] \cdot \pi_i(s_i^j, s_{-i}(t))}{N(t)}.$$

The factor ϕ is a discount factor or decay rate, which depreciates previous attraction.

2.2. Choice Reinforcement

In early reinforcement models (and some recent ones) choice probabilities are updated directly (e.g., Bush and Mosteller (1955); cf. Cross (1983)). In more recent models (Harley (1981); Roth and Erev (1995)), strategies have levels of reinforcement or propensity that are incremented cumulatively by received payoffs (and perhaps normalized; Arthur (1991)). We emphasize the latter cumulative form, which gives more modelling freedom³ and avoids some clumsy technical features (e.g., imposing boundary conditions so probabilities do not grow above one or below zero).

The initial reinforcement level of strategy s_i^j of player i is $R_i^j(0)$. These initial reinforcements can be assumed a priori (based on a theory of first-period play) or estimated from the data. Reinforcements are updated according to two principles:

$$(2.3) \quad R_i^j(t) = \begin{cases} \phi \cdot R_i^j(t-1) + \pi_i(s_i^j, s_{-i}(t)) & \text{if } s_i^j = s_i(t), \\ \phi \cdot R_i^j(t-1) & \text{if } s_i^j \neq s_i(t). \end{cases}$$

The two principles can be reduced to a single updating equation:

$$(2.4) \quad R_i^j(t) = \phi \cdot R_i^j(t-1) + I(s_i^j, s_i(t)) \cdot \pi_i(s_i^j, s_{-i}(t)).$$

It is easy to see that this updating formula is a special case of the EWA rule, when $\delta = 0$, $N(0) = 1$, and $\rho = 0$. Thus, cumulative choice reinforcement in this form is a special case of experience-weighted attraction learning.⁴

Other reinforcement models assume that previous payoffs are averaged, rather than cumulated (McAllister (1991), Mookerjee and Sopher (1994, 1997),

³In the Cross model, strategies have utilities that are weighted averages of past utilities and current payoffs (for chosen strategies), and players maximize utility. Sarin (1995) shows that when the weight on current payoff declines over time, this model behaves similarly to the Harley version in which attractions grow. The similarity reflects the fact that both models build in a declining effect of marginal reinforcements.

⁴Some reinforcement models add other parameters. Roth and Erev (1995) add a parameter that cuts off attractions close to zero, to avoid negative attractions. Erev and Roth (1997) add three parameters that allow reinforcement to depend on payoffs minus an (updated) reference point (as in Bush and Mosteller (1995), Cross (1983)), where the updating may be different for losses and gains. They also add a parameter that smears a portion of the chosen-strategy reinforcement to neighboring strategies, to reflect a kind of experimentation or generalization that is (locally) similar to our δ parameter. Camerer and Ho (1998) compare the local-generalization specification with δ updating in the EWA model and find that local-generalization fits much worse.

Sarin and Vahid (1997)). Then reinforcements are updated according to

$$(2.5) \quad R_i^j(t) = \phi \cdot R_i^j(t-1) + (1-\phi)I(s_i^j, s_i(t)) \cdot \pi_i(s_i^j, s_{-i}(t)).$$

When $\delta = 0$, $N(0) = 1/(1-\rho)$ and $\rho = \phi$, EWA reduces to this averaged-reinforcement form.

2.3. Belief-based Models

In a belief-based model, players tend to choose strategies that have high expected payoffs given beliefs formed by observing the history of what others did. While there are many ways of forming beliefs, we consider a fairly large class of weighted fictitious play models, which include familiar ones like fictitious play (Brown (1951)) and Cournot (1960) best-response as special cases (see Fudenberg and Levine (1995, 1998), Cheung and Friedman (1997)).⁵

In the weighted fictitious play model, prior beliefs of opponents' strategy combinations are expressed as a ratio of hypothetical counts of observations of strategy combination s_{-i}^k , denoted by $N_{-i}^k(0)$. These observations can then be naturally integrated with actual observations as experience accumulates. (Carnap (1962) shows an elegant set of axioms that implies this structure, which corresponds to Bayesian updating with a Dirichlet-distributed prior.) In our view, specifying prior beliefs (and computing initial expected payoffs based on the prior) is a crucial feature of belief models, though some papers have not imposed this assumption. Without specifying a prior, there is no guarantee that the updated beliefs that result from mixing initial expected payoffs with later experience will be valid beliefs (i.e., nonnegative probabilities that sum to one).

We also allow past experience to be depreciated or discounted by a factor ρ (presumably between zero and one). Formally, the prior beliefs for player i about choices of others are specified by a vector of relative frequencies of choices of strategies s_{-i}^k , denoted $N_{-i}^k(0)$. Call the sum of those frequencies (dropping the player subscript for simplicity) $N(t) = \sum_{k=1}^{m-i} N_{-i}^k(t)$. Then the initial prior $B_{-i}^k(0)$ is

$$(2.6) \quad B_{-i}^k(0) = \frac{N_{-i}^k(0)}{N(0)},$$

with $N_{-i}^k(0) \geq 0$ and $N(0) > 0$. Beliefs are updated by depreciating the previous counts by ρ , and adding one for the strategy combination actually chosen by the other players. That is,

$$(2.7) \quad B_{-i}^k(t) = \frac{\rho \cdot N_{-i}^k(t-1) + I(s_{-i}^k, s_{-i}(t))}{\sum_{h=1}^{m-i} [\rho \cdot N_{-i}^h(t-1) + I(s_{-i}^h, s_{-i}(t))]}.$$

⁵When the description 'fictitious play' is used below, we mean traditional fictitious play in which all past observations are weighted equally. Also, Crawford (1995) and Camerer and Ho (1998) estimate models in which ϕ varies across periods, which generalizes weighted fictitious play to include cases where the weight rises or falls over time. In both papers, allowing time-varying weight does not improve fit very much, so assuming a fixed ϕ seems reasonable.

Expressing beliefs in terms of previous-period beliefs,

$$(2.8) \quad B_{-i}^k(t) = \frac{\rho \cdot B_{-i}^k(t-1) + \frac{I(s_{-i}^k, s_{-i}(t))}{N(t-1)}}{\rho + \frac{1}{N(t-1)}} \\ = \frac{\rho \cdot N(t-1) \cdot B_{-i}^k(t-1) + I(s_{-i}^k, s_{-i}(t))}{\rho \cdot N(t-1) + 1}.$$

This form of belief updating weights observations from one period ago ρ times as much as the most recent observation. This includes Cournot dynamics ($\rho = 0$; only the most recent observation counts) and fictitious play ($\rho = 1$; all observations count equally) as special cases. The general case $0 \leq \rho \leq 1$ is a compromise in which all observations count but more recent observations count more.

Expected payoffs in period t , $E_i^j(t)$, are taken over beliefs according to

$$(2.9) \quad E_i^j(t) = \sum_{k=1}^{m_{-i}} \pi_i(s_i^j, s_{-i}^k) \cdot B_{-i}^k(t).$$

The crucial step is to express period t expected payoffs as a function of period $t-1$ expected payoffs. Substituting equation (2.8) into (2.9) and rearranging yields:

$$(2.10) \quad E_i^j(t) = \frac{\rho \cdot N(t-1) \cdot E_i^j(t-1) + \pi(s_i^j, s_{-i}(t))}{\rho \cdot N(t-1) + 1}.$$

This equation makes the kinship between the EWA and belief approaches transparent. Formally, suppose initial attractions are equal to expected payoffs given initial beliefs that arise from the ‘experience-equivalent’ strategy counts $N_{-i}^k(0)$, so $A_i^j(0) = E_i^j(0) = \sum_{k=1}^{m_{-i}} \pi_i(s_i^j, s_{-i}^k) \cdot B_{-i}^k(0)$. Then substituting $\delta = 1$ and $\rho = \phi$ into the attraction updating equation (2.2) gives attractions that are exactly the same as updated expected payoffs in (2.10). Hence, the weighted belief models are a special case of EWA.

The contrast with EWA makes clear that belief models actually make three separate assumptions: Players’ initial attractions are expected payoffs based on some prior; players update attractions using EWA with $\delta = 1$; and attractions are a weighted average of lagged attractions and payoffs ($\phi = \rho$). We think the most intuitively appealing assumption is the best-responsiveness to foregone payoffs embodied in $\delta = 1$, rather than the weighted-average restriction $\phi = \rho$ or the restriction on first-period play. EWA allows one to separate the three features of belief learning: Players could have attractions that begin and grow differently than belief models assume, but update those attractions in a belief-learning way. Such players are a special kind of EWA learner.

The close relation between reinforcement and belief learning is surprising because the two approaches have generally been treated as fundamentally different (e.g., Selten (1991, p. 14)). Some authors have extended choice reinforcement models to include reinforcement using all foregone payoffs (McAllister (1991)) or the highest foregone payoffs (Roth (1995, pp. 37–40), Roth and Erev (1995), on market games), without noticing that these extensions make reinforcement like belief learning.

Some connection between reinforcement and belief learning was recognized very recently by others (unbeknownst to us). Fudenberg and Levine (1995, pp. 1084–1085) and Cheung and Friedman (1997, p. 54–55) both pointed out that expected payoffs computed using fictitious play beliefs, and based on history, are asymptotically the same as histories of actual payoffs. But their arguments are based on long-run asymptotic equivalence between a stationary distribution (possible payoffs) and a sample from it (actual payoffs). Neither explicitly recognized that even in the short run, there is an exact equivalence between a general kind of reinforcement learning (EWA) and weighted fictitious play.⁶

The nonlinear interplay of parameters in the EWA updating rules is why, as a model of human learning, EWA is potentially superior to simply running a regression of choices against reinforcements and expected payoffs or combining the two in a weighted average. Reinforcements and expected payoffs differ in three crucial dimensions—initial attractions and experience weight $N(0)$, the weight δ on foregone payoffs in updating attractions, and whether attractions can grow outside the bounds of possible payoffs (which depends on ϕ and ρ). EWA is not a convex combination of reinforcement and belief models because these three dimensions are controlled by separate parameters. For instance, a weighted average in which expected payoffs are given weight δ and reinforcements have weight $1 - \delta$ will update attractions like EWA does, but that weighted average will not allow the wide range of initial attractions, experience rates, and growth rates available in EWA.⁷

2.4. Choice Probabilities

Attractions must determine probabilities of choosing strategies in some way. $P_i^j(t)$ should be monotonically increasing in $A_i^j(t)$ and decreasing in $A_i^k(t)$ (where $k \neq j$). Three forms have been used in previous research: Exponential (logit), power, and normal (probit). In estimation reported below we use the

⁶For example, Cheung and Friedman (1997) make their point by “assum[ing] for the moment (very counterfactually!), that the player somehow managed to play both strategies each period.” Then “dropping the counterfactual,” they show that the average experienced payoffs will correspond, up to some noise, to expected payoffs. Counterfactual simulation of foregone payoffs is precisely the mental process invoked by δ in EWA. However, the ‘noise’ is correlated with past observations that are included explicitly in EWA, so the relation between EWA and weighted fictitious play is exact rather than approximate.

⁷Indeed, Camerer and Ho (1998) show that EWA fits much better than a convex combination of belief and reinforcement learning, in two coordination games.

logit function, which is commonly used in studies of choice under risk and uncertainty, brand choice, etc. (Ben-Akiva and Lerman (1985), Anderson, Palma, and Thisse (1992)), and is given by

$$(2.11) \quad P_i^j(t+1) = \frac{e^{\lambda \cdot A_i^j(t)}}{\sum_{k=1}^{m_i} e^{\lambda \cdot A_i^k(t)}}.$$

The parameter λ measures sensitivity of players to attractions. Sensitivity could vary due to the psychophysics of perception or whether subjects are highly motivated or not. In this probability function, the exponent in the numerator is just the weighted effect of strategy s_i^j 's attraction, $\lambda \cdot A_i^j(t)$, on the probability of choosing strategy s_i^j . Models in which cross-effects of attractions on other strategies' choice probabilities are allowed have been estimated (Mookerjee and Sopher (1997)) but we do not have the degrees of freedom to do so.⁸

The logit, power, and probit probability functions each have advantages and disadvantages. The exponential form has been used to study learning in games by Mookerjee and Sopher (1994, 1997), Ho and Weigelt (1996), and Fudenberg and Levine (1998), and in 'quantal response equilibrium' models by Chen, Friedman, and Thisse (in press) and McKelvey and Palfrey (1995, 1998). Cheung and Friedman (1997) used the probit form. The exponential form is invariant to adding a constant to all attractions.⁹ As a result, negative values of $A_i^j(0)$ are permissible, which means one can avoid the difficult question of how to update attractions when payoffs are negative.¹⁰

The power probability form is given by

$$(2.12) \quad P_i^j(t+1) = \frac{(A_i^j(t))^\lambda}{\sum_{k=1}^{m_i} (A_i^k(t))^\lambda}.$$

⁸In Mookerjee and Sopher (1997), the exponent in the probability equation numerator is the sum of weighted effects of all the attractions, $\sum_{k=1}^{m_i} \lambda_{jk} \cdot A_i^k(t)$, where λ_{jk} is the cross-effect of strategy s_i^k 's attraction on strategy s_i^j 's score. This model allows cross-effects in which one strategy's attraction can affect other strategies' choice probabilities differently. These cross-effects are hard to interpret without knowing more about similarity of strategies or some other basis for one strategy's attraction to affect others differently. Nonetheless, they have some significance as a whole in the Mookerjee-Sopher analysis of constant-sum games. Estimating them for our median-action and p -beauty contest data uses up far too many degrees of freedom because there are too many strategies. Including cross-effects could proceed particularly efficiently if some structural considerations were used to restrict coefficients a priori (as in Sarin and Vahid's (1997) use of strategy similarity).

⁹As a result, one must normalize $A_i^j(0)$ to equal a constant for one value of j in order to identify parameters. There is some evidence that adding a constant to payoffs does matter (Bereby-Meyer and Erev (1997)) but there is also evidence that logit fits better than power, so we regard the choice of proper form as a matter of one's purpose and yet-unresolved empirical debate.

¹⁰Borgers and Sarin (1996) avoid this problem by adding x to all other strategies when a chosen strategy loses x .

The power form is invariant to multiplying all attractions by a constant. Because of this invariance, the parameters $N(0)$ and ρ make no difference when the power form is used (i.e., they are not identified).¹¹

Depending on one's purpose, being able to ignore $N(0)$ and ρ can be an advantage or disadvantage. For the purpose of distinguishing different models, it is a big disadvantage because models impose different restrictions on $N(0)$ and ρ . By using the power form, the difference between belief-based, reinforcement, and EWA models, besides initial attractions, is only one parameter, δ , rather than three parameters. For the purposes of estimating any one model reliably, however, conserving degrees of freedom is good so the power form is better. Since our main purpose in this paper is comparing models, having the extra tools to distinguish theories is a large advantage so we use the logit form rather than the power form. This choice of probability rule is, of course, not an essential part of the EWA model.

Ultimately, it is an empirical question whether the logit, probit, or power forms fit better (adjusting for degrees of freedom). Previous studies show roughly equal fits of logit and power (Tang (1996), Chen and Tang (1998), Erev and Roth (1997)) or better fits for the logit form over the power form (Camerer and Ho (1998)).

3. INTERPRETING EWA PARAMETERS

We think it is crucial to ask how a learning model's parameters can be interpreted, what general behavioral principles of learning they capture, and, for EWA, how they reveal the assumptions implicit in reinforcement and belief learning. Asking these questions about any learning theory avoids the danger of adding parameters just to improve statistical fit, without adding new insight or respecting what is known in other disciplines. In addition, if parameters have natural psychological interpretations they can be measured in other ways (e.g., response times and attention measures) and used in psychological modelling.

3.1. *Learning Principles, Aspiration Levels, and δ*

The parameter δ measures the relative weight given to foregone payoffs, compared to actual payoffs, in updating attractions. This is the most important parameter in EWA because it shows most clearly the different ways in which EWA, reinforcement and belief models capture two basic principles of learning—the law of actual effect and the law of simulated effect.

¹¹The parameter ρ disappears because it only appears in the updating equation denominator $\rho \cdot N(t-1) + 1$ that is common to all attractions and thus cancels out in the power form. Then EWA attractions at time t depend only on recent payoffs and the product $A_i^j(0) \cdot N(0)$. While initial choice probabilities depend on $A_i^j(0)$ only, these probabilities are the same as those that depend on $A_i^j(0) \cdot N(0)$ (for $N(0) > 0$). As a result, multiplying the initial attractions by an arbitrary constant makes no difference (econometrically, $N(0)$ is not identifiable).

Many decades of learning experiments, mostly with (nonhuman) animal subjects, show that successful chosen strategies are subsequently chosen more often. Behaviorist psychologists call this the ‘law of effect’ (Thorndike (1911), Herrnstein (1970)). We relabel this the ‘law of *actual* effect’ because behaviorists took it for granted for years that the only effect on subsequent choices was produced by rewards for actual choices. The behaviorists eschewed ‘mentalist’ constructs like imagination, which allowed the possibility that foregone rewards could affect the probability of choosing new strategies, until a series of demonstrations showed that those cognitive constructs are necessary. When applied to humans playing games with a known payoff matrix, it is sensible to propose a corollary general principle, the ‘law of *simulated* effect’. The law of simulated effect states that unchosen strategies that would have yielded high payoffs—simulated successes—are more likely to be chosen subsequently. Many experiments on reinforcement learning are consistent with this principle.¹²

Furthermore, most research on human and machine learning assumes that the basic process driving learning is not reinforcement, per se, but the reduction of errors. Since errors are measured by the difference between what players received and what they could have received, error-reduction algorithms effectively use both actual payoffs and foregone payoffs.

This error reduction idea also lies behind learning direction theory (Selten and Stoecker (1986), Selten (1997)). Learning direction theory presumes players have a causal understanding of the game that enables them to tell in which direction they should switch strategies.¹³ If players know strategies’ foregone payoffs, then direction learning predicts they will move (weakly) in the direction of higher-forgone payoffs, and away from low foregone payoffs. This is essentially the same prediction as EWA with $\delta = 1$, except that the direction learning allows inertia in responses, which corresponds to $\delta < 1$ when payoffs are positive.¹⁴ Thus, in our view EWA incorporates the intuition behind direction learning in a precise way, when the causal structure is known, while direction learning can more generally apply to situations with known causal structure (but unknown foregone payoffs) in a way yet to be fully specified.

The empirical strengths of the law of effect and the law of simulated effect are the key to distinguishing different models of learning in games, and are

¹²For example, anxious patients can be taught to fear a picture of a triangle (a conditioned stimulus, or CS) when it is followed by a loud annoying noise (an unconditioned stimulus, or UCS). When patients are told to simply imagine the UCS several times, their imagination increases the strength of their conditioned fear response to the triangle CS (Davey and Matchett (1994)). A related phenomenon is ‘incubation’, in which presentation of the CS itself increases the fear response (Eysenck (1979)). In these cases, people are not learning by direct reinforcement. They ‘learn’ by simply imagining either the UCS’s reinforcement, or the reinforcement that typically follows a CS.

¹³The players’ understanding could be expressed in a causal diagram or map, but this central part of the theory is not yet developed.

¹⁴When payoffs are negative, $\delta < 1$ in EWA implies players will be likely to move away from money-losing chosen strategies (which are reinforced fully, and negatively), even moving to strategies with larger negative (foregone payoffs).

calibrated by δ . Reinforcement insists that only actual effects matter ($\delta = 0$). Belief models implicitly require that actual and simulated effects are equally strong ($\delta = 1$). EWA takes the middle ground.

The parameter δ could also be interpreted as creating an endogenous aspiration level against which payoffs are compared. Including an aspiration level is sensible because many studies show that the reinforcement value of a fixed payoff depends on the aspiration level to which the payoff is compared (e.g., Erev and Roth (1997)).¹⁵

In EWA, δ creates an adjustable aspiration level endogeneously. It is easy to show that reinforcing strategies according to foregone payoffs means the probability of a chosen strategy $s_i(t)$ only increases if its payoff is larger than δ times the average foregone payoff, holding previous attractions constant (see our working paper for details). Therefore, one can interpret δ times the average foregone payoff in each period as a kind of aspiration level. A larger δ creates a higher aspiration level. Furthermore, the aspiration level adjusts automatically over time, because it depends on the foregone payoffs in each period. EWA therefore creates an endogenous, adjustable aspiration level at no extra parametric cost.

If δ is interpreted as the weight placed on foregone payoffs, many generalizations spring to mind. The size of the weight δ could depend on the size of the foregone payoff or on its sign, to allow the possibilities that unusually large or small foregone payoffs catch a player's attention, or that players are more sensitive to losses than to gains (cf. loss-aversion in risky choices, e.g., Tversky and Kahneman (1992)). If players are more sensitive to foregone payoffs for strategies that are closer to the chosen strategy, or more similar, then δ will depend on the distance or similarity between each strategy and the chosen strategy $s_i(t)$ (cf. Sarin and Vahid (1997)).

If δ is applied to others' actual payoffs instead of own foregone payoffs, EWA can be used to capture learning by imitation. Imitation is obviously common, especially among animals, children, and impressionable teenagers. Payoff-dependent imitation is also a sensible heuristic behavior in low-information environments where players do not know what their foregone payoffs are, but can observe success of other firms. EWA and imitation learning will be approximately the same when (i) games are symmetric, so that another player's payoffs are the same as one's own foregone payoffs, and (ii) when there are many players who choose different strategies, so that a player who imitates others according to how successful they were is effectively reinforcing a wide range of her own strategies according to their foregone payoffs. We conjecture that these are the conditions under which imitation is most common. If so, then imitation is just a heuristic way to implement foregone-payoff-based updating (a la EWA),

¹⁵ Players who tend to repeat previously-chosen strategies, regardless of their outcomes, reveal a 'status quo bias' or 'habit' (Majure (1994), Tang, (1996)). This habitual behavior of chosen strategies can be captured by having an aspiration level that is always lower than the actual payoffs and by a lack of simulated effect.

using payoffs of others as the best available proxy for one's own unknown payoffs. In this sense, EWA captures some of the intuition underlying imitation, and perhaps much of its empirical force as well.

3.2. *Growth of Attractions, ρ and ϕ*

The parameter ϕ depreciates past attractions, $A_i^j(t)$.¹⁶ The parameter ρ depreciates the experience measure $N(t)$. It captures decay in the strength of prior beliefs, which can be different than decay of early attraction (captured by ϕ). These factors combine cognitive phenomena like forgetting with a deliberate tendency to discount old experience when the environment is changing.

One way to interpret ρ and ϕ is by considering the numerator and denominator of the main EWA updating equation (2.2) separately, and thinking about how reinforcement and belief-based models use these two terms differently. The numerator is $\phi \cdot N(t-1) \cdot A_i^j(t-1) + [\delta + (1-\delta) \cdot I(s_i^j, s_i(t))] \cdot \pi_i(s_i^j, s_{-i}(t))$. This term is a running total of (depreciated) attraction, updated by each period's payoffs. The denominator is $\rho \cdot N(t-1) + 1$. This term is a running total of (depreciated) periods of experience-equivalence. Reinforcement models essentially keep track of the running total in the numerator, and do not adjust for the number of periods of experience-equivalence (since $\rho = 0$, the denominator is always one). Belief-based models also keep track of the attraction total but divide by the total number of periods of experience-equivalence. By depreciating the two totals at the same rate ($\rho = \phi$), the belief-based models keep the 'per-period' attractions (expected payoffs) in a range bounded by the game's payoffs.

EWA allows attractions to grow faster than an average, but slower than a cumulative total. An analogy might help illustrate. Instead of determining attractions of strategies, think about evaluating a person (for example, an athlete, or a senior colleague you might hire) based on a stream of lifetime performances. The reinforcement model evaluates people based on (depreciated) lifetime performance. The belief-based models evaluate people based on 'average' (depreciated) performance. Both statistics are probably useful in evaluation—in hiring a colleague or an athlete, you would want to know lifetime performance *and* some kind of performance averaged across experience. One way to mix the two is to normalize depreciated cumulative performance by depreciated experience, but depreciate the amount of experience more rapidly. Then if two people perform equally well on average every year, the person with 10 years of experience is rated somewhere between equally as good and twice as good as the person with five years of experience. When $\phi > \rho$, EWA models players who use something in between 'lifetime' performance and 'average' performance to evaluate strategies.

¹⁶A 'primary effect' (or 'imprinting'; Cheung and Friedman (1997)) in which early observations are remembered more strongly than recent ones, can be expressed by $\phi \geq 1$.

The depreciation rate parameters ϕ and ρ can also be understood by how they control slowdown in learning rate or sharpness of convergence. Solving recursively for steady-state attraction levels shows that those levels equal the ratio $(1 - \rho)/(1 - \phi)$ times the steady-state average payoff. Thus, when $\rho = 0$ as in reinforcement learning, attractions can end up outside the bounds of payoff levels (and they grow as large as possible, holding ϕ constant). When $\rho = \phi$, as in belief-learning, steady-state attraction levels are equal to steady-state average payoffs. The implication of these two possibilities depends on how attractions determine probabilities. In the logit probability form, only differences in attraction levels affect choice probabilities. Therefore, given a fixed value of λ , attractions that can grow outside the bounds of payoff levels have a wider range across strategies. This allows the possibility of sharper convergence in the sense that choice probabilities can converge closer to the boundaries at zero and one. When attractions are bounded to be close to payoff levels, convergence cannot be as sharp. In the power probability form, only ratios of attraction levels matter. Therefore, if attractions grow, the relative impact of new reinforcements falls; learning slows down. *Ceteris paribus*, reinforcement learning requires convergence to be as sharp as possible (in the logit form) or requires learning to slow down as quickly as possible (in the power form), while belief learning requires the opposite. EWA is able to choose an intermediate value of ρ that tailors the sharpness of convergence or rate of learning to the data.

3.3. Initial Attractions $A_i^j(0)$ and their Strength $N(0)$

The term $A_i^j(0)$ represents the initial attraction, which might be derived from an analysis of the game, from surface similarity between strategies and strategies that were successful in similar games, etc. Belief models restrict the $A_i^j(0)$ strongly by requiring initial attractions to be derived from prior beliefs. This requires, for example, that weakly dominated strategies will always have (weakly) lower initial attractions than dominant strategies. EWA allows more flexibility.

For example, suppose players make first-period choices randomly, by choosing what was chosen previously in a different game, by setting each strategy's initial attraction equal to its minimum payoff (the maximin rule) or maximum payoff (the maximax rule),¹⁷ or by choosing stochastically among selection principles like payoff-dominance, risk-dominance, loss-avoidance, etc. All these decision rules are plausible models of first-period play, but none of them generate initial attractions that are always expected payoffs given some prior beliefs.

We consider the scientific problem of figuring out how people choose their initial strategies as fundamentally different than explaining how they learn.

¹⁷ Making a strategy's initial attraction equal to its minimum payoff, for example, is implicitly putting all the belief weight on the choices by others that yield that minimum. But the choices by others that lead to minima for different strategies are likely to be different. So the implicit beliefs underlying each attraction will be different.

Leaving initial attractions unrestricted makes them numerical placeholders that can be filled by a theory of first-period play that supplies attractions as an input to EWA. That combination would be a complete theory of behavior in games, from start to finish.

The initial-attraction weight $N(0)$ appears in the EWA model to allow players in belief-based models to have an initial prior that has a certain strength (measured in units of actual experience). In EWA, $N(0)$ is therefore naturally interpreted as the strength of initial attractions, relative to incremental changes in attractions due to actual experience and payoffs. Fixing $N(0) = 1$ means that, unit for unit, initial attractions $A_i^j(0)$ and chunks of reinforcement from payoffs are weighed equally when attractions are updated. This is easiest to see by fixing $\delta = 1$ for simplicity and directly computing the attraction after two periods, $A_i^j(2)$, which gives

$$(3.1) \quad A_i^j(2) = \frac{\phi^2 \cdot A_i^j(0) \cdot N(0) + \phi \cdot \pi_i(s_i^j, s_{-i}(1)) \cdot \pi_i(s_i^j, s_{-i}(2))}{\rho^2 \cdot N(0) + \rho + 1}.$$

The parameter ϕ captures the declining weight placed on payoffs from more distant periods of actual experience, compared to more recent periods. (That is, the older period 1 payoff $\pi_i(s_i^j, s_{-i}(1))$ is weighted by ϕ but the recent period 2 payoff $\pi_i(s_i^j, s_{-i}(2))$ is not.) Like previous payoffs, the initial attraction is also weighted by a power of ϕ (ϕ^2 , because it ‘happened’ two periods earlier), but is also weighted by $N(0)$. Thus, the parameter $N(0)$ captures the special weight placed on the initial attractions, compared to increments in attraction due to payoffs. $N(0)$ can therefore be thought of as a ‘pre-game (introspective) experience’ weight. If $N(0)$ is small the effect of the initial attractions is quickly displaced by experience. If $N(0)$ is large then the effect of the initial attractions persists.

Notice that updating the experience-weight by $N(t) = \rho \cdot N(t-1) + 1$ implies a steady-state value of $N^* = 1/(1-\rho)$. In estimation, we have found it useful to restrict $N(0)$ to be less than N^* . This implies $N(t-1) \leq N(t)$; the experience weight is (weakly) rising over time. Since the relative weight on decayed attractions, compared to recent reinforcement, is always increasing, the relative weight on observed payoffs is always declining. This implies a ‘law of declining effect’ that is widely observed in research of learning.

The flexibility of initial attractions and experience weight allows one to fit a variety of models. Theories of equilibrium behavior are special cases in which all ‘learning’ occurs before the game starts. For example, a ‘stubborn’ game-theoretically-minded player sets $A_i^j(0)$ equal to the equilibrium payoffs of each strategy and act as if $N(0)$ is infinite (meaning that no amount of game-playing experience can outweigh the prior calculation). An adaptive game theorist assumes $A_i^j(0)$ are equilibrium payoffs but has a small $N(0)$, so she learns from experience. A player who does not begin with prior beliefs, but updates accord-

ing to experience as a belief learner does, has $\phi = \rho$ and $\delta = 1$ with arbitrary $A_i^j(0)$.

4. PREVIOUS RESEARCH

In this section we briefly summarize previous research (see Camerer (in progress) for more details).

Several papers investigate only belief learning. Cheung and Friedman (1997) (CF) estimated a weighted fictitious play model on individual-level data from four games (hawkdove, stag hunt, 'buyer seller' and battle-of-the-sexes). They find substantial heterogeneity across subjects but stability across games in parameters that are like our ϕ and λ . A more general belief model, allowing idiosyncratic shocks in beliefs and time-varying weights, was developed by Crawford (1995) to fit data from coordination games, extended by Broseta (1995) to allow ARCH error terms, and applied by Crawford and Broseta (1998) to coordination with preplay auctions. Brandts and Holt (in press) and Cooper, Garvin, and Kagel (in press) simulate fictitious play in signaling games. Boylan and El-Gamal (1992) compare fictitious play and Cournot learning in coordination and dominance-solvable games; they find overwhelming relative support for fictitious play.

Other studies concentrate only on reinforcement learning. Versions of reinforcement in which probabilities were reinforced directly, or cumulative payoffs normalized, were used by Bush and Mosteller (1955), Cross (1983), and Arthur (1991). Harley (1981) posited a reinforcement model using cumulative payoffs and simulated its behavior in several games. The Harley model was later extended by Roth and Erev (1995) to include spillover of reinforcement to neighboring strategies. Their model fits the time trends in ultimatum, public good, and responder-competition games but converges much too slowly. McAlister (1991) shows that a modified Cross model that uses foregone payoff information fits weak-link data modestly well. Sarin and Vahid (1997) show that a modified Cross model with distance-weighted spillover of reinforcement to similar strategies fits data on coordination experiments with low information fairly well.

These studies of belief and reinforcement learning find that each approach, evaluated separately, has some explanatory power. Other studies compared models.¹⁸ Erev and Roth (1997) add an adjustable reference point to their earlier model (cf. Cross (1983)). The extended model fits slightly better than fictitious play, at the individual level, in constant-sum games played for 100 or

¹⁸In still another approach, models in which players learn to shift weight across various rules (or 'methods'), rather than across strategies, were studied by Tang (1996) and by Stahl (1996, 1997). In Tang's comparison 'method-learning' does slightly worse than reinforcement. Stahl (1997) finds that players seem to weight rules that mimic choices of others or best-respond given diffuse priors.

more periods. Mookerjee and Sopher (1994, 1997) (MS) compare average-payoff reinforcement and fictitious play in constant-sum games; reinforcement does somewhat better. Ho and Weigelt (1996) compare modified versions of fictitious play and choice reinforcement (the MS 'vindication' model) in coordination games with multiple Nash equilibria. Fictitious play fits better. Battalio, Samuelson, and Van Huyck (1997) compare average-payoff reinforcement and fictitious play in three variants of stag hunt with the same equilibria but different disequilibrium incentives to converge. Fictitious play does better but a model that uses both reinforcements and expected payoffs does better still.

Many variants of weighted fictitious play and reinforcement (and other models) were compared by Tang (1996) in games with mixed-strategy equilibria. Reinforcement does better in most games. Chen and Tang (1998) fit models to data from two public goods games. In one game equilibration is so fast that Nash equilibrium outpredicts the learning models. In the other game reinforcement does better.

The overall picture from previous research is somewhat blurry. Comparisons appear to favor reinforcement in constant-sum games and belief learning in coordination games. However, specifications of the models and estimation techniques vary across studies. Our approach allows one to compare models more systematically by including all features that have been used differently in different studies. Two general features are especially notable.

First, most papers assume equal initial attractions or, for belief models, uniform priors. Some papers estimate initial attractions using data from early periods (which does not generally optimize overall fit). Our procedure is more general because we estimate initial attractions and experience weight as part of an overall maximization of fit. Estimating initial experience weight $N(0)$ allows belief models to express a prior strength. This is an important feature of belief learning; omitting it may explain why belief models have sometimes fit relatively poorly (in Mookerjee and Sopher (1997), Tang (1996), Erev and Roth (1997), Chen and Tang (1998)).

Second, some reinforcement models assume averaged payoffs affect choices, while others assume reinforcements cumulate. This difference can be captured by allowing ρ to vary between ϕ (for averaging) and 0 (for maximum cumulation), as EWA does. In addition, many studies of belief learning did not allow weighted fictitious play, as EWA does. Including ϕ and ρ therefore allows us to determine whether previous mixed results depend on whether reinforcements are averaged or cumulated, and on whether belief models are weighted.

Our methodology for model estimation is more general than most earlier papers in four ways. First, we compare across three classes of games using the same estimation technique (only Cheung and Friedman (1997) have done this in one paper). Second, our method uses standard statistical tests to judge whether differences in fit are due to chance, or put differently, to decide whether simple models are too simple or not. Third, we calibrate models on the first 70% of the periods in each sample and predict the rest of the sample to validate the estimates and avoid overfitting (no previous paper has done this). Fourth, we

allow heterogeneity across individuals by comparing a model with a single class of agents with a two-segment model, which has not been done before.¹⁹

5. PARAMETER ESTIMATION FROM EXPERIMENTAL DATA

5.1. *Estimation Strategy*

We estimated the values of model parameters using three samples of experimental data²⁰ and validated the models by predicting behavior out of sample. The games are: Constant-sum games with unique mixed-strategy equilibria (and one weakly dominated strategy); a ‘median-action’ coordination with multiple Pareto-ranked equilibria; and a dominance-solvable ‘*p*-beauty contest’ game with a unique equilibrium. We chose these games for several reasons.

First, the games have a range of different structural features (as in Cheung and Friedman (1997) and Stahl (1997)). This avoids the possible mistake of concluding that a model generally fits well because it happens to fit one class of games.

Second, the games have different spans—the constant-sum games last 40 periods and the others last 10 periods. Longer spans provide more data and more power for estimating individual differences. But a mixture of long and short spans are valuable too, because some games—like the coordination games—converge quickly. Learning models should be able to explain why convergence is quick in those games and slow in others.

Third, most previous studies have reported results that are favorable to either reinforcement or belief learning. The games we use each present some new challenges to these models. The presence of dominated strategies in the constant-sum games is a challenge for belief models, which predict those strategies will be played relatively rarely. Rapid convergence in the coordination and dominance-solvable games is a challenge for reinforcement learning (see also Van Huyck, Battalio, and Rankin (1996)), which tends to be sluggish.

Next we describe some general features of the estimation method. For simplicity we assume that players’ strategies are the stage-game strategies, and denote player *i*’s strategy choice in period *t* by $s_i(t)$. (Of course, in general

¹⁹The only paper that estimates individual-level parameters on these kinds of models is Cheung and Friedman (1997). While the median parameter estimates are reasonable and similar across games when expected to be, the individual-level estimates are variable (e.g., a third of the ϕ estimates are negative and a sixth are above one). This reflects some imprecision in individual-level estimation that suggests that multiple-segment estimation, which lies between single-segment estimation and individual-level estimation, may be a reasonably parsimonious compromise between the desires to allow heterogeneity and to estimate reliably.

²⁰Our working paper includes two other samples of data, on weak-link coordination games and matching pennies (Mookerjee and Sopher (1994)). We dropped these because the weak-link data did not have a long enough span to permit both calibration and validation; calibration is reported in Camerer and Ho (1999). The matching pennies data did not distinguish models from each other or from Nash equilibrium.

strategies could be history-dependent or could be decision rules; we say more about this in the conclusion.)

We use a ‘latent class’ approach in which there are one or two segments of players, and all players in a segment are assumed to have the same parameter values. This technique is standard in some fields (e.g., analyses of brand choice in marketing) and was also suggested by Crawford (1995).²¹ The single-class estimation provides a representative-agent benchmark. Allowing a second class gives a clue about how important it is to allow heterogeneity. While allowing heterogeneity should obviously improve fit in most cases, it does not always. For example, in constant-sum games, allowing a second class hardly improves the fit at all.

The two-class procedure makes sense for these data sets because there are not enough observations per subject to reliably estimate many more classes.²² And while including more segments would be desirable, there is no reason to think that assuming all players have the same parameters favors some models over others, so this simplification should not alter our conclusions about which models fit better.

Estimating initial attractions $A^i(0)$ (suppressing the player subscript), as we do, has three advantages. First, assuming initial attractions are equal (or assuming equal priors in belief models) saves degrees of freedom but fits poorly in some of these data sets. Second, estimating initial attractions creates raw material that may be useful for constructing a good theory of first-period play; in a sense, any such theory is trying to *predict* the values that we estimate. Third, forcing a model to maximize fit by estimating learning parameters *and* initial attractions at the same time allows the possibility that the model will misspecify the initial attractions ‘on purpose’. That is, if a model cannot easily explain how players move away from initial behavior as they learn, it will misestimate the initial behavior. Therefore, estimating the initial attractions and comparing them with first-period play serves as an indirect test for gross specification error. For example, we see below that reinforcement models fail this test for coordination games.

Let the stage game be repeated for T rounds. Recall that the indicator function $I(s_i^j, s_i(t))$ is equal to 1 if $s_i^j = s_i(t)$ and 0 otherwise. Define the vector of initial attractions for player i to be $A_i(0) \stackrel{\text{def}}{=} (A_i^1(0), A_i^2(0), \dots, A_i^m(0))$. Since we study symmetric games and assume all players have the same parameter values, for this paper there is a common set of initial attractions $A(0) = A_i(0) \forall i$. Define the number of subjects by N . The overall sample size for calibration,

²¹Note that even though all agents in a class have the same parameter values, after the first period they will be predicted to behave differently because their actual choices and experiences vary.

²²Two segments are also useful because one can then compare a two-segment EWA model with a two-segment model in which one segment consists of reinforcement learners and the other segment contains belief learners, to see whether EWA is fitting better than a ‘population mixture’ of belief and reinforcement learning. We did this in Camerer and Ho (1998) for weak-link and median-action data and found that EWA does fit much better than the mixture model.

$.7 \cdot T \cdot N$, is denoted by M . Then the log-likelihood function, $LL(A(0), N(0), \phi, \rho, \delta, \lambda)$, is

$$(5.1) \quad LL(A(0), N(0), \phi, \rho, \delta, \lambda)$$

$$(5.2) \quad \begin{aligned} &= \sum_{t=1}^{0.7 \cdot T} \sum_{i=1}^N \ln \left(\sum_{j=1}^{m_i} I(s_i^j, s_i(t)) \cdot P_i^j(t) \right) \\ &= \sum_{t=1}^{0.7 \cdot T} \sum_{i=1}^N \ln \left(\sum_{j=1}^{m_i} I(s_i^j, s_i(t)) \cdot \frac{e^{\lambda \cdot A_i^j(t-1)}}{\sum_{k=1}^{m_i} e^{\lambda \cdot A_i^k(t-1)}} \right). \end{aligned}$$

Keep in mind that in the exponential form, attractions are only identified up to a constant, so we must fix one of the $A^j(0)$ to equal a constant. We searched over parameter values to maximize the LL function using the MAXLIK routine in GAUSS, which uses a gradient method. To avoid converging to local optima we tried a variety of starting points. We restricted ϕ, λ to be positive, $0 \leq \delta, \rho \leq 1$.

Generally, violations of these restrictions should be interpreted as either signs of misspecification, or evidence that the model is trying to reach outside reasonable parameter values to explain an unusual feature of data. Fortunately, there are only two cases where the EWA restrictions bind— $\hat{\delta} = 0$ in constant-sum game G1, which we can't easily explain, and $\hat{\rho} = 0$ in coordination games, which we can explain.²³

In order to make the value of $N(0)$ interpretable as a weight on initial attractions relative to reinforcing payoffs, we restricted the range of $A^j(0)$ to be less than or equal to the difference between the minimum and maximum payoffs in the entire game (while also setting one of the attractions equal to zero for identifiability).²⁴ Since this restriction is naturally satisfied in belief models, in order to compare EWA to belief and reinforcement learning we imposed it in EWA and reinforcement as well.²⁵ We also restricted $0 \leq N(0) \leq 1/(1 - \rho)$ to guarantee that the weights $N(t)$ rise over time.

²³In the coordination games convergence is very sharp after substantial initial dispersion. To explain this in the logit model, EWA needs attractions to grow apart as rapidly as possible, so high attraction strategies are chosen very frequently (and more frequently than in early periods). Since the denominator of the attraction equation is $\rho \cdot N(t-1) + 1$, attractions grow faster if this denominator shrinks, which means a negative value of ρ can make growth even faster. Therefore, the restricted estimate $\hat{\rho} = 0$ in the coordination games means the model is trying to make ρ negative to capture the remarkable pace of equilibration seen in those data.

²⁴If the attractions are not restricted in this way, then the experience weight $N(0)$ expresses both the relative weight on initial attractions and payoffs, and a scaling factor which puts attractions and payoffs on the same scale. By restricting attractions to have the same range as payoffs, we can then interpret $N(0)$ as a relative weight.

²⁵In our working paper we allowed initial attractions to have arbitrary scale, which made MLE convergence slower and identification worse. Allowing arbitrary attractions helps reinforcement a bit in constant-sum games but does not help much in median-action and beauty-contest games.

Standard errors of parameters were estimated using a jackknife procedure. In each run of the jackknife, one subject was excluded from the analysis and the model was estimated using all remaining subjects.²⁶ Doing this sequentially produces N vectors of estimates (where N is the number of subjects). The parameter standard errors are then the standard deviations of parameter estimates across the N runs. (Correlations between parameters can also be computed this way, and help detect identification problems.)

Since EWA is always more general than the special cases, it will necessarily fit the data better so there is some danger of overfitting. To guard against this we use five ways to penalize theories that have more degrees of freedom. We calibrate the models by deriving MLE estimates using the first 70% of the observations in each sample, and measure goodness of fit three ways. Then we validate the models by using the derived estimates to predict the path of play in the remaining 30% of the sample and measure fit two ways. This procedure uses enough data to estimate parameters reliably, but also forecasts out-of-sample to ensure models are not being overfit.

Obviously, deriving estimates on one part of the sample and predicting the rest of the sample implicitly assumes stationarity in parameters throughout the sample. In games that converge within an experiment, it is an open question how much power out-of-sample forecasting has to distinguish theories when behavior in the later, holdout-sample periods does not vary much. As is shown below for coordination games, the power can be substantial. The reason is that some theories *do* predict a lot of variation in later periods. If there is little variation in the holdout sample, the *absence* of variation counts against a theory that predicts variation. Also, we use the first 70% of the periods in an experimental session to calibrate parameter estimates and the final 30% of the periods to validate (forecast). A different approach is to derive estimates for some fraction of the *subjects* (using their entire samples), and use those estimates to forecast behavior for a holdout sample of subjects across the entire experimental session. Future work could use that method and compare the results to ours.

To evaluate model accuracy in the calibration phase, the three criteria we report are: Log likelihoods (these are used in χ^2 tests, which effectively penalize for EWA), and Akaike and Bayesian information criteria, which penalize theories according to the number of free parameters in different ways.²⁷ We also report a pseudo- R^2 , denoted ρ^2 , based on the Akaike measure, so one can see how much better the models do than random choice.²⁸ For the validation sample we report the log likelihood and a mean squared deviation

²⁶ For the constant-sum games, with only twenty subjects per game (ten row players and ten column players), every pair of row and column players were excluded, giving 100 jackknife runs.

²⁷ The Akaike criterion (AIC) is $LL - k$ and the Bayesian criterion (BIC) is $LL - (k/2) \cdot \log(M)$ where k is the number of degrees of freedom and M is the size of the calibration sample.

²⁸ The measure ρ^2 is the difference between the Akaike measure and the log likelihood of a model of random choices, normalized by the random-model log likelihood.

(MSD), which is defined as

$$(5.3) \quad MSD = \sum_{t=.7 \cdot T+1}^T \sum_{i=1}^N \sum_{j=1}^{m_i} \frac{[P_i^j(t) - I(s_i^j, s_i(t))]^2}{.3 \cdot T \cdot N \cdot m_i}.$$

(Note that this MSD does not average observations across individuals.) Model fits are also compared to a random choice model in which all strategies are chosen equally often in each period.²⁹

Since out-of-sample forecasting is a tougher test than in-sample fitting, and levels the playing field among theories with different numbers of parameters, why report in-sample fits at all? One reason is that both measures of fit are needed to judge whether a model does much better in-sample than out-of-sample, which is the telltale sign of egregious over-fitting. Furthermore, most previous studies have relied solely on in-sample fits for comparing models. Reporting them allows readers familiar with such comparisons to weigh the results against their own standards.

For each game, we describe the game and basic details of how the experiments were conducted. Then we compare models and discuss parameter estimates.

Table I previews and summarizes the results. Within each game and measure, other than LL and ρ^2 , the best fit statistic is printed in italics and marked with an asterisk. In both the calibration and validation phases, EWA fits substantially better in four of six games; in two cases the belief models fit a little better. (If EWA was overfitting, it would do relatively better in calibration than in validation, but this isn't the case.) Belief models do better than reinforcement in constant-sum games and worse in the median-action game. In the beauty contest game, the belief model does worse than reinforcement during calibration and better during validation. The two-segment models generally fit a little better during both validation and calibration, but the improvement in fit over one-segment models is small.

5.2. Constant-sum Games with Dominated Actions

We fit data from four constant-sum games: two are 4×4 (G1 and G3) and the others are 6×6 (G2 and G4) from Mookherjee and Sopher (1997). Tables IIa–b show the payoff matrices.³⁰ The 4×4 games essentially collapse three of the undominated actions (actions 3–5) of the 6×6 games into a single action (action 3).

²⁹We do not compare results with Nash equilibrium, as many studies do, because it does very poorly in constant sum games and beauty-contest games (in which iteratively-dominated strategies predicted to have zero probability are often played) and does not exclude any choices in the coordination games.

³⁰The fractional payoffs (e.g., $2/3W$), denote probabilistic chances of winning W . These present a complication for reinforcement models, including EWA—do you reinforce the actual payoff (which has a one-third chance of being zero if $2/3W$ is the payoff) or the expected payoff? We reinforce according to the expected payoff.

TABLE I
MODEL CALIBRATION AND VALIDATION IN MEDIAN-ACTION, CONSTANT-SUM,
AND P-BEAUTY CONTEST GAMES

Game Model	No. of Parameters	Calibration			Validation		
		LL	AIC	BIC	ρ^2	LL	MSD
Constant-sum 1 (M = 560)							
1-Segment							
Random	0	− 803.69	− 803.69	− 803.69	0.0000	− 332.71	0.1875
Choice Reinforcement	5	− 681.97	− 686.97	− 697.96	0.1452	− 335.50	0.1888
Belief-based	6	− 680.23	− 686.23	− 699.42	0.1461	− 285.60*	0.1688*
EWA	8	− 653.07	− 661.07*	− 678.66*	0.1775	− 326.38	0.1883
2-Segment							
Random	0	− 803.69	− 803.69	− 803.69	0.0000	− 332.71	0.1875
Choice Reinforcement	11	− 681.69	− 692.69	− 716.87	0.1381	− 328.71	0.1905
Belief-based	13	− 680.12	− 693.12	− 721.70	0.1376	− 285.96	0.1689
EWA	17	− 652.46	− 669.46	− 706.83	0.1670	− 317.66	0.1865
Constant-sum 2 (M = 560)							
1-Segment							
Random	0	− 1003.39	− 1003.39	− 1003.39	0.0000	− 430.02	0.1435
Choice Reinforcement	7	− 853.61	− 860.61	− 876.00	0.1423	− 359.74	0.1222
Belief-based	8	− 797.72	− 805.72	− 823.31	0.1970	− 350.09	0.1212
EWA	10	− 790.61	− 800.61	− 822.59*	0.2021	− 341.71	0.1179*
2-Segment							
Random	0	− 1003.39	− 1003.39	− 1003.39	0.0000	− 430.02	0.1435
Choice Reinforcement	15	− 853.50	− 868.50	− 901.48	0.1344	− 363.10	0.1513
Belief-based	17	− 790.10	− 807.10	− 844.47	0.1956	− 347.50	0.1205
EWA	21	− 776.83	− 797.83*	− 844.00	0.2049	− 335.95*	0.1195
Constant-sum 3 (M = 560)							
1-Segment							
Random	0	− 803.69	− 803.69	− 803.69	0.0000	− 332.71	0.1875
Choice Reinforcement	5	− 710.14	− 715.14	− 726.13	0.1102	− 308.47	0.1801
Belief-based	6	− 681.63	− 687.63	− 700.82*	0.1444	− 296.28	0.1728
EWA	8	− 678.50	− 686.50*	− 704.09	0.1458	− 301.70	0.1767
2-Segment							
Random	0	− 803.69	− 803.69	− 803.69	0.0000	− 332.71	0.1875
Choice Reinforcement	11	− 710.14	− 721.14	− 745.32	0.1027	− 308.30	0.1800
Belief-based	13	− 681.07	− 694.07	− 722.65	0.1364	− 295.85*	0.1722
EWA	17	− 677.60	− 694.60	− 731.97	0.1357	− 302.88	0.1712*
Constant-sum 4 (M = 560)							
1-Segment							
Random	0	− 1003.39	− 1003.39	− 1003.39	0.0000	− 430.02	0.1435
Choice Reinforcement	7	− 901.60	− 908.60	− 923.99	0.0945	− 375.94	0.1284
Belief-based	8	− 857.19	− 865.19*	− 882.78*	0.1377	− 371.18	0.1262
EWA	10	− 855.29	− 865.29	− 887.27	0.1376	− 362.26	0.1241
2-Segment							
Random	0	− 1003.39	− 1003.39	− 1003.39	0.0000	− 430.02	0.1435
Choice Reinforcement	15	− 901.60	− 916.60	− 949.58	0.0865	− 375.62	0.1283
Belief-based	17	− 856.97	− 873.97	− 911.34	0.1290	− 372.12	0.1265
EWA	21	− 854.00	− 875.00	− 921.17	0.1280	− 361.15*	0.1239*

TABLE I

Continued

Game Model	No. of Parameters	Calibration			ρ^2	Validation	
		LL	AIC	BIC		LL	MSD
Median Action (M = 378)							
1-Segment							
Random Choice	0	−677.29	−677.29	−677.29	0.0000	−315.24	0.1217
Choice Reinforcement	8	−341.70	−349.70	−365.44	0.4837	−80.27	0.0301
Belief-based	9	−438.74	−447.74	−465.45	0.3389	−113.90	0.0519
EWA	11	−309.30	−320.30	−341.94*	0.5271	−41.05	0.0185
2-Segment							
Random	0	−677.29	−677.29	−677.29	0.0000	−315.24	0.1217
Choice Reinforcement	17	−331.25	−348.25	−381.70	0.4858	−66.32	0.0245
Belief-based	19	−379.24	−398.24	−435.62	0.4120	−70.31	0.0250
EWA	23	−290.25	−313.25*	−358.51	0.5375	−34.79*	0.0139*
p-beauty contests (M = 1372)							
1-Segment							
Random	0	−6318.29	−6318.29	−6318.29	0.0000	−2707.84	0.0099
Choice Reinforcement	12	−5910.99	−5922.99	−5954.33	0.0626	−2594.37	0.0101
Belief-based	13	−6083.04	−6096.04	−6129.99	0.0352	−2554.21	0.0097
EWA	15	−5878.20	−5893.20	−5932.38	0.0673	−2381.28	0.0098
2-Segment							
Random	0	−6318.29	−6318.29	−6318.29	0.0000	−2707.84	0.0099
Choice Reinforcement	25	−5910.98	−5935.98	−6001.28	0.0605	−2594.17	0.0101
Belief-based	27	−6083.02	−6110.02	−6180.54	0.0330	−2554.11	0.0097*
EWA	31	−5771.46	−5802.46*	−5883.43*	0.0816	−2355.00*	0.0098

Note that these games each have a weakly dominated action (action 4 in G1 and G3 and 6 in G2 and G4). Dominated actions are useful for model discrimination because belief-based models always predict these actions will be chosen (weakly) less frequently than dominant actions, whereas the arbitrary initial attractions allowed by EWA and choice reinforcement can allow frequent choices of dominated strategies.

All these games have a unique mixed strategy equilibrium that is symmetric (even though the games are not symmetric). In games G1 and G3, in equilibrium actions 1–4 are played with probabilities $3/8$, $2/8$, $3/8$, 0 respectively. In games G2 and G4, equilibrium proportions are $3/8$, $2/8$, $1/8$, $1/8$, $1/8$, 0 for actions 1–6.

Each game was played by 10 different pairs of subjects playing with the same partner 40 times. At the end of each period players were told their partner's choice and their own payoff. In games G1 and G2 a win paid 5 rupees; in games G3 and G4 the payoffs were doubled to 10 rupees. (A typical student's monthly room and board cost 600 rupees.)

We derived MLE parameter estimates using the first 28 periods, and validated by predicting the last 12 periods. Because the payoff matrix is not symmetric (even though the equilibrium mixed-strategy proportions are), we estimate separate initial attractions $A_i^j(0)$ and separate initial experience-weights $N_i^j(0)$

TABLE IIa
4 × 4 CONSTANT SUM GAMES, G1 ($W = 5$ RUPEES);
AND G3 ($W = 10$ RUPEES)

		Column			
		S1	S2	S3	S4
Row	S1	W, L	L, W	L, W	W, L
	S2	L, W	L, W	W, L	W, L
	S3	L, W	W, L	1/3W, 2/3W	1/3W, 2/3W
	S4	L, W	L, W	2/3W, 1/3W	W, L

TABLE IIb
6 × 6 CONSTANT-SUM GAMES, G2 ($W = 5$ RUPEES) AND G4 ($W = 10$ RUPEES)

		Column					
		S1	S2	S3	S4	S5	S6
Row	S1	W,L	L,W	L,W	L,W	L,W	W,L
	S2	L,W	L,W	W,L	W,L	W,L	W,L
	S3	L,W	W,L	L,W	L,W	W,L	L,W
	S4	L,W	W,L	W,L	L,W	L,W	L,W
	S5	L,W	W,L	L,W	W,L	L,W	W,L
	S6	L,W	L,W	W,L	L,W	W,L	W,L

for row and column players (though we restrict the total experience weight $N(0)$ to be the same for both types of players). Tables IIIa–b show the MLE parameter estimates of the models, and χ^2 tests of the belief and reinforcement restrictions (along with p -values and degrees of freedom). We report only the one-segment results because the two-segment results do not improve much and offer no special insights.

Tables IIIa–b shows that for one-segment models, belief-based models and choice reinforcement restrictions are weakly and strongly rejected by χ^2 tests, respectively, in the calibration phase. In the validation phase, the reinforcement model is worst. The belief model is better than EWA in the four-strategy games G1 and G3, and worse in the six-strategy games G2 and G4. These differences are not large, however, and seem to be due to an idiosyncrasy in game G1.³¹

Tables IIIa–b report parameter estimates and jackknifed standard errors. The initial conditions $\hat{A}^j(0)$ are encouragingly similar in pairs of low- and high-stakes games (G1–G3 and G2–G4), and put low initial attraction on the dominated

³¹In game G1, EWA overfits the first 28 periods because it detects some upward trend in strategies S1 and S3, and downward trend in S2. These trends are reversed in the last 12 periods so EWA predicts poorly there. The belief model estimates differences in initial expected payoffs but has a huge value of $\hat{N}(0) = 300$, so it doesn't predict much movement at all.

TABLE IIIa
4 × 4 CONSTANT-SUM GAMES (G1 AND G3) (M = 560)

Parameters	EWA		Choice Reinforcement		Belief-Based Models	
	G1	G3	G1	G3	G1	G3
Initial Values						
ROW						
$A^1(0) [N^1(0)]$	1.320 (0.059)	5.237 (0.222)	0.000 (0.012)	10.000 (0.000)	1.780 [107.370] (0.029) [1.758]	3.914 [22.703] (0.073) [0.423]
$A^2(0) [N^2(0)]$	1.145 (0.092)	1.172 (0.265)	0.000 (0.000)	0.000 (0.000)	1.422 [107.290] (0.070) [4.687]	1.739 [25.214] (0.290) [1.403]
$A^3(0) [N^3(0)]$	1.913 (0.066)	7.187 (0.300)	5.000 (0.000)	10.000 (0.000)	2.262 [85.335] (0.057) [4.187]	4.927 [10.084] (0.149) [1.682]
$A^4(0) [N^4(0)]$	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.948 [0.000] (0.047) [0.000]	1.159 [0.000] (0.193) [0.000]
COLUMN						
$A^1(0) [N^1(0)]$	2.681 (0.091)	6.790 (0.640)	5.000 (0.000)	10.000 (0.000)	3.385 [96.887] (0.049) [2.923]	6.674 [19.294] (0.092) [0.536]
$A^2(0) [N^2(0)]$	2.583 (0.102)	6.859 (0.580)	5.000 (0.000)	10.000 (0.000)	3.079 [87.932] (0.043) [2.582]	6.423 [17.962] (0.089) [0.935]
$A^3(0) [N^3(0)]$	2.345 (0.092)	5.359 (0.735)	5.000 (0.000)	8.408 (0.000)	2.896 [115.280] (0.040) [2.563]	5.711 [20.744] (0.136) [0.514]
$A^4(0) [N^4(0)]$	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	1.281 [0.001] (0.029) [0.000]	2.384 [0.000] (0.059) [0.000]
$N(0)$	19.630 (0.065)	18.391 (0.713)	1.000 (0.000)	1.000 (0.000)	300.000 [0.000]	58.000 [0.001]
Decay Parameters						
ϕ	1.040 (0.010)	1.005 (0.009)	1.012 (0.006)	0.978 (0.008)	1.000 (0.001)	1.000 (0.005)
ρ	0.961 (0.014)	0.946 (0.011)	0.000 (0.000)	0.000 (0.000)	1.000 (0.001)	1.000 (0.005)
Imagination factor						
δ	0.000 (0.035)	0.730 (0.103)	0.000 (0.000)	0.000 (0.000)	1.000 (0.000)	1.000 (0.000)
Payoff sensitivity						
λ	0.508 (0.048)	0.182 (0.015)	0.053 (0.004)	0.033 (0.002)	1.168 (0.067)	0.459 (0.063)
Log-likelihood						
–LL	653.072	678.496	681.968	710.136	680.232	–681.632
χ^2	—	—	57.792	63.280	54.320	5.328
(p-value, dof)			(0.000, 3)	(0.000, 3)	(0.000, 2)	(0.067, 2)

strategies. The initial experience weight $\hat{N}(0)$ varies between about 10–20 and is close to its steady-state value of $1/(1 - \rho)$. This means that initial attractions are weighted quite heavily, which is reasonable given the slow convergence in these 40-period games. The decay parameters $\hat{\phi}$ and $\hat{\rho}$ are close to one, with $\hat{\phi} > \hat{\rho}$. These numbers imply that attractions grow only slightly on average. By forcing $\rho = 0$, in contrast, the reinforcement model forces attractions to grow and ‘locks in’ initial behavior too quickly. Finally, $\hat{\delta}$ is between .4 and .7 and

TABLE IIIb
 6×6 CONSTANT-SUM GAMES (G2 AND G4) ($M = 560$)

Parameters	EWA		Choice Reinforcement		Belief-Based Models	
	G2	G4	G2	G4	G2	G4
Initial Values						
ROW						
$A^1(0) [N^1(0)]$	2.996 (0.068)	9.491 (0.145)	5.000 (0.000)	10.000 (0.000)	2.309 [41.566] (0.045) [0.810]	5.335 [16.005] (0.053) [0.283]
$A^2(0) [N^2(0)]$	2.434 (0.043)	7.964 (0.090)	5.000 (0.000)	10.000 (0.000)	2.077 [11.048] (0.036) [0.886]	4.665 [0.001] (0.053) [0.000]
$A^3(0) [N^3(0)]$	0.000 (0.015)	0.027 (0.081)	0.000 (0.000)	0.000 (0.000)	1.122 [10.227] (0.086) [1.111]	0.734 [1.090] (0.050) [0.083]
$A^4(0) [N^4(0)]$	0.000 (0.036)	0.004 (0.001)	0.000 (0.000)	0.000 (0.000)	1.182 [18.018] (0.080) [0.850]	0.364 [10.704] (0.026) [0.298]
$A^5(0) [N^5(0)]$	1.338 (0.034)	6.105 (0.060)	0.229 (0.000)	8.501 (0.000)	1.615 [9.141] (0.030) [0.993]	3.568 [2.200] (0.071) [0.115]
$A^6(0) [N^6(0)]$	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	1.076 [0.000] (0.038) [0.000]	0.097 [0.000] (0.040) [0.000]
COLUMN						
$A^1(0) [N^1(0)]$	4.998 (0.020)	8.733 (0.198)	5.000 (0.000)	10.000 (0.000)	3.595 [25.296] (0.044) [0.797]	7.769 [6.692] (0.072) [0.253]
$A^2(0) [N^2(0)]$	4.047 (0.046)	6.306 (0.080)	4.852 (0.000)	9.996 (0.000)	3.218 [32.627] (0.034) [1.151]	6.383 [10.301] (0.131) [0.496]
$A^3(0) [N^3(0)]$	2.201 (0.050)	1.572 (0.049)	0.001 (0.000)	0.001 (0.000)	2.444 [16.539] (0.068) [0.922]	3.954 [5.141] (0.102) [0.352]
$A^4(0) [N^4(0)]$	3.852 (0.043)	6.565 (0.070)	4.973 (0.000)	9.832 (0.000)	3.068 [13.391] (0.053) [0.914]	6.557 [5.684] (0.150) [0.462]
$A^5(0) [N^5(0)]$	1.737 (0.047)	1.851 (0.267)	0.000 (0.000)	0.001 (0.000)	2.269 [2.147] (0.085) [0.389]	4.135 [0.028] (0.213) [0.000]
$A^6(0) [N^6(0)]$	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	1.663 [0.000] (0.034) [0.000]	3.608 [2.155] (0.131) [0.373]
$N(0)$	15.276 (0.009)	9.937 (0.017)	1.000 (0.000)	1.000 (0.000)	90.000 (0.012)	30.000 (0.448)
Decay Parameters						
ϕ	0.986 (0.005)	0.991 (0.011)	0.960 (0.005)	0.962 (0.005)	0.989 (0.004)	1.000 (0.002)
ρ	0.935 (0.006)	0.926 (0.024)	0.000 (0.000)	0.000 (0.000)	0.989 (0.004)	1.000 (0.002)
Imagination factor						
δ	0.413 (0.082)	0.547 (0.054)	0.000 (0.000)	0.000 (0.000)	1.000 (0.000)	1.000 (0.000)
Payoff sensitivity						
λ	0.646 (0.030)	0.218 (0.019)	0.098 (0.005)	0.046 (0.002)	1.812 (0.123)	1.501 (0.019)
−LL	790.608	855.288	853.608	901.600	797.720	842.968
χ^2	—	—	126.000	92.264	14.224	21.280
(p -value, dof)			(0.000,3)	(0.000,3)	(0.001,2)	(0.000,2)

significantly different from both zero and one, except in game G1 where it is estimated to be zero.

Notice how the EWA estimates reflect a hybridization of elements of reinforcement and belief learning. First, the initial EWA attractions place much less relative weight on the dominated strategies (the highest-numbered strategies 4 or 6) than the corresponding expected payoffs in belief models. In the belief model the gap between the initial expected payoffs of strategy 2 (the dominant strategies) and the dominated strategies cannot be too large because the strategies are only weakly dominated. For example, in game G1 the estimated EWA attractions on row strategies 2 and 4 are 1.14 and .00, while the corresponding estimated expected payoffs are 1.42 and .95, a gap less than half as large. Thus, EWA exploits the flexibility of initial attractions from reinforcement models to squash the likelihood of playing weakly dominated strategies further down than belief models can. Second, EWA borrows the belief-model property that attractions do not grow much, since $\hat{\phi}$ and $\hat{\rho}$ are very close. Third, the estimates of δ around .5 (except G1) reflect both the law of simulated effect ($\delta > 0$) and stronger effects of actual payoffs than foregone payoffs ($\delta < 1$).

Our conclusions about the relative performance of reinforcement and belief models are different from the findings of Mookerjee and Sopher (1997). Their analysis differed in a couple of important ways.³² They allowed cross-effects so that the attraction $A_i^j(t)$ can affect the probability of playing other strategies differentially, which is more general than our approach and seems to favor reinforcement in their estimation. Their version of reinforcement also used 'average achieved earnings' rather than (weighted) cumulative earnings. The fact that $\hat{\phi}$ was very close to $\hat{\rho}$ in the EWA estimates indicates that MS took the right tack by using average earnings rather than cumulative earnings, because the cumulative-earnings assumption predicts a sharpness of convergence that is not evident in the data. In addition, their version of the belief model (which uses time-averaged expected payoffs) does not begin with an initial pre-game experience count expressing prior beliefs. Our estimates of $N(0)$ range from 30 to 300, which means that the belief model does best when it starts with a strong prior and updates very little. Thus, the difference between our results and theirs is either due to their use of cross-effects of attractions on other-strategy probabilities, to the fact that they use averaged reinforcements rather than cumulative ones (which improves reinforcement relative to our method), or to the fact that they did not allow strong prior beliefs (which handicaps the belief model relative to our method).

Finally, notice that these constant-sum games simply do not distinguish models empirically very well (as shown also by Erev and Roth (1997)), and the

³²Their analysis used logit estimation of strategy choices to judge whether choices depended more strongly on a player's own average past earnings (a kind of choice reinforcement) or on expected earnings based on opponent's past history (fictitious play). They also compared models based on the entire previous history, weighting all observations equally, with models based on a five-period moving average. (The entire-history models fit better.)

TABLE IV
THE MEDIAN EFFORT GAME

		7	6	5	Median { X_i } 4	3	2	1
	7	1.30	1.15	0.90	0.55	0.10	-0.45	-1.10
	6	1.25	1.20	1.05	0.80	0.45	0.00	-0.55
	5	1.10	1.15	1.10	0.95	0.70	0.35	-0.10
X_i	4	0.85	1.00	1.05	1.00	0.85	0.60	0.25
	3	0.50	0.75	0.90	0.95	0.90	0.75	0.50
	2	0.05	0.40	0.65	0.80	0.85	0.80	0.65
	1	-0.50	-0.05	0.30	0.55	0.70	0.75	0.70

pseudo- R^2 's are low. Coordination games, in which players converge quickly, may prove to be a better domain in which to distinguish theories.

5.3. Median-action Games

In median-action order-statistic coordination games, the group payoff depends on the median of all players' actions.³³ Table IV shows the payoff matrix used by Van Huyck, Battalio, and Beil (VHBB (1990)), whose data we use.

Players earn a payoff that increases in the median, and decreases in the (squared) deviation from the median. The median-action games capture social situations in which conformity pressures induce people to behave like others do, but everyone prefers the group to choose a high median.

We estimate EWA, choice reinforcement, and belief models using sessions 1–6 from VHBB (game Γ). In their experiments groups of nine subjects each play ten periods together, so the sample has 54 subjects.³⁴ In each round players choose an integer from 1 to 7, inclusive. At the end of each round the median is announced (but not the full distribution of choices) and players compute their payoffs. Since the groups are large (and players do not know what the median would be if their own choice was different), we assume that players form beliefs over the median of all players, ignoring their own influence on the median and treating the group as a composite single player.

Figure 1a shows the actual frequencies across the six sessions, pooled together. Initial choices are concentrated around 4–5, with a dip at 6 and small spikes at 3 and 7. Later choices move sharply toward the initial medians, which were always 4 or 5. A striking feature, which is masked by pooling sessions, is that the 10th-round median in *every* session was equal to the first-round median.

³³Camerer and Ho (1999) also report estimates from 'weak-link' coordination games in which the group payoff depends on the minimum. The parameter estimates are similar to those reported here—for example \hat{p} is .65 and $N(0)$ is around two.

³⁴They compared two treatments using nine-person groups and 'dual market' (dm) treatments in which players play with a nine-person group and a twenty-seven person group simultaneously. There is no apparent or statistically-significant difference between these treatments so we pool them together.




FIGURE 1A.—Actual frequencies, median-effort games.

In three sessions the median began at 4 and stayed there; in the other three sessions the median began at 5 and stayed there. Figure 1a shows three key features of the data that any learning model should account for: The initial spikes at 4–5 roughly double in size (as players converge fully toward them); disequilibrium choices of 3 and 7 are quickly extinguished after the first period; and there is a “dip” in initial choices at 6 (fewer players choose 6 than choose neighboring strategies 5 or 7).

From a learning point of view, median-action games are interesting because the penalty for deviating is fairly small if the players are close to equilibrium. Yet sharp convergence occurs within a couple of periods. Learning models that assume choices are reinforced must explain why players move quickly to equilibrium despite the large reinforcement if they are close to equilibrium and the small extra gain from moving precisely to equilibrium. The EWA model can account for this swift convergence if δ is close to one, incorporating the best-responsiveness inherent in belief learning.




FIGURE 1B.—Predictive errors of EWA model.

Table V shows estimation results for the median-action games. First we focus on one-segment results. EWA fits better than the reinforcement model ($\chi^2 = 64.8$) and much better than the belief model ($\chi^2 = 258.9$). The sources of EWA's improved fit are evident from looking at the data and plots of prediction errors.

Figure 1a shows that in the actual data, there are two large spikes in initial choices at 4–5, smaller spikes (about 15% of the observations) at 3 and 7, and few observations at 6. The estimated EWA initial attractions basically reflect this pattern in the data. The accuracy of the reflection can be judged from Figure 1b, an EWA error plot. This figure shows the difference between (MLE) predicted frequencies of the EWA model and the actual frequencies. The largest error is that EWA underpredicts the frequency of choices of 3 by about .06; predictions of 6 and 7 are too high by .03 and .01.

Reinforcement and belief learning cannot fit the initial conditions as well as EWA, but for different reasons. Reinforcement learning underpredicts the




FIGURE 1c.—Predictive errors of choice reinforcement model.

actual initial frequencies of 3 and 7 by about .08. Players who chose strategy 7 in the first period quickly switch to lower numbers in period 2, as Figure 1a shows. (The same is true for players who chose strategy 3, but this cannot be seen in Figure 1a.) Reinforcement learning cannot predict how quickly this convergence occurs. Since the initial medians are 4–5, choices of 3 or 7 earn between \$.55 and \$.95, while ex-post best responses earn \$1.00 to \$1.10. Since the initial choices are positively reinforced, reinforcement learning cannot explain why subjects will abandon these strategies so quickly and switch in the direction of the observed median. (EWA explains convergence with a high estimate of $\hat{\delta} = .85$.) Since choice reinforcement does not adjust chosen strategies quickly enough, to maximize overall fit it deliberately misestimates the first-period choices by assuming the initial frequencies are close to frequencies in later periods, thereby underpredicting choices of 3 and 7 (and overpredicting 4–5).

Figure 1c shows that the belief model underpredicts 3 and 7 also, but for a different reason. In the belief-based framework it is hard to explain why players




FIGURE 1D.—Predictive errors of belief-based model.

would play 6 less than they play 5 or 7. The problem is that initial beliefs that give a high expected payoff to 4–5 (expecting a median of 4–5) also give an expected payoff to 6 that is nearly as large, and larger than the expected payoff to 7. Beliefs that give a large expected payoff to 7, because there is a high probability that the median will be 7, will also give a high expected payoff to 6. Thus, it is difficult to find a single set of beliefs that can explain the spikes at 4–5 and 7, without also predicting a spike at 6. As a result, Table V shows that the one-segment model generates initial expected payoffs that are higher for 6 (\$.78) than for 3 or 7 (\$.71 and \$.60), so it overpredicts 6 and underpredicts 3 and 7 (and also overpredicts 5).

Adding a second segment of players improves the belief-model fit dramatically. As Table V shows, the log likelihood improves a lot (the χ^2 statistics for the two-segment results compare one- and two-segment fits within each model). The two belief-model segments correspond naturally to a large (78%) segment with high expected payoffs for 4–5 generated by high initial beliefs in 4–5, and a

TABLE V
MEDIAN EFFORT GAMES, ($M = 378$)

Parameters	EWA		Choice Reinforcement		Belief-Based Models	
	1-Segment	2-Segment	1-Segment	2-Segment	1-Segment	2-Segment
Initial values/Segment Size						
$A^1(0)$ [$N^1(0)$]	0.000 (0.000)	0.659 (0.012)	0.000 (0.000)	0.799 (0.008)	0.168 [0.075] (0.007 [0.018])	0.785 (0.007)
$A^2(0)$ [$N^2(0)$]	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.457 [0.096] (0.007 [0.010])	0.215 (0.007)
$A^3(0)$ [$N^3(0)$]	0.000 (0.000)	0.095 (0.028)	0.034 (0.175)	0.000 (0.028)	0.725 [0.037] (0.005 [0.012])	0.050 [0.000] (0.000 [0.000])
$A^4(0)$ [$N^4(0)$]	0.337 (0.145)	0.000 (0.000)	0.240 (0.175)	0.877 (0.028)	0.713 [0.059] (0.005 [0.012])	0.500 [0.000] (0.000 [0.000])
$A^5(0)$ [$N^5(0)$]	0.561 (0.143)	2.387 (0.005)	2.400 (0.328)	2.400 (0.000)	0.002 [0.014] (0.003 [0.019])	0.000 [0.000] (0.000 [0.000])
$A^6(0)$ [$N^6(0)$]	0.364 (0.148)	2.313 (0.007)	1.797 (0.255)	1.293 (0.049)	0.836 [0.073] (0.003 [0.019])	0.961 [1.124] (0.003 [0.063])
$A^7(0)$ [$N^7(0)$]	0.431 (0.146)	2.348 (0.008)	2.000 (0.111)	0.986 (0.041)	0.859 [0.093] (0.003 [0.022])	0.928 [0.490] (0.003 [0.069])
$N(0)$	0.647 (0.059)	0.205 (0.027)	0.149 (0.034)	1.000 (0.000)	0.787 [0.123] (0.006 [0.028])	1.250 [0.00] (0.006 [0.000])
Decay Parameters					0.604 [0.170] (0.008 [0.038])	1.300 [0.358] (0.000 [0.051])
ϕ	0.800 (0.018)	0.746 (0.030)	0.800 (0.020)	0.979 (0.011)	0.642 (0.150)	0.358 (0.057)
ρ	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.494 (0.030)	0.999 (0.000)
Imagination factor					0.738 (0.054)	0.999 (0.000)
δ	0.853 (0.005)	0.947 (0.002)	0.497 (0.014)	0.000 (0.000)	0.564 [0.019] (0.008 [0.002])	1.000 (0.000)
Payoff sensitivity					1.000 (0.000)	1.000 (0.000)
λ	6.827 (0.251)	17.987 (0.869)	2.969 (0.127)	1.002 (0.032)	16.726 (0.366)	43.600 (0.837)
- LL	309.304	290.255	341.697	331.348	438.748	379.236
χ^2	—	38.098	64.785	20.698	258.888	119.024
(p -value, dof)		(0.000,12)	(0.000,3)	(0.014,9)	(0.000,2)	(0.000,10)

smaller (22%) segment with belief only in 7, which generates the highest expected payoff for 7. While testing the restriction that the second segment does not improve fit rejects strongly ($\chi^2 = 119.0$), the two-segment belief model still does not fit as well as the one- or two-segment EWA model.

Finally, Table I shows that EWA performs much better than the other models in out-of-sample validation, using either the LL or MSD measure. The error plots show that reinforcement overpredicts the number of people choosing 6 in periods eight and ten, because these choices are positively reinforced and should persist; EWA does a little better by predicting that these late-period choices move toward 4–5. The belief model predicts a growth in strategies 3 and 6 over time, because the expected payoffs of those strategies will be almost as high as the expected payoffs from strategies 4–5 as equilibration occurs. Since those strategies are rarely used in later periods, the belief model badly overpredicts them in those periods.

Besides fitting initial conditions, a good learning model must explain why convergence in the first couple of periods is fast and sharp. EWA does this by estimating a large value of δ (.85) and $\hat{\phi}$ much larger than $\hat{\rho}$, which allows attractions to grow rapidly so that choice probabilities move toward zero and one swiftly. The low value of $N(0)$, .65, also allows players to learn quickly from payoff reinforcement relative to initial attractions.

The estimates show how EWA mixes and matches the best features of belief and reinforcement learning: It allows near-best response (δ close to one) as in belief models, which explains why players choosing near-equilibrium strategies move quickly toward equilibrium. But as in reinforcement, it can allow arbitrary initial attractions, which explains the relative paucity of choices of 6 in the first period, and allows attractions to grow (because $\rho = 0$) to explain the sharpness of convergence. As a result, the EWA errors (Figure 1b) are generally much smaller than those in reinforcement (Figure 1c) and belief learning (Figure 1d).

The results shown in the error plots are for one-segment models. Adding a second segment improves fits significantly for all three models. In EWA, the larger segment (with proportion 66%) has an estimate $\hat{\delta} = .95$, very close to the belief restriction of one, while the smaller second segment has $\hat{\delta} = .50$. This corresponds to a segment of people with belief-type equal weighting of actual and foregone payoffs, and another segment who weight actual payoffs twice as heavily. Notice that these two segments do not particularly correspond to one segment of reinforcement learners and another segment of belief learners, so EWA is not simply capturing a mixture of these two special cases.

In reinforcement, the larger segment (80%) has parameter values which are similar to those in the single segment, except the estimates of initial attractions for 3 and 7 are zero. The smaller second segment (20%) is the opposite—strategies 3 and 7 have the largest possible initial attractions and all the others are close to zero—except that $\hat{\phi} = 0$.³⁵ This means the two-segment structure is

³⁵The estimate of zero for ϕ is the full-sample MLE estimate. The jackknifed standard error of .235 means that in many jackknife samples ϕ is estimated to be positive. Indeed, the mean of the jackknife estimates is .18, but this does not substantially affect the point we make in the text.

trying to solve the problem of explaining first-period choices of 3 and 7 that are quickly extinguished by creating a second segment of players who choose only 3 or 7 initially, then immediately decay their initial attraction. But adding this segment does not improve log likelihood much and the two-segment reinforcement model still fits worse than the one-segment EWA model.

The two-segment belief model improves fit substantially but still does not capture initial attractions flexibly enough (compared to EWA). We think the problem is that the belief model requires initial behavior to be consistent with prior beliefs *and* requires beliefs to be updated using weighted fictitious play. The latter assumption requires $\delta = 1$ and $\phi = \rho$. In games like the median-action game, the $\delta = 1$ assumption is a reasonable approximation but $\phi = \rho$ does not allow sharp enough convergence.³⁶ More importantly, forcing initial attractions to spring from expected payoffs does not explain behavior of players who use certain decision rules. For example, a player who randomizes among different selection principles will not necessarily choose according to expected payoffs given prior.

5.4. Dominance-solvable *p*-Beauty Contest Games

In a *p*-beauty contest game, n players simultaneously choose numbers x_i in some interval, say $[0,100]$. The average of their numbers $\bar{x} = \sum_i x_i / n$ is computed, which establishes a target number, τ , equal to $p \cdot \bar{x}$. The player whose number is closest to the target wins a fixed prize $n \cdot \pi$ (and ties are broken randomly³⁷).

P-beauty contest games were first studied experimentally by Nagel (1995) and extended by Ho, Camerer, and Weigelt (1998) and Duffy and Nagel (in press). 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.

There are two behavioral regularities in beauty contest games (see Nagel (1999) for a review). First, initial choices are widely dispersed and centered somewhere between the interval midpoint and the equilibrium. This basic result has been replicated with students on three continents and with several samples of sophisticated adults, including economics Ph.D.'s and a sample of CEOs and

³⁶The fact that $\hat{\rho} = 0$ in EWA (and never varies across the jackknife runs) also suggests that adding more segments to the belief model will not improve fit substantially compared to EWA models with the same number of segments, because the belief models are always constrained to have $\rho = \phi$.

³⁷Formally, $\pi(x_i, x_{-i}) = n \cdot \pi \cdot I(x_i, \arg\min_{x_j} |x_j - \tau|) / \sum_i I(x_i, \arg\min_{x_j} |x_j - \tau|)$ where $I(x, y)$ is the indicator function that equals one if $x = y$ and 0 otherwise.

corporate presidents (see Camerer (1997)). Second, when the game is repeated, numbers gradually converge toward the equilibrium.

Explaining beauty contest convergence is a challenge for adaptive learning models. Standard choice reinforcement is likely to converge far too slowly, because only one player wins each period and the losers get no reinforcement. Belief models with low values of ϕ , which update beliefs very quickly, may track the learning process reasonably well, but earlier work suggests Cournot dynamics do not converge fast enough either (Ho, Camerer, and Weigelt (1998)).

The three models were estimated on a subsample of data collected by Ho, Camerer, and Weigelt (1998). 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 time, about \$2.33 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 'high experience' condition (the game is the second one subjects play, following a game with a value of $p > 1$) and the 'low experience' condition (the game is the first they play). The experience conditions were pooled to create enough data to get reliable estimates.

Several design choices were necessary 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 quickly use too many degrees of freedom estimating the initial attractions. Rather than imposing too many structural requirements on the distribution of $A^j(0)$, we assumed initial attractions were equal in ten-number intervals $[0, 9]$, $[10, 19]$, etc.³⁸

To implement EWA we assumed subjects knew the winning number, $w = \arg\min_{x_j} [|x_j - \tau|]$, and neglected the effect of their own choice on the target number.³⁹ Define the distance between the winning number and the target number as $d = |\tau - w|$. All subjects reinforced numbers in the intervals $(\tau - d, \tau + d)$ by δ times the prize, and numbers in the intervals $[0, \tau - d)$ and $(\tau + d, 100]$ received no reinforcement. Winners reinforced the boundary number they chose, either $\tau - d$ or $\tau + d$, by the prize divided by the number of winners, and reinforced the other boundary number by δ times the prize divided by the

³⁸In our working paper we assumed the distribution of the values of $A^j(0)$ came from a beta distribution but the basic results were not much different. We also tried fitting asymmetric triangular distributions, in which $A^{100}(0) = 0$, $A^{50}(0) = c$, $A^0(0) = b$, and $A^j(0)$ was piecewise linear between 0 and 50, and 50 and 100, with slopes $(c - b)/50$ and $-c/50$, respectively, and tried normal distributions but the basic results were unchanged.

³⁹Since subjects were not told the winning number (unless their number won), the fact that we must assume they do to estimate the model could be considered a handicap for the EWA and belief-based models, and a possible advantage for choice reinforcement, which does not require this assumption.

number of winners. Losers reinforced both boundary numbers $\tau - d$ and $\tau + d$ by δ times the prize, divided by the number of winners plus one.

Implementing the belief model is not straightforward because subjects were told only the target number, and whether they won, so they do not have enough information to form beliefs about what other subjects will do, and use these updated beliefs to calculate expected payoffs. Reinforcing numbers in some intervals, as in the EWA updating, will not necessarily correspond to belief learning based in information about all others' numbers (which they do not know anyway). As a result, we estimate a restricted form of EWA with belief-type parameters by setting $\delta = 1$, $\phi = \rho$, estimating initial belief counts in the ten-number intervals, and taking initial expected payoffs to be normalized belief counts multiplied by the prize. Numbers in the winning interval $(\tau - d, \tau + d)$ are reinforced by one times the prize. This corresponds to a special kind of belief learning in which players are learning what the target number will be and best-responding given their beliefs.

Table I reports overall results. Generally the fit is not very impressive; ρ^2 values are only around 7%. In the calibration sample, EWA is slightly better than reinforcement, which is better than the belief model. Out of sample, the belief model and EWA model are about equally good (and reinforcement is clearly worst); the belief model is slightly better on MSD and much worse in log likelihood than EWA.

Table VI reports results of parameter estimates.

The EWA model fits the data as best it can in an odd way: It assumes there is a general tendency to pick lower numbers, which grows stronger over time. This can be seen in the initial attractions, which are largest for the lowest number intervals,⁴⁰ even though the first-period choices are clustered around 40–49 (i.e., attraction category $A^5(0)$). Then the model assumes these initial attractions 'inflate' over time ($\hat{\phi} = 1.33$). The model is not capturing learning from experience well because lagged attractions are weighted heavily compared to payoff reinforcement ($N(0)$ is 16.82), and the estimate $\hat{\delta}$ is small (.23).

Choice reinforcement uses the same ingredients—high initial attractions for lower numbers, inflated by $\hat{\phi} = 1.38$ —but fits substantially worse because $N(0)$ is forced to be one and there is little reinforcement from direct payoffs (since most players lose and get nothing). The belief model, in contrast, fits best by assuming initial expected payoffs are highest for choices in the interval [40, 49], responding to payoff experience strongly (δ is fixed at one), and decaying attractions fairly quickly ($N(0) = 1.67$ and $\hat{\phi} = .40$).

The two-segment analysis of EWA improves calibration substantially, compared to the one-segment model, and improves on the validation log-likelihood modestly. The two-segment reinforcement and belief models add very little to fit, especially in validation.

⁴⁰The exception is that attractions are high for the interval [90,100]. This is to account for the occasional outlying choices of 100, which are discussed at length in Ho, Camerer, and Weigelt (1998).

TABLE VI
P-BEAUTY CONTESTS ($M = 1372$)

Parameters	EWA	Choice Reinforcement	Belief-Based Models
Initial values			
$A^1(0) [N^1(0)]$	3.348 (0.002)	3.500 (0.000)	0.000 [0.000] (0.000) [0.000]
$A^2(0) [N^2(0)]$	3.311 (0.002)	3.240 (0.002)	0.000 [0.000] (0.000) [0.000]
$A^3(0) [N^3(0)]$	3.301 (0.002)	3.160 (0.002)	0.000 [0.000] (0.000) [0.000]
$A^4(0) [N^4(0)]$	3.269 (0.001)	2.891 (0.002)	0.000 [0.000] (0.000) [0.000]
$A^5(0) [N^5(0)]$	3.227 (0.001)	2.442 (0.002)	0.350 [0.167] (0.000) [0.000]
$A^6(0) [N^6(0)]$	3.180 (0.001)	2.006 (0.001)	0.000 [0.000] (0.000) [0.000]
$A^7(0) [N^7(0)]$	3.052 (0.003)	0.591 (0.001)	0.000 [0.000] (0.000) [0.000]
$A^8(0) [N^8(0)]$	2.912 (0.006)	0.000 (0.000)	0.000 [0.000] (0.000) [0.000]
$A^9(0) [N^9(0)]$	2.871 (0.005)	0.000 (0.000)	0.000 [0.000] (0.000) [0.000]
$A^{10}(0) [N^{10}(0)]$	3.060 (0.004)	0.700 (0.001)	0.000 [0.000] (0.000) [0.000]
$N(0)$	16.815 (0.000)	1.000 (0.000)	1.672 [0.002]
Decay Parameters			
ϕ	1.330 (0.004)	1.375 (0.002)	0.402 (0.001)
ρ	0.941 (0.000)	0.000 (0.000)	0.402 (0.001)
Imagination factor			
δ	0.232 (0.013)	0.000 (0.000)	1.000 (0.000)
Payoff sensitivity			
λ	2.579 (0.002)	0.223 (0.002)	0.942 (0.001)
Log-likelihood			
–LL	5878.197	5910.988	6083.036
χ^2		65.582	409.679
(p -value, dof)	—	(0.000,3)	(0.000,2)

The two EWA segments that emerge (not reported in Table VI) are interesting. The larger segment (66%) is very much like the one-segment EWA estimate: Estimated initial attractions increase for smaller-number intervals, $\hat{\phi}$ is 1.61, $\hat{\delta}$ is zero, and the experience weight $N(0)$ is 16.83. The smaller segment (34%) is remarkably like the one-segment belief model estimate: Initial attractions are highest for choices in the middle interval [50, 59], $\hat{\phi}$ and $\hat{\rho}$ are small and very close (.50 and .43), $\hat{\delta}$ is estimated to be 1.0, and $N(0) = 1.76$.

None of these models captures the nature of learning well. The reinforcement and one-segment EWA models simply pretend that the first period is like later periods and inflate initial attractions to gradually reproduce the latter-period data. Belief models converge too slowly.

We think the results suggest two possible kinds of misspecification which we plan to remedy in future research (and have been studied by others).

One possibility is that players detect trends and so, for example, use an extrapolative rule of the form “expect that $x(t+1)$ is the same multiple of $x(t)$ that $x(t)$ was of $x(t-1)$.” (Our model does not explicitly allow this since players are assumed to use stage-game strategies only.) Such a model could be adaptive, in the sense of using only past observations and neglecting information about payoffs of others, but could allow subjects to choose strategies that are best responses to behavior they have not observed before.

Another possibility is this: Since these models are adaptive, they only use information about previous payoffs (including previous foregone payoffs). Adaptive models of this sort cannot account for learning when players sophisticatedly realize that other players are learning as well (cf. Milgrom and Roberts (1991)). Our earlier work (Ho, Camerer, and Weigelt (1998)) and Stahl (1996) showed that a fraction of players seems to ‘iteratively best-respond’ in the sense that they choose numbers that are not best responses to observed history (as in weighted fictitious play), but instead choose numbers that are best responses to anticipated best-responding by others. Because the belief and reinforcement models do not have this kind of sophistication, the hybrid EWA does not either. Including sophistication in some parsimonious way may improve the fit.

We are agnostic about whether the extrapolative or sophisticated approach is generally better. There are empirical reasons to pursue both: Some experiments have shown that players behave differently when they know the payoffs of others (e.g., Partow and Schotter (1993)), indicating sophistication. Still other experiments (besides p -beauty contests) suggest sensitivity to time trends or differences in previous results (Huck, Normann, and Oechssler (1997)).

5.5. *Identification of Parameters and Model Diagnostics*

The results generally show that EWA fits better than either of the special cases, both adjusting for extra parameters and predicting out of sample. A further test for model specification is to ask whether there are regular correlations among the three added parameters, δ , $N(0)$, and ρ , and other parameters. Because the EWA model is highly nonlinear, it is possible that certain parameters covary so closely that it is difficult to identify them econometrically. (By definition, a nonidentified parameter could be dropped from the model without reducing fit.) It is easy to show algebraically that the parameters *are* identified, in the sense that for arbitrary data sets and MLE parameter estimates, no other set of parameter values would fit equally well. However, it is possible that parameters are nearly nonidentified in some data sets.

One way to check the severity of nonidentifiability is to compute correlations among parameter estimates across jackknife runs. Two parameters that cannot be disentangled will be perfectly correlated across runs. Low or modest correlations across runs indicate that parameters have detectably separate influences. By inspecting the intercorrelations of the three important added parameters we can check whether each parameter contributes to predictive power.

A good overall statistic is the mean absolute correlation of the estimates of a parameter with all the other parameters with which it might be misidentified. We exclude initial attractions and compute correlations among ϕ , δ , ρ , $N(0)$, and λ .

For δ the mean absolute correlation with the other parameters is .31, .39, and .23 across the constant-sum, median-action, and beauty-contest games. None of the correlations with a specific parameter are consistent in magnitude and sign across games. This indicates that δ is well-identified. The same statistics for $N(0)$ are .19, .22, and .32. The latter number excludes the correlation between $N(0)$ and ρ in the beauty-contest game, which is nearly one because the declining-effect constraint is binding.⁴¹ These figures show that $N(0)$ is well-identified too (except when the constraint binds). The mean absolute correlations for ρ are .48, .30, and .32 (the latter again excludes the high correlation with $N(0)$). These correlations are somewhat higher than for δ and $N(0)$, especially in constant-sum games, indicating possible identification problems. The most systematic large correlation is between ρ and ϕ , which have an average correlation of .88 in the constant-sum games (and the correlations are nearly equal in all four games). They are also correlated .50 in the median-action game and uncorrelated ($-.03$) in the beauty contest game. This pattern of correlations is a hint that the two depreciation parameters may be fundamentally related, in some games, in a way we hope to explore in further research.

The fact that the intercorrelations among estimates are modest and unsystematic (with noted exceptions) confirms that the parameters added in EWA contribute separately to its fit. We can also ask whether adding these parameters helps solve identification problems that arise in the belief and reinforcement special cases. For the reinforcement model, λ and ϕ are correlated $-.79$, $-.68$, and $.05$ in the three classes of games. The large negative correlations arise because when ϕ is lower attractions decay more rapidly, so λ must be larger to magnify small differences in attractions into large differences in choice probabilities. (The same effect does not seem to happen across runs of the beauty-contest game, where $\hat{\phi}$ is 1.38 and none of the models captures learning well.) Therefore, it is difficult to identify separate influences of the two parameters.

⁴¹When the declining-effect constraint $N(0) \leq 1/(1 - \rho)$ is binding, $N(0)$ and ρ are not identified separately. (The same is true in the belief model.) We regard this as a shred of evidence about the way in which parameters may vary systematically across classes of games (see Cheung and Friedman (1997)). It may be that dominance-solvable games in which observed strategy choices are constantly shifting location have this general property so the restriction $N(0) = 1/(1 - \rho)$ can be safely imposed.

Adding ρ and $N(0)$ in the EWA model reduces the correlations between λ and ϕ in magnitude, to .15, $-.40$, and $-.20$, eliminating any possible identification problem.

In the belief model the only apparent identification problem is between $N(0)$ and ϕ , which are correlated .20, $-.86$, and .99 in the three games. When ρ is included in the EWA model, these correlations become .23, .31, and .99, so the identification problem is partly eliminated.

Overall, there are modest identification problems in all three models. Problems in the reinforcement and belief models are largely alleviated by introducing ρ , $N(0)$, and δ , in EWA. These new parameters are fairly well identified, except for modest-to-strong correlation between ρ and ϕ in two of three games. EWA therefore solves minor identification problems in the simpler models at the expense of creating another minor one, which could be explored in further research.

6. DISCUSSION AND CONCLUSION

We proposed a general ‘experience-weighted attraction’ (EWA) learning model in which the probability of choosing a strategy is determined by its relative attraction. A strategy’s attractions are updated by weighting lagged attractions by the number of periods of ‘experience-equivalence’ they contain, adding the payoffs actually received or a fraction of the payoffs that would have been received, then normalizing by an experience weight.

The paper makes two basic contributions.

First, we show that belief learning is not fundamentally different from reinforcement learning because both are special examples of one general learning rule—EWA. By showing their common basis, EWA lays bare the essential components of reinforcement and belief learning, and shows how those components can be combined to make a better model. Comparing choice reinforcement to EWA makes it clear that reinforcement assumes players ignore foregone payoffs, and attractions either cumulate as quickly as possible (in some models) or average past reinforcements (in other models). Comparing weighted fictitious play to EWA makes it clear that belief models assume initial attractions are consistent with prior beliefs, foregone and actual payoffs are equally reinforcing, and attractions are weighted averages of past attractions and payoffs.

Indeed, as an empirical matter there is no reason to think that the clusters of parametric restrictions embodied in cumulative choice reinforcement and weighted fictitious play are the clusters most likely to arise in human behavior. For example, there is no empirical reason to think that players who ignore foregone payoffs also cumulate reinforcements, or that players who weight foregone payoffs necessarily use weighted averages of past attractions. EWA shows that there are many more clusters of parameter values than the two special kinds embodied in cumulative reinforcement and weighted fictitious play. For this reason, perhaps EWA should be studied more extensively in modern empirical work before these special cases.

Second, by estimating the more general EWA model, along with reinforcement and belief-learning restrictions, our study combines methodological strengths of earlier studies while avoiding weaknesses. All earlier studies did one or more of the following: concentrated on only one or two models, focussed on one class of games, ignored player heterogeneity, restricted the generality of models, derived parameter values using methods that do not guarantee best-fits, or did not report inferential statistics testing relative fit. Our paper had none of these limits because we compared three general models, on three classes of games, allowed some heterogeneity, derived parameter values optimally, and reported both test statistics (adjusting for free parameters three ways) and out-of-sample predictive accuracy (measured two ways).

EWA fits better than the reinforcement models in all cases, and better than belief learning in most cases, both adjusting for degrees of freedom within-sample and in out-of-sample prediction. Belief models are more accurate than reinforcement in some games, and by some measures, and less accurate in others.

The foregone payoff weight δ is estimated to be .42 (averaging across the four constant-sum games), .85 in median-action games, and .23 in beauty contests. The raw average of these numbers, .50, suggests that players generally weight foregone payoffs about half as much as actual payoffs. This result incorporates the intuitions underlying both reinforcement (actual payoffs are stronger) and belief learning (foregone payoffs matter). Put differently, players seem to obey both the law of actual effect and a corollary law of simulated effect.

In the three games, the decay parameters ϕ and ρ average 1.00 and .94, .80 and 0, and 1.33 and .94. The first two games indicate that sometimes attractions are approximately averages (as in belief models) and other times they cumulate as rapidly as possible (as in reinforcement). The value of ϕ above one in beauty contests, as discussed above, reflects a likely misspecification because the adaptive models with stage-game strategies do not allow players to extrapolate trends or have sophisticated beliefs (a shortcoming the weighted fictitious play and reinforcement models also share).

The initial experience weight $N(0)$ averages 15.80, .65, and 16.82. The large values in constant-sum and beauty-contest games imply that players learn slowly, because they give much more weight to lagged attractions than to payoffs. The low value of .65 in median-action games means players respond more strongly to payoffs, learning faster.

EWA also exploits the flexibility of initial attractions shared by reinforcement models, compared to belief models in which initial attractions must be expected payoffs based on some prior. This flexibility is particularly helpful in the coordination games.

The results show how EWA is able to ‘gene-splice’ the best features of belief and reinforcement learning while avoiding their weaknesses. For example, in the median-action games players begin with dispersed choices that seem to reflect different selection principles, and converge quickly. Explaining this pattern well requires initial attractions that are flexible and cumulate (as in reinforcement),

rather than belief-based initial attractions that are averages, but also requires players to respond strongly to foregone payoffs (as in belief learning).

The fact that parameter values vary widely across the data sets is not surprising, because all studies that have looked for differences in parameters across games have found them (e.g., Crawford (1995), Cheung and Friedman (1997), Erev and Roth (1997), Chen and Tang (1998)). Furthermore, the parameters capture different features of the data—speed of learning and sharpness of convergence. Since these features are different across games, parameter values should differ. Nonetheless, our understanding of learning will not be complete until there is a theory of how parameter values depend on game structure and experimental conditions (see Cheung and Friedman (1997) for important progress). Our estimates provide raw material for such theorizing.

6.1. *EWA Extensions*

There are many directions for future research.

Theorizing about the kinds of equilibria to which EWA learning rules converge would be extremely useful. Hart and Mas-Colell (1996) provide a clue. They study a process in which players shift probability toward strategies they wish they had played, in proportion to the difference between foregone and actual payoffs (“regret”). This process is similar to EWA learning because EWA also shifts probability toward high-regret strategies when δ is close to one. Hart and Mas-Colell prove that their process converges almost surely to correlated equilibrium for finite normal-form games. Similar results might be derived for EWA (perhaps for a restricted class of games). Interested theorists might keep in mind that mapping attractions into probabilities using the ratio form means the denominator of the updating equation vanishes, so ρ disappears and the experience weight only enters if the scale of the initial attractions is restricted. Then the model can be sensibly reduced to two parameters— δ and ϕ (which can be set to one for some theoretical purposes)—which should make theorizing easier.

An empirical direction for further research is measurement of model parameters using psychological methods. For example, if δ is interpreted as attention to foregone payoffs from unchosen alternatives, then values of δ should correlate with direct measures of attention, such as the amount of time subjects spend looking at different numbers in a payoff matrix (see Camerer et al (1993)). (In general, measuring attention to information provides a direct way to test theories that assume certain kinds of information are not used.⁴²) Or if $N(0)$ is

⁴² For example, choice reinforcement predicts that players do not use information other than their own payoff history. Experiments that vary the information subjects are given have shown this prediction is clearly wrong (Mookerjee and Sopher (1994), Van Huyck, Battalio, and Rankin (1996), Huck, Normann, Oechssler (1997)). Direct measures of attention provide a more direct test: if players look at foregone payoffs frequently, then reinforcement models have some explaining to do. Similarly, all adaptive models predict that players do not use information about others' payoffs; looking at those payoffs is evidence of sophistication.

the number of pregame ‘trials’ a player simulates that form prior beliefs, then $N(0)$ should be related to the ratio of initial response times to later-period response times.

EWA will also have to be upgraded to cope with three modelling challenges—sophistication, imperfect payoff information, and specification of strategies—before it is generally applicable.

Incorporating sophistication is important because EWA players only use information about their opponents’ past choices, ignoring information about payoffs of others. Using this information in an expanded learning rule that incorporates sophistication could help explain data like those from the beauty-contest games. Iterating sophistication might also link sophisticated-EWA to equilibrium theories like quantal-response equilibrium.

Incorporating imperfect payoff information is important because any general model should be able to explain learning in low-information environments, where players do not know everything about their own payoffs, opponents’ strategies, etc. EWA can obviously be applied in these settings by fixing $\delta = 0$ (which means EWA can apply to any environment choice to which reinforcement applies). A more general approach would use imperfect information in some other way, rather than just giving it zero weight.

Incorporating a richer specification of strategies is important because stage-game strategies are not always the most natural candidates for the strategies that players learn about. For example, players may learn by extrapolating from sequences of observations, or learn about history-dependent repeated-game strategies or a wide variety of decision rules (like minimax, Nash equilibrium, or imitation; e.g., Stahl (1997)). Once a set of richer strategies is specified, of course, EWA can still model learning about those strategies. The open question, therefore, is what rules to specify a priori, and how a model can winnow down a very large set of possible rules as quickly as humans probably do.

Adding these difficult extensions to EWA, and a theory of first-period play to supply initial attractions, might eventually create a unified way to predict how people play games in the lab and, eventually, how they play outside as well.

Div. of Social Sciences, 228–77, California Institute of Technology, Pasadena, CA 91125, U.S.A.; camerer@hss.caltech.edu

and

Marketing Dept., The Wharton School, Philadelphia, PA 19104, U.S.A.; hoteck@wharton.upenn.edu

Manuscript received January, 1997; final revision received August, 1998.

REFERENCES

- ANDERSON, S., A. DE PALMA, AND J.-F. THISSE (1992): *Discrete Choice Theory of Product Differentiation*. Cambridge: MIT Press.
- ARTHUR, B. (1991): “Designing Economic Agents That Act Like Human Agents: A Behavioral Approach to Bounded Rationality,” *American Economic Review Proceedings*, 81, 353–359.

- BATTALIO, R., L. SAMUELSON, AND J. VAN HUYCK (1997): "Risk Dominance, Payoff Dominance and Probabilistic Choice Learning," Working Paper, Department of Economics, Texas A&M University.
- BEN-AKIVA, M., AND S. LERMAN (1985): *Discrete Choice Analysis: Theory and Application to Travel Demand*. Cambridge: MIT Press.
- BEREBY-MEYER, Y., AND I. EREV (1997): "On Learning to Become a Successful Loser: A Comparison of Alternative Abstractions of Learning Processes in the Loss Domain," Working Paper, Technion-Israel Institute of Technology.
- BOYLAN, R. T., AND M. A. EL-GAMAL (1992): "Fictitious Play: A Statistical Study of Multiple Economic Experiments," *Games and Economic Behavior*, 5, 205-222.
- BORGERS, T., AND R. SARIN (1996): "Naive Reinforcement Learning With Endogenous Aspirations," Working Paper, Texas A&M University.
- BRANDTS, J., AND C. HOLT (in press): "Naive Bayesian Learning and Adjustment to Equilibrium in Signalling Games," *Journal of Economic Behavior and Organization*, forthcoming.
- BROSETA, B. (1995): "Estimation of an Adaptive Learning Model in Experimental Coordination Games: An ARCH(1) Approach," Working Paper, Department of Economics, University of Arizona.
- BROWN, G. (1951): "Iterative Solution of Games by Fictitious Play," in *Activity Analysis of Production and Allocation*. New York: John Wiley & Sons.
- BUSH, R., AND F. MOSTELLER (1955): *Stochastic Models for Learning*. New York: John Wiley & Sons.
- CAMERER, C. F. (1997): "Progress in Behavioral Game Theory," *Journal of Economic Perspectives*, 11, 167-188.
- (in progress): "Game Theory Experiments," Manuscript in progress.
- CAMERER, C. F., AND T.-H. HO (1998): "EWA Learning in Games: Probability Form, Heterogeneity, and Time Variation," *Journal of Mathematical Psychology*, 42, 305-326.
- (1999): "EWA Learning in Games: Preliminary Estimates from Weak-Link Games," in *Games and Human Behavior: Essays in Honor of Amnon Rapoport*, ed. by D. Budesu, I. Erev, and R. Zwick. New Jersey: Lawrence Erlbaum Assoc., Inc., 31-52.
- CAMERER, C. F., E. JOHNSON, S. SEN, AND T. RYMON (1993): "Cognition and Framing in Sequential Bargaining for Gains and Losses," in *Frontiers of Game Theory*, ed. by K. Binmore, A. Kirman, and P. Tani. Cambridge: MIT Press, 27-48.
- CARNAP, R. (1962): *The Logical Foundations of Probability (2nd edition)*. Chicago: University of Chicago Press.
- CHEN, H., J. W. FRIEDMAN, AND J. F. THISSE (in press): "Boundedly Rational Nash Equilibrium: A Probabilistic Choice Approach," *Games and Economic Behavior*, forthcoming.
- CHEN, Y., AND F.-F. TANG (1998): "Learning and Incentive Compatible Mechanisms for Public Goods Provision," *Journal of Political Economy*, 106, 633-662.
- CHEUNG, Y. W., AND D. FRIEDMAN (1997): "Individual Learning in Normal Form Games: Some Laboratory Results," *Games and Economic Behavior*, 19, 46-76.
- COOPER, D., S. GARVIN, AND J. KAGEL (in press): "Signaling and Adaptive Learning in an Entry Limit Pricing Game," *RAND Journal of Economics*, forthcoming.
- COURNOT, A. (1960): *Recherches sur les Principes Mathématiques de la Théorie des Richesses*. Translated into English by N. Bacon as *Researches in the Mathematical Principles of the Theory of Wealth*. London: Haffner.
- CRAWFORD, V. P. (1995): "Adaptive Dynamics in Coordination Games," *Econometrica*, 63, 103-143.
- CRAWFORD, V. P., AND B. BROSETA (1998): "What Price Coordination? Auctioning the Right to Play as a Form of Preplay Communication," *American Economic Review*, 88, 198-225.
- CROSS, J. G. (1983): "A Theory of Adaptive Economic Behavior," New York/London: Cambridge University Press.
- DAVEY, G. C. L., AND G. MATCHETT (1994): "Unconditioned Stimulus Rehearsal and the Retention and Enhancement of Differential 'Fear' Conditioning: Effects of Trait and State Anxiety," *Journal of Abnormal Psychology*, 103, 708-718.
- DUFFY, J., AND R. NAGEL (in press): "On the Robustness of Behavior in Experimental Guessing Games," *Economic Journal*, forthcoming.

- EREV, I., AND A. ROTH (1997): "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed-Strategy Equilibria," Working Paper, University of Pittsburgh.
- EYSENCK, H. J (1979): "The Conditioning Model of Neurosis," *Behavioral and Brain Sciences*, 2, 155–199.
- FRIEDMAN, D. (1991): "Evolutionary Games in Economics," *Econometrica*, 59, 637–666.
- FUDENBERG, D., AND D. K. LEVINE (1995): "Consistency and Cautious Fictitious Play," *Journal of Economic Dynamics and Control*, 19, 1065–1090.
- (1998): *Theory of Learning in Games*. Cambridge: MIT Press.
- HARLEY, C. B. (1981): "Learning the Evolutionary Stable Strategies," *Journal of Theoretical Biology*, 89, 611–633.
- HART, S., AND A. MAS-COLEL (1996): "A Simple Adaptive Procedure Leading to Correlated Equilibrium," Hebrew University, Department of Economics, December.
- HERRNSTEIN, J. R. (1970): "On the Law of Effect," *Journal of Experimental Analysis of Behavior*, 13, 342–366.
- HO, T-H., C. CAMERER, AND K. WEIGELT (1998): "Iterated Dominance and Iterated Best-response in *p*-Beauty Contests," *American Economic Review*, 88, 947–969.
- HO, T-H., AND K. WEIGELT (1996): "Task Complexity, Equilibrium Selection, and Learning: An Experimental Study," *Management Science*, 42, 659–679.
- HUCK, S., H.-T. NORMANN, AND J. OECHSSLER (1997): "Learning in Cournot Oligopoly—An Experiment," Humboldt University, Department of Economics Working Paper.
- MAJURE, W. R. (1994): "Disequilibrium Game Theory," Unpublished dissertation, Department of Economics, Massachusetts Institute of Technology.
- MCCALLISTER, P. H. (1991): "Adaptive Approaches to Stochastic Programming," *Annals of Operations Research*, 30, 45–62.
- MCKELVEY, R. D., AND T. R. PALFREY (1995): "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 10, 6–38.
- (1998): "Quantal Response Equilibria for Extensive Form Games," *Experimental Economics*, 1, 9–41.
- MILGROM, P., AND J. ROBERTS (1991): "Adaptive and Sophisticated Learning in Repeated Normal Form Games," *Games and Economic Behavior*, 3, 82–100.
- MOOKERJEE, D., AND B. SOPHER (1994): "Learning Behavior in an Experimental Matching Pennies Game," *Games and Economic Behavior*, 7, 62–91.
- (1997): "Learning and Decision Costs in Experimental Constant-sum Games," *Games and Economic Behavior*, 19, 97–132.
- NAGEL, R. (1995): "Unraveling in Guessing Games: An Experimental Study," *American Economic Review*, 85, 1313–1326.
- (1999): "A Review of Beauty Contest Games," in *Games and Human Behavior: Essays in Honor of Amnon Rapoport*, ed. by D. Budesu, I. Erev, and R. Zwick. New Jersey: Lawrence Erlbaum Assoc., Inc., 105–142.
- PARTOW, Z., AND A. SCHOTTER (1993): "Does Game Theory Predict Well for the Wrong Reasons: An Experimental Investigation," New York University C.V. Starr Center for Applied Economics RR#93–46, December.
- ROTH, A. E. (1995): "Introduction," in *Handbook of Experimental Economics*, ed. by A. E. Roth and J. H. Kagel. Princeton: Princeton University Press.
- ROTH, A., AND I. EREV (1995): "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," *Games and Economic Behavior*, 8, 164–212.
- SARIN, R. (1995): "Learning Through Reinforcement: The Cross Model," Working Paper, Department of Economics, Texas A&M University.
- SARIN, R., AND F. VAHID (1997): "Payoff Assessments Without Probabilities: Incorporating 'Similarity' Among Strategies," Working Paper, Department of Economics, Texas A&M University.
- SELTEN, R. (1991): "Evolution, Learning, and Economic Behavior," *Journal of Risk Uncertainty*, 3, 3–24.
- (1997): "Features of Experimentally Observed Bounded Rationality," Universitat Bonn Sonderforschungsbereich 303, Discussion Paper No. B-421, November.

- SELTEN, R., AND R. STOECKER (1986): "End Behavior in Sequences of Finite Prisoner's Dilemma Supergames," *Journal of Economic Behavior and Organization*, 7, 47–70.
- SMITH, A. (1981): *An Inquiry into Nature and Causes of the Wealth of Nations*. Indianapolis: Liberty Classics.
- STAHL, D. (1996): "Boundedly Rational Rule Learning in a Guessing Game," *Games and Economic Behavior*, 16, 303–330.
- (1997): "Rule Learning in Symmetric Normal-Form Games: Theory and Evidence," Working Paper, Department of Economics, University of Texas.
- TANG, F. (1996): "Anticipatory Learning in Two-person Games: An Experimental Study," Working Paper, University of Bonn Working Paper.
- THORNDIKE, E. L. (1911): *Animal Intelligence*. New York: Macmillan.
- TVERSKY, A., AND D. KAHNEMAN (1992): "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty*, 5, 297–323.
- VAN HUYCK, J., R. BATTALIO, AND R. BEIL (1990): "Strategic Uncertainty, Equilibrium Selection, and Coordination Failure in Average Opinion Games," *Quarterly Journal of Economics*, 106, 885–909.
- VAN HUYCK, J., R. BATTALIO, AND F. W. RANKIN (1996): "Selection Dynamics and Adaptive Behavior Without Much Information," Working Paper, Department of Economics, Texas A & M University.
- WEIBULL, J. W. (1997): "What Have We Learned From Evolutionary Game Theory So Far?" Stockholm School of Economics Research Institute of Industrial Economics Working Paper No. 487.