

Stated Beliefs Versus Inferred Beliefs: A Methodological Inquiry and Experimental Test

by

E. Elisabet Rutström

and

Nathaniel T. Wilcox

Abstract:

If asking subjects their beliefs during repeated game play changes the way those subjects play, using those stated beliefs to evaluate and compare theories of strategic behavior is problematic. We experimentally verify that belief elicitation can alter paths of play in a repeated asymmetric matching pennies game. In this setting, belief elicitation improves the goodness of fit of structural models of belief learning, and the prior beliefs implied by such structural models are both stronger and more realistic when beliefs are elicited than when they are not. These effects are, however, confined to the player type who sees a strong asymmetry between payoff possibilities for her two strategies in the game. We also find that “inferred beliefs” (beliefs estimated from past observed actions of opponents) can be better predictors of observed actions than the “stated beliefs” resulting from belief elicitation.

Draft: September 2008

Keywords: Stated beliefs, inferred beliefs, repeated games, experimental methods

Rutström: Department of Economics, College of Business Administration, University of Central Florida, Orlando, FL 32816; erutstrom@bus.ucf.edu. **Wilcox:** Department of Economics, University of Houston, Houston TX 77204; nwilcox@mail.uh.edu. We appreciate the comments, advice and/or help we have received from Colin Camerer, Rachel Croson, Ido Erev, Nick Feltovich, Dan Friedman, Glenn W. Harrison, Ernan Haruvy, Teck-hua Ho, Chris Murray, Yaw Nyarko, Jack Ochs, David Papell, Andrew Schotter, Dale Stahl and Peter Thompson. The usual disclaimer applies, of course. Rutström thanks the U.S. National Science Foundation for research support under grants NSF/IIS 9817518, NSF/MRI 9871019 and NSF/POWRE 9973669. Wilcox thanks the U.S. National Science Foundation for research support under grant NSF/SES 0350565. All supporting documents can be found at the ExLab Digital Library: <http://exlab.bus.ucf.edu>.

Game theory and common sense suggest that beliefs shape strategic actions. Beliefs are at the center of many game-theoretic solution concepts, and many behavioral models of game play also place large explanatory burdens on beliefs that players may construct through introspection, experience or both (Stahl and Wilson 1995; Cheung and Friedman 1997; Fudenberg and Levine 1998). Many experimenters now elicit participants' beliefs about partner play during the course of game play, using these "stated beliefs" to test hypotheses, identify models and/or test model specifications; McKelvey and Page (1990) is just one relatively early example. Manski (2002, 2004) argues forcefully that, in many cases, strong identification of models of choice under risk requires strictly exogenous measures of beliefs. Put differently, "inferred beliefs" estimated from an assumed belief updating process and observed actions of participants and their partners (e.g. Cheung and Friedman) cannot provide strong theory tests. However, stated beliefs have their own potentially serious drawback: Belief elicitation procedures may alter the very strategic actions we wish to explain or predict using stated beliefs. In such an instance, stated beliefs cannot be safely regarded as strictly exogenous, and Manski's cogent point about strong identification loses some of its methodological force.

It is tempting to view beliefs stated within an incentive-compatible mechanism for truthful revelation as a "gold standard" in the universe of potential empirical approaches to beliefs. Yet stated beliefs, like inferred beliefs, may only be estimates of underlying true or "latent" beliefs. If beliefs and strategic actions both result from substantially nonconscious cognition, agents may have to consciously estimate their own latent beliefs in order to state them. In such a case, stated beliefs need not be better predictors of behavior than inferred beliefs. In fact, we show here that inferred beliefs can predict game play better than stated beliefs, contrary to the well-known (and striking) results of Nyarko and Schotter (2002).

If belief elicitation requires players to estimate their own beliefs, it may also engage them in a more conscious, deliberative consideration of the likely play of partners. This might alter game play itself—in the direction of more belief-based play. Our main purpose here is to test the assumption that belief elicitation does not alter the strategic actions we wish to explain with stated beliefs. We reject this hypothesis. This has already been observed in games with a unique dominance-solvable equilibrium or pure strategy Nash equilibrium (Erev, Bornstein and Wallsten 1993; Croson 1999, 2000; Nelson 2003; Gächter and Renner 2006). To our knowledge, we are the first to show that this can also occur in games with a unique mixed strategy Nash equilibrium—in particular, a repeated asymmetric matching pennies game like those studied by Ochs (1995) and McKelvey, Palfrey and Weber (2000).

Asymmetric matching pennies (or AMP) games are interesting for several reasons. Nash Equilibrium often does not predict play very well in these games: Probabilistic best-response generalizations of Nash Equilibrium, such as Quantal Response Equilibrium or QRE (McKelvey and Palfrey 1995), describe play much better. More importantly for our experimental design purposes, action sequences generated by belief-based and non-belief-based learning models can differ quite significantly in suitably chosen AMP games (Erev and Roth 1998, and our appendix). Importantly, though, our findings are player-specific: It is only the players who have strongly asymmetric payoff opportunities who show all these strong belief elicitation effects. One explanation for this may be that with asymmetric payoff opportunities, relatively automatic emotional or “affective” processes attract a player to actions with relatively high payoffs. If belief elicitation prompts more conscious deliberation about the likely actions of partners, the relative impact of those affective predispositions may be dampened. We conjecture that this is what happens with belief elicitation. This suggests that belief elicitation only changes strategic

actions when relatively automatic strategic predispositions and more deliberative strategic judgments are in conflict. This may help explain why some studies (e.g. Croson) find belief elicitation effects on game play while others (e.g. Nyarko and Schotter) do not.

1. Three assumptions about beliefs and three kinds of beliefs.

When we ask subjects to state beliefs that we intend to use as predictors of subsequent behavior, we usually make three tacit assumptions. First, it is of course a hypothesis, not an established fact, that play is belief-based. Models that are free of beliefs, such as reinforcement learning models and variants of them (Erev and Roth 1998, Sarin and Vahid 2001), have been proposed and compared with belief-based models.¹ Although we mostly stay within the realm of belief-based theory in our thinking and our data analysis, the existence of well-developed theoretical alternatives to belief-based play, as well as various supportive empirical results, suggest that some subjects may not normally play games in a belief-based manner. Consequently, belief elicitation procedures could move such subjects toward belief-based thinking and play. This is an important motivation for testing the second assumption—that belief elicitation does not alter the strategy choice behavior we wish to explain with stated beliefs. Testing that assumption is our central planned purpose here.

The third assumption is that stated beliefs are better predictors of strategic actions than are inferred beliefs.² This assumption can be formulated rigorously in the following terms.

Suppose the first assumption (play is belief-based) is true, and let “latent beliefs” refer to the

¹ See Churchland (1981) for a radically skeptical view of the ultimate status of all “propositional attitudes” such as “belief” and “desire” in a completed behavioral science.

² Notwithstanding Nyarko and Schotter’s (2002) specific support for this, it may not be a general fact. Experimental studies of stated beliefs in nonstrategic settings document apparent biases (Lichtenstein, Fischhoff and Phillips 1982); while some of those studies did not use incentivized truth-telling mechanisms like a scoring rule procedure, some do and the biases, while reduced, do not disappear (e.g. Wright and Aboul-Ezz 1989). Palfrey and Wang (2007) offer some evidence that stated beliefs may not be regarded as exogenous to the action choices and may therefore not be good predictors of the latter.

theoretical belief object—that is, the true but unobserved “beliefs in the head” that actually determine a player’s actions according to a belief-based theory of strategic action. Let “inferred beliefs” instead refer to estimators of latent beliefs that are based on an assumed latent belief updating process, a stochastic model of the impact of those latent beliefs on actions, and the observed action history of a subject and her partner(s). Cheung and Friedman (1997) is a pioneering example of the inferred beliefs approach.

Cognitive processes that construct and update latent beliefs, as well as processes that combine them with payoff or value information to determine action probabilities, may not be wholly conscious ones. For many subjects, the main conscious product of an encounter with a strategic choice situation may simply be an inclination toward a particular action without much awareness of any latent beliefs. If so, asking a subject to state her latent beliefs is, in part, asking her to make inferences about the causes of her own inclinations to action: It is, in effect, a request for an estimate of her beliefs based on possibly incomplete data and whatever theoretical identifying restrictions she adopts (about parts of her own cognitive processes that she cannot directly observe) to draw such inferences. If so, appropriate language for talking about both stated and inferred beliefs is the language of estimators and their relative statistical properties: In our particular context, we focus on the estimators’ predictive content for strategic actions. Therefore, we distinguish below between latent beliefs (thought of as the theoretical “beliefs in the head” specified by some belief-based theory of strategic actions), and two classes of estimators of latent beliefs—inferred beliefs and stated beliefs.

Returning to the second assumption, we suspect that in some games, or for some players in a game, automatic affective processes and deliberative judgment processes will produce similar strategic inclinations. If belief elicitation procedures increase use of the latter processes

relative to the former, then in such cases one should expect little or no elicitation effects on game play. For other games, or players, the opposite may hold and the effects on game play from belief elicitation may be substantial. Some elicitation procedures, particularly complex ones like scoring rules, require nontrivial instruction to subjects, and interrupt the flow of subject attention and game play in a potentially significant way. Such procedures are cognitively “intrusive” and discovering the consequences of them is a necessary step toward confidently exploiting the full inferential potential of stated beliefs.

This view gets some circumstantial support from the decidedly mixed effects reported in the literature. Croson (1999, 2000), Gächter and Renner (2006), Nelson (2003), and Erev, Bornstein and Wallsten (1993) all report effects on game play from eliciting beliefs. On the other hand, Costa-Gomes and Weizsäcker (2008), Wilcox and Feltovich (2000) and Nyarko and Schotter (2002) report no significant effects. There are, of course, more prosaic explanations for the mixed results: Poor statistical power could cause this too. So we take steps (described shortly) to design our own experiment with adequate statistical power.

2. Experimental design.

2.1 Overview.

Our subjects play a repeated 2x2 asymmetric matching pennies game. We implement a control condition without any belief elicitation procedure, the “no beliefs” or NB treatment, and two experimental conditions with different belief elicitation procedures thought by us to vary in their cognitive intrusiveness. The most “intrusive” procedure uses a proper scoring rule (Aczel and Pfanzagl 1966): Subjects report a probability concerning the partner’s play, and are rewarded for its accuracy according to the scoring rule, as in Nyarko and Schotter (2002). The less

intrusive procedure simply asks subjects to state which strategy they believe their partner is most likely to play, without any reward for accuracy, as in Croson's (1999, 2000) experiments with one-shot dominance-solvable 2x2 games. We view this as a minimally intrusive procedure.

We begin with Monte Carlo power planning of our design to improve our chance of finding significant treatment effects (when present) using conservative nonparametric two-sample tests. These Monte Carlo simulations also illustrate the role played by the "degree of asymmetry" in an AMP game's payoff structure. Power to detect whether the data generating process is belief based or not appears to be higher in games with greater payoff asymmetry. This was originally suggested to us by Erev and Roth's (1998) simulation results on Och's (1995) AMP games (and our appendix illustrates it further). When testing our hypothesis that belief elicitation procedures change game play, such power is important because one potential cause for such change in play is that players estimate and use beliefs in a more deliberative way.

We first use conservative nonparametric tests (those tests examined in power planning) to establish significant treatment differences in our data. Since these are found, we view this as a warrant to proceed to a more parametric econometric analysis of treatment differences. We use an extended version of Cheung and Friedman's (1997) "Gamma-Weighted Belief" or GWB model (described below) for this analysis. This allows us to assess differences in model fit across the treatments and parametrically locate likely sources of treatment differences.

2.2 Power planning.

We use Monte Carlo simulations to choose payoffs, total periods T played by each pair of players in each treatment, and the number of player pairs M in each treatment, to provide

adequate statistical power to detect between-treatment effects (when present).³ The simulations (and our experiment) use a fixed pairing protocol where players are anonymously but permanently matched to a single partner for the duration of repeated play. This makes each row (column) player's time series of strategy choices independent of all other row (column) players' time series of strategy choices,⁴ so that nonparametric tests based on treating these time series as independent are justified. It also matches two of Nyarko and Schotter's (2002) treatments for greater comparability of results. Let "AMPX" denote an asymmetric matching pennies game, where "X" refers to the degree of asymmetry in the row player's payoffs. Our simulations searched over values of X , T and M in each of two treatments with different data-generating processes or DGPs to find values giving an acceptable probability of detecting the difference using a Wilcoxon two-sample test on summary measures of game play.

Power planning is always partially heuristic, since an unknown treatment difference must be specified in order to carry it out. We wish our design to have good power to detect the difference between some belief-based learning model and an alternative model of some sort. Our belief-based DGP is a 3-parameter version of Fudenberg and Levine's (1998) weighted fictitious play model, with parameters we estimated from Ochs' (1995) data on these games. The alternative DGP is Erev and Roth's (1998) preferred 3-parameter reinforcement learning model. We do not expect that belief elicitation procedures will produce a wholesale switch from reinforcement-based to belief-based learning; nor is our hypothesis "belief learning with belief elicitation, reinforcement learning without it." Rather, these two alternative DGPs for the simulations represent a large, theoretically important difference between a belief-based model

³ Salmon (2001) shows that power is an important and perhaps neglected aspect of learning experiments.

⁴ If different row players meet the same column partner in different periods of play (as would be true with a random rematching design) that column player carries information between those different row players and this undermines statistical independence between these different row players. This does not occur with a fixed pairing design.

and an alternative. Assuming that behavior with and without belief elicitation produces a similarly large effect (for whatever reason), the simulations tell us which designs yield a good chance of detecting that effect. The simulations suggested that $T = 36$ repetitions of an AMP19 game with $M = 40$ subject pairs in each treatment would give good power against the null of no treatment difference. For more detail on the simulations, see Rutström and Wilcox (2008), and for a summary see the appendix. Here is the AMP19 game (in the experiment, a payoff of “1” in this table is \$0.20, so that the actual dollar payoff associated with “19” in this table is \$3.80):

	<i>cl</i> (column left)	<i>cr</i> (column right)
<i>ru</i> (row up)	(19,0)	(0,1)
<i>rd</i> (row down)	(0,1)	(1,0)

One consistent finding in our experiment is that it is the row players—the players in this game with the strong payoff asymmetry—whose behavior changes in the presence of belief elicitation procedures. Anticipating this, we couch much of the following discussion of notation, belief elicitation procedures, theoretical models and econometrics in terms of row players’ situations, beliefs and decisions; generalization to the column players is straightforward.

2.3 Belief elicitation.

Our maximum expected contrast to the control NB treatment (no belief elicitation procedure) is provided by the SR (scoring rule) treatment. A scoring rule links monetary outcomes to subjects’ stated beliefs about future events in a manner that motivates expected utility maximizers to report their latent beliefs. Nyarko and Schotter used a quadratic scoring rule (and it is widely used by experimenters) so we also use it for our SR treatment.⁵ Couching this in

⁵ Strictly speaking, quadratic scoring rules are only incentive compatible for risk neutral judges. Separate elicitation of the risk attitudes in the population from which we recruited our subjects indicate that they are indeed moderately risk averse (Harrison, Johnson, McInnes and Rutström 2005). Alternative procedures that account for differences in

terms of a row players' situation, the payoff for accuracy of beliefs is specified as a quadratic function of the row player's stated belief \tilde{B}_t (stated just prior to period t play) that her column partner will play left in period t , and a dummy $l_t = 1$ or 0 indicating whether or not this subsequently occurs. In U.S. dollars, the function is $0.1[1 - (\tilde{B}_t - l_t)^2]$, so that the payoff to the stated belief \tilde{B}_t is $0.1\tilde{B}_t(2 - \tilde{B}_t)$ if $l_t = 1$ and $0.1(1 - \tilde{B}_t^2)$ if $l_t = 0$.

From a cognitive perspective, a proper scoring rule is particularly intrusive because it requires relatively fine-grained belief reports and employs a rather complex motivational scheme. Holt (1986) points out that alternative procedures may be adequate when an experimenter can make do with a monotone relationship between the resulting belief measure and latent beliefs, and that other procedures may be more transparent or natural for subjects. Therefore, we also examine a second EC (expected choice) treatment, in which subjects simply guess which strategy their partner will play in that period without any reward for accuracy—the same procedure examined by Croson (1999, 2000) in the 2x2 games she examined. This is the simplest possible belief elicitation procedure, and produces such coarse information (if any) about latent beliefs that it would not usually be useful to experimenters. Yet because of its very simplicity and the fact that no rewards are present, the EC treatment is minimally intrusive and may draw a negligible amount of conscious attention from subjects. We view it as providing a lower bound of sorts on any unintended consequences of any belief elicitation procedure. In summary, we expect that any difference in game play between the EC and NB treatments will be smaller than between the SR and NB treatments.

risk attitudes are conceivable (Andersen, Fountain, Harrison and Rutström 2007). One good feature of the small 10-cent range of possible payoffs at stake in our quadratic scoring rule (many other experimenters use a similarly small range) is that risk attitudes are unlikely to matter much over such a small payoff range.

2.4 Software, subjects, sessions and other procedures.

Volunteer student subjects were recruited at the University of South Carolina; sessions lasted about ninety minutes. Subjects were seated at visually isolated computer stations. In addition to earnings from game play (and, in the SR treatment, stated beliefs,) subjects were paid a standard \$5 show-up fee. Each subject was randomly and anonymously matched to a single partner for all 36 periods of the session. The planned minimum sample size was 80 subjects (N=40 pairs) in all three treatments, but actual final session sizes in the SR and EC treatments slightly exceeded what was required to meet the planned minimum sample sizes. As a result, we have 80 subjects (N=40 pairs) in the NB treatment, 92 subjects (N=46 pairs) in the EC treatment and also 92 subjects (N=46 pairs) in the SR treatment.

The experimental interface was programmed using the z-Tree software (Fischbacher 2007). The software presented the game as a pair of 2x2 payoff tables, with the subject's own payoffs in the left table and her partner's payoffs in the right table.⁶ Subjects were prompted to choose one of their strategies using radio buttons next to their two possible strategies, and were then asked to confirm their choice. When both subjects had done so, a new screen appeared, reporting both subjects' choices and the resulting earnings of both members of the pair. After finishing reading the review screen, subjects clicked "continue" and the software took them to the next period of play until all thirty-six periods were done.

Prior to the strategy choice in each period, EC treatment subjects simply guessed which strategy their partner would play in that period, without any reward for accuracy. SR treatment subjects instead stated a belief that their partner would play one of her two available strategies, with scoring rule rewards for accuracy. The screen showed a table with eleven rows representing

⁶ All subjects viewed themselves as row players and their partner as column player. All supporting documents, including the instructions, the code, and the data are available at the ExLab Digital Library: <http://exlab.bus.ucf.edu>.

eleven possible probabilities presented in a frequentist manner (in 0.1 units, as “Z in 10 chances,” from “0 in 10 chances” to “10 in 10 chances”). The two columns of the table showed scoring rule payoffs associated with each possible choice as a function of the realized strategy choice of the partner. Subjects chose a row of the table and confirmed their choice; when both had done so, the software proceeded to the strategy choice screen as described above.

Before starting the 36 periods of play, subjects practiced their tasks in a non-interactive setting, where they were each required to enter choices for both players in an imaginary pair. This practice allowed them to gain some familiarity with the choice tasks, and also to experience some consequences for both players of different decisions made. A similar procedure was used to allow SR treatment subjects to become familiar with the scoring rule procedure.

3. Statistical Tests and Econometric Models.

We first use Wilcoxon two-sample tests to establish significant treatment effects. Nonparametric tests do this conservatively, but are relatively uninformative about the many possible sources of observed treatment differences. Therefore, once treatment differences are established with nonparametric tests, we estimate structural, belief-based models in each treatment, and then test for significant differences in estimated parameters across the treatments.

We begin with a common class of logit-based models that map a subjective expected payoff difference between strategies, based on latent beliefs, into strategy choice probabilities. Couched in terms of row player choice probabilities in our AMP19 game, the probability P_{t+1}^u that row plays up in period $t+1$ is

$$(1) \quad P_{t+1}^u = \Lambda[\lambda E_t(\Delta\Pi_{t+1})], \text{ where } \Lambda(x) \equiv (1 + e^{-x})^{-1} \text{ is the logistic c.d.f.}$$

We call $\lambda \in \mathbb{R}^+$ the sensitivity parameter. It denotes the row player's sensitivity to the expected payoff difference between her two strategies—here, $E_t(\Delta\Pi_{t+1}) = 19B_{t+1}^l - (1 - B_{t+1}^l)$, where B_{t+1}^l is row's latent belief that her column partner will play left in period $t+1$. Note carefully that so far, B_{t+1}^l is simply the theoretical entity we call a latent belief—the row player's causally crucial belief (in a belief-based theory) that her column partner will play left in period $t+1$.

Completion of such models requires an estimator of latent beliefs. Stated beliefs \tilde{B}_{t+1}^l are one such estimator: Setting $B_{t+1}^l = \tilde{B}_{t+1}^l$ yields the stated belief model that plays a central role in Nyarko and Schotter (2002), and we later examine a version of this model for comparative purposes. However, we first estimate and discuss inferred belief models that set $B_{t+1}^l = \hat{B}_{t+1}^l$, where \hat{B}_{t+1}^l is what we call an inferred belief. Inferred beliefs are estimated with the aid of some assumed specification of latent belief updating, usually based solely on a player's observations of her partner's past behavior and a handful of parameters. Those parameters become extra estimable parameters in the resulting econometric model of the dynamics of strategy choice.

Our specification of belief updating begins as a straightforward generalization of Cheung and Friedman's (1997) widely used gamma-weighted belief or GWB process:

$$(2) \quad \hat{B}_{t+1}^l = \frac{\Gamma_t^l}{\Gamma_t}.$$

Beliefs are formed as a ratio of two “discounted sums of experience:” The numerator Γ_t^l is the discounted sum of experience of the partner playing left, while the denominator Γ_t is the discounted sum of all experience of partner play (both through the end of period t). The sums are updated throughout the game as $\Gamma_t^l = \mathcal{X}_{t-1}^l + l_t$ and $\Gamma_t = \mathcal{X}_{t-1} + 1$. The initial value Γ_0 (for the denominator sum Γ_t) may be interpreted as the initial “strength” (relative to a period of

experience) with which an initial prior belief $\hat{B}_1^l \in [0,1]$ is held. Once one specifies Γ_0 and \hat{B}_1^l , (2) implies an initial value $\Gamma_0^l = \hat{B}_1^l \Gamma_0$ (for the numerator sum Γ_t^l). This completes the GWB process as a function of three parameters: $\gamma \in [0,1]$, $\hat{B}_1^l \in [0,1]$ and $\Gamma_0 \in [0, (1-\gamma)^{-1}]$.⁷

If players regard partners as nonstationary stochastic processes, which is sensible if partners are learning too, the discounting of past experience (at some rate γ) is a reasonable assumption (Cheung and Friedman 1997; Fudenberg and Levine 1998); it might also simply reflect decaying memory. Though estimable prior beliefs \hat{B}_1^l were not part of Cheung and Friedman's own estimation, they discussed the possibility and it has been added to GWB by others (e.g., Battalio, Samuelson and Van Huyck 2001). Adding the separate initial strength parameter Γ_0 for the prior is novel in GWB, but this is a well-known feature of other adaptive learning models with a belief-based component such as EWA (Camerer and Ho 1999). So we regard all of this (so far) as conventional.

Nyarko and Schotter (2002) found that the high-frequency variance of stated beliefs is much greater than that of inferred beliefs estimated using the GWB process. This could be an artefact of belief elicitation procedures, and it also may reflect the well-known phenomenon of overconfidence in judgment (Lichtenstein, Fischhoff and Phillips 1982). Nevertheless, this finding is so striking that we think belief updating processes ought to be generalized to allow for persistent high-frequency variability of beliefs. In the received GWB model, the γ parameter can produce persistent high-frequency variability only at the expense of rapid discounting of history:

⁷ Obviously, any two of the three terms Γ_0 , Γ_0^l and \hat{B}_1^l could be estimated, with the remaining term determined by equation (2). We think it most natural to estimate the prior \hat{B}_1^l and its strength Γ_0 . The restriction $\Gamma_0 \leq (1-\gamma)^{-1}$ is usually imposed since the denominator sum Γ_t^l can grow no larger than $(1-\gamma)^{-1}$ through the accumulation of experience, according to its updating formula. It would seem odd to allow the prior strength of belief to be larger than what the posterior strength of belief could possibly be through infinite experience.

For instance, $\gamma=0$ implies a Cournot belief process in which high-frequency variability of beliefs is maximal but past observations of partner play beyond a single lag have no influence whatever on beliefs. On the other hand, values of γ close to 1 generate persistent attention to the history of partner play, but correspondingly little persistent high-frequency variability of beliefs. This means that the standard gamma-weighted belief process forces a tradeoff between high-frequency and low-frequency components of an assumed belief updating process.

One plausible way to relax this tradeoff, particularly in fixed pairing environments, is to permit high-frequency “state dependence” of an otherwise gamma-weighted belief process. As in the GWB model, a row player’s own strategy choice probability P_{t+1}^u is determined by her belief that her column partner will play left in period $t+1$ (see equation 1), but now we allow that belief to depend on the “state of game play” she observes in period t . Our 2x2 game has four possible states of game play $s \in S = \{ul,ur,dl,dr\}$: For instance, $ul \in S$ is “row plays up and column plays left.” Let $s_t \in S$ be the state that occurs in period t . A player observes her partner’s choice, and hence this state, at the conclusion of period t . Therefore, a player could condition her beliefs about the future action of her partner (in period $t+1$) on the state s_t (observed in period t). To do that, a row player could keep track of four different kinds of information about her column partner’s past actions: What column did in periods j immediately following observations of state ul in periods $j-1$; what column did in periods k immediately following observations of state ur in periods $k-1$; and so forth. This is the essence of a very simple state-dependent belief process.

Formally, define the updating of conditional discounted experience sums as follows:

$$(3) \quad \Gamma_t^{cl}(s) = \gamma \Gamma_{t-1}^{cl}(s) + I_t(s_{t-1} = s), \text{ and } \Gamma_t^c(s) = \gamma \Gamma_{t-1}^c(s) + I(s_{t-1} = s), \forall s \in S = \{ul,ur,dl,dr\},$$

where $I(s_{t-1} = s) = 1$ if state s is observed in period $t-1$, zero otherwise. Equation (3) divides each of the two unconditional experience sums (Γ_t^l and Γ_t^c) into four conditional sums. The updating in equation (3) depends both on column's action l_t in period t , and on the state s_{t-1} observed in period $t-1$. Once this updating is finished, these eight conditional sums create four conditional beliefs for period $t+1$, in a manner similar to the unconditional belief in equation (2):

$$(4) \quad \hat{B}_{t+1}^{cl}(s) = \frac{\Gamma_t^{cl}(s)}{\Gamma_t^c(s)} \forall t > 1, \forall s \in S = \{ul, ur, dl, dr\}.$$

Notice that these conditional beliefs do not change between periods t and $t+1$ if state s is not observed in period $t-1$ since, in this instance, equation (3) implies that

$$(5) \quad \hat{B}_{t+1}^{cl}(s) = \frac{\Gamma_t^{cl}(s)}{\Gamma_t^c(s)} = \frac{\gamma \Gamma_{t-1}^{cl}(s)}{\gamma \Gamma_{t-1}^c(s)} = \frac{\Gamma_{t-1}^{cl}(s)}{\Gamma_{t-1}^c(s)} = \hat{B}_t^{cl}(s) \quad \forall s \neq s_{t-1}.$$

This is sensible: The observation l_t in period t only alters the belief conditioned on the state that was actually observed in $t-1$. Still, the strength of the other three conditional beliefs (those conditioned on states not observed in $t-1$) do decay at rate γ .

After the updating of experience sums in equation (3), and the determination of conditional beliefs in equation (4), the state s_t observed in period t selects one of the four conditional beliefs to determine the row player's belief (and her choice probability P_{t+1}^u): This is simply $\hat{B}_{t+1}^{cl}(s_t)$. Notice that the states observed in $t-1$ and t play quite distinct roles in determining beliefs (and hence choice probabilities). The distinction is exactly that between estimation and prediction. The player updates her conditional belief estimates on the basis of once-lagged states: During period t , this updating of estimates is based on l_t and the lagged state s_{t-1} . Then, for the purpose of predicting how column will play in period $t+1$, which determines

row's choice probability, the updated estimates are projected ahead one period using the latest conditioning information—the state observation s_t in period t .

If players completely condition beliefs on past states, $\hat{B}_{t+1}^{cl}(s_t)$ would be the belief at period $t+1$. However, players may only partially condition beliefs on past states, and we formalize this by assuming that beliefs at $t+1$ are a mixture of conditional (equations 3 and 4) and unconditional (equation 2) updating. This can be accommodated with a parameter $\theta \in [0,1]$ that mixes the two processes:

$$(6) \quad \hat{B}_{t+1}^l(s_t) = \theta \hat{B}_{t+1}^{cl}(s_t) + (1-\theta) \hat{B}_{t+1}^l = \theta \left(\frac{\Gamma_t^{cl}(s_t)}{\Gamma_t^c(s_t)} \right) + (1-\theta) \frac{\Gamma_t^l}{\Gamma_t}.$$

With such a mixture, the observation l_t in period t alters beliefs across all states through the last term, but accentuates the effect the observation has on the actual state observed in $t-1$. We call the belief updating process defined by equations (3), (4) and (6) gamma-theta beliefs. Notice that $(1-\theta)$ regresses conditional beliefs toward the ordinary gamma-weighted belief in (2). This has some statistical justification since the unconditional belief reflects all past observations (and hence is a relatively low variance estimator) whereas each conditional belief reflects just a subset of past observations (those following the state on which it is conditioned). When $\theta=0$, this gamma-theta belief process is the standard unconditional gamma-weighted belief process. When $\theta=1$, beliefs are fully conditioned on the most recently observed state of game play.

From a behavioral perspective, gamma-theta beliefs allow for a kind of conditional learning not present in the standard gamma-weighted belief model. Modeling the updating process as contingent on the combination of actions of both players, rather than on the actions of the partner alone, allows players to learn to anticipate high frequency reactions of partners to one's own immediate history of play (if such reactions exist): Players can learn that new actions

by the partner are in part reactions to one's own recent play. For example, a row player's beliefs about column playing left no longer depends simply on the history of left play, but also on whether the row player herself just played up or down. At the very least, since up/left results in a very asymmetric earnings consequence compared to any other play, it is not unreasonable to expect reactions and beliefs in response to that particular play to differ in a qualitative way.

From a general strategic viewpoint, $\theta \in (0,1)$ is not merely a regression parameter. If a partner is highly predictable on the basis of past states of play (a serious mistake in AMP games), $\theta > 0$ can be profitable. Yet a player with $\theta > 0$ will herself be a more predictable, and hence more exploitable, opponent in an AMP game. So while gamma-theta beliefs algebraically behave like conditional probabilities, they might also be interpreted as "strategic decision weights" since θ may be as much a product of strategic as statistical reasoning; but we will refer to them as beliefs in spite of this subtlety. Our intent is to allow for high-frequency variance of these beliefs (or decision weights), as observed by Nyarko and Schotter (2002) for stated beliefs, in a manner that is also useful to our purpose. If belief elicitation causes relatively more deliberate and conscious modeling of opponents, this could result in either larger or smaller values of θ , depending on how that modeling balances the possibility that the partner may be predictable against the possibility that the partner may exploit one's own predictability.

While the belief updating of equation (6) adds just one parameter θ , a complete model needs a specification of prior beliefs and initial strengths (as with gamma-weighted beliefs). In principle, each conditional prior belief could have its own initial value and strength, so a maximally "fat" specification could add (in net) six more parameters to the model (remove the unconditional initial belief and its initial strength, and replace them with four initial conditional beliefs $\hat{B}_1^{cl}(s)$ and their four associated initial strengths $\Gamma_0^c(s)$). We opt for the opposite in the

interest of parsimony—a maximally “lean” specification that adds no extra parameters. First, we assume that $\hat{B}_1^{cl}(s) = \hat{B}_1^l \forall s$: We let all four conditional prior beliefs equal a single unconditional prior belief parameter. In words, we assume that the player applies something like the principle of insufficient reason to conditional prior beliefs (she sets them all equal to \hat{B}_1^l). This could be a reasonable assumption if players believe that other players have highly heterogeneous dynamic behavior. But our main reason for this assumption is parsimony, not any theoretical conviction.

We also assume that the four initial strengths $\Gamma_0^c(s)$ are proportional to four implicit “prior state likelihoods” $\pi_1(s)$. Equation (1) and the row player’s prior belief \hat{B}_1^l are sufficient to determine her initial choice probability P_1^u . This implies that row’s “prior likelihood of state ul ” is $\pi_1(ul) = P_1^u \hat{B}_1^l$. Now, equation (3) has long-run implications about conditional strengths and state likelihoods: In particular, the four conditional strengths have asymptotic expected values $E[\Gamma_\infty^c(s)] = \text{Pr}_\infty(s)/(1-\gamma)$, where $\text{Pr}_\infty(s)$ is the asymptotic likelihood of state s . That is, asymptotic conditional strengths are proportional to asymptotic state likelihoods. We impose the same relationship on the players’ prior state likelihoods and initial conditional strengths: For instance, for state $s = ul$, we set $\Gamma_0^c(ul) = P_1^u \hat{B}_1^l \Gamma_0$. Determined this way, the initial conditional strengths also sum to the initial unconditional strength as they must (also an implication of equation 3), since the four prior state likelihoods $\pi_1(s)$ sum to one. Finally, for updating in the first period, we let prior state likelihoods take the place of the “period zero state indicator” $I(s_0 = s)$ in equation (3): For instance, for state $s = ul$, we have $\Gamma_1^c(ul) = \gamma \Gamma_0^c(ul) + P_1^u \hat{B}_1^l$. In these ways, initial conditions are determined by \hat{B}_1^l and Γ_0 alone without additional parameters.

Another parameter $\alpha \in \mathbb{R}$ frequently appears in models resembling equation (1), usually in the form $P_{t+1}^u = \Lambda[\lambda E_t(\Delta \Pi_{t+1}) + \alpha]$. Cheung and Friedman (1997) interpret α as a bias in strategy choice not attributable to subjective expected payoff differences between strategies. Battalio, Samuelson and Van Huyck (2001) interpret α as allowing for players' attempts to encourage selection of (or an aspiration toward) a preferred equilibrium in future repetitions of a game with two equilibria. We add a similar term for a reason resembling that of Battalio, Samuelson and Van Huyck—to allow for forward-looking behavior of players. Since we use a fixed pairing design, forward-looking reasoning (that anticipates a partner's future reactions to decisions made now) may be common. In many games, models that specify only stage-game strategies as objects of learning may be inappropriate for fixed pairing designs, where forward-looking strategies may be reasonable. For this reason, many experimental protocols meant to test these “simple” adaptive learning models use random rematching protocols or protocols where each player meets every possible partner in no more than one period. We use a fixed pairing protocol so that our simple nonparametric tests, which treat each row player's time series of strategy choices as independent of all other row players' strategy choices and abstract from learning processes, are statistically justified. However, forward-looking or “sophisticated” learning models can also be constructed to handle fixed-pairing environments; Camerer, Ho and Chong's (2002) “sophisticated EWA model” is such a modification of “adaptive EWA.”

To represent forward-looking behavior in a parsimonious way, let α denote a row player's perception of the change in her expected payoff difference in the next period brought about by playing up (rather than down) in the current period.⁸ Suppose also that the row player

⁸ We are assuming that α is either constant across periods, or that its variability across periods is small enough that approximating it by a constant is econometrically inconsequential. However, we do not assume that the row player

believes that her column player partner learns according to some model that discounts experience at the same rate γ as her own, and believes that this implies a roughly geometrically declining impact of her period t decision on changes in expected future payoffs in periods $t+k$ as k grows. This implies that her total perceived change in expected future payoffs over all remaining periods (from playing up in period t) is $\alpha + \gamma\alpha + \gamma^2\alpha + \dots + \gamma^{T-t}\alpha = \alpha(1 - \gamma^{T-t})/(1 - \gamma)$. So, to account for forward-looking behavior with one extra parameter α , we modify equation (1) to be

$$(7) \quad P_{t+1}^u = \Lambda[\lambda E(\Delta\Pi_{t+1}) + \alpha(1 - \gamma^{T-t})/(1 - \gamma)].$$

We expect $\alpha < 0$ since players' incentives are strictly opposed in the AMP19 game.⁹

Nevertheless, belief elicitation may influence forward-looking behavior and significant treatment effects on α would reflect this.

We will call equation (7), with the gamma-theta belief updating process given by equations (3), (4) and (6), an “extended gamma-weighted belief” or EGWB model. We do not offer this as a new competitor in the universe of learning models. Instead, we view it as an econometric model that is sufficiently flexible to capture empirical regularities and theoretical possibilities discussed above that cannot be captured within the simple GWB model.

Before examining the results, a few remarks on estimation and alternative specifications are in order. Wilcox (2006) shows that pooled estimation of learning models, that is, estimation that specifies a single shared parameter vector for all players, can result in severely biased and inconsistent estimates when the learning model contains own lagged dependent variables and learning model parameters in fact vary across players. This is a general econometric

knows all of the column player's adaptive learning structure, nor that α is the change in one-period-ahead payoffs implied by that structure: In other words, α denotes an expectation, but not necessarily a rational one.

⁹ The row player's maximum payoff opportunity occurs when the column player plays left. But when the row player plays up to attempt to capitalize on it, the column player's belief that row players will play up in the future is strengthened. Given the column player's payoffs, this causes column players to play left less often, reducing the row player's future payoff opportunities.

phenomenon (see, e.g., Wooldridge 2002). Although reinforcement learning models and hybrid models such as EWA (Camerer and Ho 1999) contain own lagged dependent variables, belief-based models typically do not. For instance, in the received GWB model, beliefs of the row player are indeed directly conditioned on lagged strategy choices of the column partner, but not on the lagged strategy choices of the row player herself. However, the EGWB model does contain own lagged dependent variables since the gamma-theta belief process is conditioned on the lagged state of game play, which depends, in part, on a player’s own lagged choices.

We employ the “mixed random estimator” recommended by Wilcox (2006) to account for heterogeneity.¹⁰ This estimator assumes that $\ln(\lambda)$ is normally distributed with mean μ_λ and standard deviation σ_λ across subjects,¹¹ and also adds a normally distributed, mean zero subject-specific additive random effect ν with standard deviation σ_ν to the latent variable in the EGWB model. These distributions are assumed to be independent. The estimator is implemented in the usual way for random parameters: λ and ν are numerically integrated out of the EGWB likelihood for each player, conditional on μ_λ , σ_λ and σ_ν ; the log of this is summed across players in a sample; and this sum is maximized in α , γ , \hat{B}_1^l , Γ_0 , θ , μ_λ , σ_λ and σ_ν .

The gamma-theta belief process is only one of many alternative belief updating processes that might capture high-frequency variance of beliefs. We examined some alternative belief updating specifications, and two summary facts are important. First, neither standard gamma-weighted beliefs, nor the other alternatives we examined, fit our data as well as gamma-theta

¹⁰ Since it is the gamma-theta belief process that contains own lagged dependent variables, and since θ controls the degree of conditioning on them, the expectation would be that pooled estimation will bias estimates of θ upward. In fact, the pooled MLE of θ is about 15-25 percent larger than the random parameters estimate of θ in all three treatments, illustrating this expectation and suggesting that random parameters estimation is a sensible precaution.

¹¹ Wilcox (2006) shows that heterogeneity of λ is a particularly serious source of bias in pooled MLE estimates. In the interest of robustness checking, we also specified a 3-parameter gamma distribution for λ and find that this neither fits our data significantly better nor produces different parameter estimates than the lognormal specification.

beliefs. Second, the EGWB model with gamma-theta beliefs generally results in larger p-values against hypotheses concerning equality of model parameters across treatments—that is, weaker evidence of treatment effects—than gamma-weighted beliefs or any alternative we examined.¹² Relative to these alternatives, then, the EGWB model is the least “friendly” model we have examined toward our basic hypothesis that belief elicitation changes play.

4. Results

4.1 Graphical summary and nonparametric test results.

Figures 1 and 2 summarize observed patterns of play for row and column players in the three treatments, aggregated into blocks of 12 periods, and Table 1 presents associated simple non-parametric tests based on this level of aggregation. Significant treatment differences appear for row players, while they do not for column players. The Wilcoxon two-sample tests in Table 1 treat the proportion of up plays (of each row player) and left plays (of each column player) in the first and last block of 12 periods, as well as the difference between these proportions, as single observations on each subject. This is exactly the testing procedure examined in the power planning simulations discussed in Rutström and Wilcox (2008). These are conservative tests because they treat a function (averages, or differences in averages) of each subjects’ whole time series of play as a single observation on that subject, and assume nothing about underlying structural models guiding play and/or learning. For row players, significant differences between the SR and NB treatments are found; therefore we will delve into structural modeling of treatment differences for row players below, but not for column players since the nonparametric

¹² Moreover, the significant treatment effects we identify subsequently with the EGWB model are also found under all the alternative specifications we examined that include parameters with similar interpretations.

results do not warrant it. We note that our power planning concentrated on making power large for the row players, since Erev and Roth's (1998) discussion suggested that we might create the largest treatment effect for them, though our power planning suggested considerable power for detecting treatment difference in column play as well. It may be that because column player incentives are much weaker in our design, they are less responsive to treatment variations, though we discuss an alternative interpretation in our conclusions.

4.2 Estimates of the EGWB model: Parametric location of treatment differences.

Table 2 shows the estimated EGWB model for row players in each treatment along with likelihood ratio tests against various hypotheses. The top of the table shows the results of testing the equivalence of the entire parameter vector across pairs of treatments and all three treatments together. In agreement with the nonparametric test results in Table 1, there is a strongly significant difference between the SR and NB model estimates but no significant difference between the NB and EC model estimates. However, unlike the nonparametric test, this model-based test does suggest a weakly significant difference between the SR and EC treatments. In summary, the EGWB model suggests that the particularly "intrusive" scoring rule changes behavior, while the less intrusive expected choice procedure does not (cf. Croson 1999, 2000). Because of this, we now turn to a parameter-by-parameter examination of differences between the estimated SR treatment model and the estimated model in the other two treatments.

The parameters α , γ and θ do not appear to vary significantly across the treatments. Put differently, neither forward-looking calculation (represented by α) nor the dynamics of belief learning (represented by γ and θ) appear to be altered by belief elicitation. In all three treatments, α is significantly negative, as expected if row players consider the negative future

payoff consequences of playing up in any period $t < T$ due to column's expected reactions. Estimates of γ , the discount rate on observations of partner behavior, are in a range from about 0.90 to 0.95, similar to what others estimate for the GWB model with a prior belief (e.g. Battalio, Samuelson and Van Huyck 2001), and indicating a low-frequency persistence of observational history much more like fictitious play ($\gamma = 1$) than Cournot play ($\gamma = 0$). The conditionality parameter θ takes values from 0.2 to 0.35 and is highly significant in all treatments.¹³ We return to the question of the practical significance of these estimates of θ shortly. Significant parametric effects of scoring rule procedures relative to no belief elicitation are confined to three areas. These are: (1) More realistic and stronger initial conditions of beliefs (represented by \hat{B}_1^i and Γ_0); (2) greater sensitivity to expected payoff differences (represented by the expected value of λ) and an improvement of model fit (represented by log likelihoods); and (3) greater subject heterogeneity that is not explained by the structural EGWB model (represented by σ_v). We now discuss each of these in turn.¹⁴

In some respects, the significant differences between initial inferred belief conditions and their accuracy across the SR, EC and NB treatments is the most interesting effect of belief elicitation. Figure 3 shows the fit (in terms of log likelihood) of row players' estimated gamma-theta beliefs to the actual play of their column partners by 12 period block.¹⁵ It is clear that there is a radical difference in this fit in early periods (periods 1-12); however, by the late periods

¹³ While it is tempting to draw this conclusion (and similar ones) on the basis of the standard errors reported in Table 2, or on the basis of likelihood ratio tests against $\theta = 0$ (the reduction in the log likelihood is quite large), neither Wald tests nor likelihood ratio tests are asymptotically valid when the restriction being tested lies on the boundary of the allowed parameter space. The appropriate test in this case is the LM or score test (Wooldridge 2002). In the NB, EC and SR treatments, the LM test statistics against $\theta = 0$ are 24.21, 12.05 and 30.09, respectively, distributed χ^2 with one degree of freedom—all (obviously) producing miniscule p -values.

¹⁴ As a robustness test we also estimate the model with θ restricted to 0. Qualitative findings from this estimation are the same: prior beliefs in left play are still significantly lower with belief elicitation and the strength of the priors are significantly higher. In fact, all inferences are qualitatively the same with greater statistical significance.

¹⁵ In other words, these are not the maximized likelihoods of the estimated models of row player behavior, but rather the fit of estimated gamma-theta inferred belief updating process to the actual play of column partners.

(periods 25-36), there is little difference in fit across the three treatments. We believe this occurs because scoring rule procedures prompt more conscious pre-play modeling of partners, decreasing the relative contribution of unconscious processes to latent beliefs. It takes perhaps twenty periods of experience in the NB treatment to match the improved accuracy of inferred beliefs brought about at the very beginning of play by the scoring rule procedure in the SR treatment. The actual proportions of first period plays of left by column partners in the NB, EC and SR treatments are 0.40, 0.22 and 0.35, with standard errors 0.078, 0.061 and 0.071, respectively. Corresponding inferred prior beliefs \hat{B}_1^l are 0.89, 0.62 and 0.45, with standard errors of 0.29, 0.15 and 0.09, respectively. Inferred prior beliefs closely match actual first period play of column players only in the SR treatment, and are especially biased in the NB treatment.

The second effect of scoring rule procedures is, quite simply, that the EGWB model fits row player behavior much better with than without them. The last row of Table 2 shows that the estimated log likelihoods of the EGWB model in the NB and SR treatment (per 40 subjects) are -866.78 and -820.68 , respectively.¹⁶ Parametrically, the estimated expected value of λ —the sensitivity of choice to modeled expected payoff differences—is more than fifty percent larger in the SR treatment than in the NB treatment, suggesting that inferred belief updating processes have greater explanatory force in the presence of scoring rule procedures.¹⁷

Third, σ_v is significantly greater in the SR treatment than the NB treatment. We interpret v as representing relatively persistent but unmodeled differences in behavior across individual

¹⁶ To put this 5.62% poorer model fit in the NB treatment into empirical perspective, consider that the median percent reduction in log likelihoods reported by Camerer, Ho and Chong (2002) from using reinforcement learning rather than adaptive EWA across thirty games, player types and/or treatment conditions is 2.52%.

¹⁷ λ is usually viewed as sensitivity to expected payoff differences, but this interpretation implicitly assumes that the inferred belief updating process determining expected payoff differences is properly specified. If we relax this assumption, the estimated size of λ can also be viewed as reflecting (in part) specification errors in the model of expected payoff differences, including specification errors in the inferred belief updating process. This is the sense in which higher estimated values of λ can signal greater predictive force of inferred belief updating processes.

row players. Under this interpretation, larger σ_v implies greater unexplained but systematic variance of row player behavior in the SR treatment. This suggests that the belief processes encouraged by scoring rule procedures are more variable across players than the ones that prevail without them. If conscious attentional resources vary more across individuals than unconscious ones do, and if these attentional resources control the ability to inhibit or suppress relatively automatic strategic inclinations based on unconscious processes (e.g. Feldman-Barrett, Tugade and Engle 2004), this difference makes some sense.

4.3 The relative explanatory power of stated and inferred beliefs.

Nyarko and Schotter (2002) found that stated beliefs have much greater high-frequency variance than do gamma-weighted inferred beliefs, and that the stated belief model explains player choices better than the gamma-weighted inferred belief model. Are these results also true in our setting? Table 3 compares certain characteristics of three estimators of latent beliefs (gamma-weighted beliefs, gamma-theta beliefs and stated beliefs) in the SR treatment. The gamma-theta beliefs are those implied by the estimates of α , γ , θ , \hat{B}'_1 and Γ_0 shown in Table 2 for the SR treatment data. The gamma-weighted beliefs are those implied by estimates of the EGWB model in the SR treatment when the constraint $\theta=0$ is imposed. The stated belief model uses row players' stated beliefs in the SR treatment, and therefore (obviously) includes none of the three inferred belief process parameters θ , \hat{B}'_1 or Γ_0 . However, it does include the parameters α and γ to represent forward-looking calculation by the players as in equation (7). Both the stated belief model and the gamma-weighted belief model are estimated using the same "mixed random estimator" used to estimate the EGWB model (for comparability of fit).

The first row of Table 3 compares the relative high-frequency variability of the three estimators of latent beliefs. The ratio $\hat{\sigma}_{i\Delta B} / \hat{\sigma}_{iB}$ measures the variance of high-frequency changes (i.e. changes across consecutive time periods) in a series of belief estimates relative to the overall time series variability of those estimates (for row player i). In Table 3, our mean value of $\hat{\sigma}_{i\Delta B} / \hat{\sigma}_{iB}$ for stated beliefs is more than twice that for gamma-weighted beliefs. However, also notice that gamma-theta beliefs narrow this gap quite substantially. The last three columns of Table 3 report one-sample tests of the hypothesis that $\hat{\sigma}_{i\Delta B} / \hat{\sigma}_{iB}$ does not vary across pairs of belief estimators, and this is easily rejected in all cases. Thus, we replicate Nyarko and Schotter’s (2002) finding that stated beliefs have much greater high-frequency variability than gamma-weighted beliefs, but the high-frequency variability of estimated gamma-theta beliefs lies squarely and significantly between them. The estimated value of θ in the SR treatment (0.30) produces a pronounced increase in the high-frequency variability of estimated beliefs relative to the gamma-weighted belief model—though not to the level observed for stated beliefs. Recall that our motivation for the gamma-theta belief updating process was to allow for greater high-frequency variability of inferred beliefs.

The rest of Table 3 compares three measures of the predictive value of the three belief estimators in the model explaining row player’s choices. Let $\hat{\rho}_{iBu}$ denote a Spearman rank correlation coefficient between choices of “up” (measured by a dummy variable) and any one of the three estimated latent belief series \bar{B}_{it} , for each player i .¹⁸ The second row of Table 3 shows the mean of this correlation for each of the three belief estimators, as well as tests for differences

¹⁸ The Spearman correlation is a rank correlation coefficient appropriate to cases where one or both variables do not have cardinal meaning or are qualitative in nature. The latter is true here since u_{it} is a qualitative variable. Kendall’s τ could also be used, but it throws away potentially useful magnitude information represented by the rank of \bar{B}_{it} that is used by the Spearman coefficient, replacing it with purely ordinal information about \bar{B}_{it} .

between them. The third row of Table 3 shows the mean (across row players) of estimated log likelihoods of the models based on each of the three belief estimators, per player. These two evaluations of the predictive value of the estimators do not replicate Nyarko and Schotter’s (2002) findings: Both suggest that stated beliefs are a worse predictor of row player behavior than either gamma-theta or gamma-weighted beliefs, though the comparison is statistically convincing only for gamma-theta beliefs.¹⁹ Overall, this evidence contradicts the third assumption about stated beliefs we outlined earlier: Stated beliefs are sometimes worse estimators of latent beliefs than are inferred beliefs.

Do stated beliefs predict anything better than the inferred belief updating processes? The last row of Table 3 computes Spearman correlations $\hat{\rho}_{i\Delta B\Delta u}$ between changes in “up” plays across consecutive time periods and changes in belief estimates $\Delta \bar{B}_i$ from each of the three estimators, for each player i . It is here that stated beliefs show some comparative advantage, significantly (but weakly) outperforming gamma-weighted beliefs according to two of the three test statistics, and not performing significantly worse than gamma-theta beliefs. Stated beliefs would then have a predictive advantage in games with a high frequency of changes in play.

5. Discussion and Conclusions

Is belief elicitation a good tool for understanding behavior in games? Clearly, if it provided us with an unbiased and efficient estimator of latent beliefs without altering the very behavior we wish to explain, it would be extremely valuable for empirical study of behavior in

¹⁹ Nyarko and Schotter’s (2002) scoring rule procedure allows subjects to state beliefs more finely (in 0.01 probability units) than ours does (in 0.10 probability units). To see whether this matters, we used Nyarko and Schotter’s “experiment 1” data (which most closely resembles ours) to estimate their own stated belief model (that is, their model 1 reported in their Table 4) two ways. Rounding the stated beliefs to the nearest 0.1 prior to estimating the stated belief model actually improves its fit very slightly (the log likelihood per subject, over the sixty periods of play, is -36.93 without rounding versus -36.89 with rounding).

games (Manski 2002) and in many other contexts (Manski 2004). Unfortunately, we find that scoring rule belief elicitation changes row players' observed actions, and row players are the most highly motivated players in our game. Both nonparametric tests and tests based on a structural model of belief learning support that conclusion. Estimation of the structural model suggests three kinds of changes wrought by the use of scoring rules: (1) Prior beliefs become stronger and more realistic; (2) structural belief models fit the resulting data much better; and (3) unmodeled structural heterogeneity of play is more pronounced. We also find that stated beliefs can be poorer predictors of actions than inferred beliefs. This does not weigh decisively against Manski's point (which is not that stated beliefs have unique predictive value) that strictly exogenous beliefs would have a decisive econometric advantage in model identification. Unfortunately, if belief statement alters game play—as it does here and in Erev, Bornstein and Wallsten (1993), Croson (2000), and Nelson (2003)—this hoped-for exogeneity is dubious. Our findings, limited to a repeated matching protocol matching pennies game, should motivate further tests in other game domains.

As none of this matches what Nyarko and Schotter (2002) observe, what might explain the difference? With obvious caution due to the limited number of game structures that have been exposed to such tests, here are some possibilities. First, we think statistical power plays a role. The appendix illustrates how statistical power in an AMP game can vary quite strongly with the asymmetry of the payoffs, holding other design features constant. Our game was designed explicitly to generate statistical power to detect shifts from non-belief-based to belief-based learning models, and is therefore highly asymmetric. Other games designed with some other primary focus, such as Nyarko and Schotter's design, do not have this property.²⁰ Therefore, it

²⁰ Nyarko and Schotter (2002) originally set out to compare stated and inferred belief models and chose their game with that central purpose in mind.

should not be surprising that their specific game produces no evidence of belief elicitation effects, while we and Croson (1999, 2000), who deliberately chose games to look for this, do.

Beyond statistical power, we have one theoretical conjecture based on our findings and those of Croson (2000) and Nyarko and Schotter (2002). We only find belief elicitation effects for row players, but not column players, in our AMP19 game, even though our power analysis shows that statistical power should be high enough to detect early period shifts in column play as well. Indeed, although we do not show the results, we performed all of our analyses for our column players as well: For them, all results closely resemble those of Nyarko and Schotter.

This contrast may reveal something useful. Figure 4 reproduces our AMP19 game, along with Nyarko and Schotter's zero-sum game and the prisoner's dilemma where Croson finds strong belief elicitation effects on game play. It is striking that across these three games, each player type who shows a significant belief elicitation effect (our row players, and both of Croson's players) also see some arguably pronounced affective asymmetry between the two strategies available to them. For row players in our game, the asymmetry is immediately obvious. In Croson's prisoner's dilemma, the affective asymmetry is between the high and equitable payoffs of cooperation versus the temptation to defect and the regret of being cheated. By contrast, no strong asymmetry is present for player types who show no belief elicitation effect (our column players, and both of Nyarko and Schotter's players).

These results are consistent with dual-process theories of mind (Feldman-Barrett, Tugade and Engle 2004) if the payoff asymmetry invokes relatively automatic emotional predispositions favoring up over down. As belief elicitation procedures require the decision maker to estimate latent beliefs, it is possible that attention is brought to the conflict between the desirability of the high payoff to row and the lack of desirability for column to play left when row plays up.

Attention to these conflicting desires could motivate row to change how he plays the game. If this is the case, then in games where automatic predispositions lead to the same choices as attentive deliberations belief elicitation may not have an effect on game play.

Our finding that stated beliefs can be worse predictors of actions than inferred beliefs needs to be replicated in other games, but if replicated it is important. We note that plenty of other evidence points to strong dissociations between demonstrable causes of human behavior and human abilities to accurately declare those causes (Nisbett and Wilson 1977, Ross 1989). To the extent that latent beliefs are below the level of consciousness, and their conscious products are mostly inclinations to action rather than explicit beliefs, subjects' stated beliefs can in part be inferences about the causes of their own actions or inclinations. Therefore, it should not be all that surprising that there may also be a dissociation between latent and stated beliefs. Some optimism is still called for since the less intrusive belief elicitation mechanism that we used in our EC treatment, where subjects simply report which strategy partners are most likely to play, affects play to a lesser degree than does the scoring rule.²¹ Of course, this elicitation is a rather blunt instrument for latent beliefs. Yet this fact gives us hope that some belief elicitation mechanism might be informative about latent beliefs yet not so intrusive as to alter play, and thus give researchers a suitable "belief instrument" for identification and theory tests. One potentially useful possibility is a suitably fine-grained verbal response scale. This technique has been extensively studied by statisticians and psychologists for many years; Mosteller and Youtz (1990) provide a useful review. Our point is that there are many imaginable techniques for eliciting a belief instrument, if not beliefs themselves, as this example shows. Some of those may have no effect on game play and, at the same time, may be good instruments for latent beliefs.

²¹ But Croson (1999, 2000) documents belief elicitation effects even with this procedure.

References

- Aczel, J. and J. Pfanzagl. 1966. Remarks on the measurement of subjective probability and information. Metrika 11:91-105.
- Andersen, Steffen, John Fountain, Glenn W. Harrison and E. Elisabet Rutström. 2007. Eliciting beliefs: Theory and Experiments. University of Central Florida, Department of Economics Working Paper 07-08.
- Battalio, Ray, Larry Samuelson, and John Van Hyuck. 2001. Optimization incentives and coordination failures in laboratory stag hunt games. *Econometrica*, 69(3), December, pp. 749-764.
- Bracht, Juergen, and Hidehiko Ichimura. 2001. Identification of a general learning model on experimental game data. Hebrew University of Jerusalem, working paper. December.
- Camerer, Colin and Teck-Hua Ho. 1999. Experience weighted attraction learning in normal-form games. Econometrica 67:827-74.
- Camerer, Colin, Teck-Hua Ho, and Juin-Kuan Chong. 2002. Sophisticated experience-weighted attraction learning and strategic teaching in repeated games. Journal of Economic Theory
- Cheung, Yin-Wong and Daniel Friedman. 1997. Individual learning in normal form games: Some laboratory results. Games and Economic Behavior. 19 (46-76).
- Churchland, Paul M. 1981. Eliminative materialism and the propositional attitudes. Journal of Philosophy 78(2):67-90.
- Cohen, Jacob. 1988. *Statistical Power Analysis for the Behavioral Sciences*, 2nd ed. Hillsdale N.J.: Erlbaum.
- Costa-Gomes, Miguel A. and Georg Weizsäcker. 2008. Stated beliefs and play in normal-form games. *Review of Economic Studies*, forthcoming.
- Croson, Rachel. 1999. The disjunction effect and reason-based choice in games. Organizational Behavior and Human Decision Processes 80(2):118–33.
- _____. 2000. Thinking like a game theorist: Factors affecting the frequency of equilibrium play. Journal of Economic Behavior and Organization 41(3):299-314.
- Erev, Ido, Gary Bornstein and Thomas S. Wallsten. 1993. The negative effect of probability assessments on decision quality. Organizational Behavior and Human Decision Processes 55:78-94.

- Erev, Ido and Alvin E. Roth. 1998. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. American Economic Review 88:848-81.
- Feldman-Barrett, Lisa, Michele M. Tugade and Randall W. Engle. 2004. Individual differences in working memory capacity and dual-process theories of the mind. Psychological Bulletin 130:553-73.
- Fischbacher, Urs. 2007. Z-tree: Zurich Toolbox for Ready-made Economic Experiments, Experimental Economics 10(2), 171-178.
- Fudenberg, D. and D. K. Levine. 1998. The Theory of Learning in Games (Economics Learning and Social Evolution). Cambridge: MIT Press.
- Gächter, S. and E. Renner. 2006. The Effects of (Incentivized) Belief Elicitation in Public Good Experiments. The Centre for Decision Research and Experimental Economics, School of Economics, University of Nottingham, Discussion Paper 2006-16.
- Goeree, J. and C. A. Holt. 2001. Ten little treasures of game theory and ten intuitive contradictions. American Economic Review 91(5):1402-1422.
- Harrison, Glenn W., Eric Johnson, Melayne M. McInnes, and E. Elisabet, Rutström. 2005. Risk aversion and incentive effects: Comment. American Economic Review 95(3):897-901.
- Holt, Charles A. 1986. Scoring-rule procedures for eliciting subjective probability and utility functions. In Prem K. Goel and A. Zellner, eds., Bayesian Inference and Decision Techniques: Essays in Honor of Bruno de Finetti. Amsterdam: North Holland Press.
- Ichimura, H., and J. Bracht. 2001. Estimation of learning models on experimental game data. Hebrew University of Jerusalem working paper.
- Lichtenstein, Sarah, Baruch Fischhoff and Lawrence D. Phillips. 1982. Calibration of subjective probabilities: The state of the art to 1980. In D. Kahneman, P. Slovic and A. Tversky, eds., Judgment Under Uncertainty: Heuristics and Biases. Cambridge, U.K.: Cambridge University Press.
- Manski, Charles F. 2002. Identification of decision rules in experiments on simple games of proposal and response. European Economic Review 46(4-5):880-891.
- _____. 2004. Measuring expectations. Econometrica 72(5):1329-76.
- McKelvey, R. D. and T. Page. 1990. Public and private information: An experimental study of information pooling. Econometrica 58:1321-39.

- McKelvey, R., T. Palfrey, and R. A. Weber. 2000. The effects of payoff magnitude and heterogeneity on behavior in 2x2 games with unique mixed strategy equilibria. Journal of Economic Behavior and Organization 42:523-548.
- Mookerjee, D. and B. Sopher. 1997. Learning and decision costs in experimental constant-sum games. Games and Economic Behavior 19:97-132.
- Mosteller, Frederick and Cleo Youtz. 1990. Quantifying probabilistic expressions. Statistical Science 5:2-12.
- Nelson, William Robert. 2003. Dictators, expectations, and the order of decisions: An experiment. State University of New York at Buffalo School of Management working paper.
- Nisbett, Richard E. and Timothy D. Wilson. 1977. Telling more than we can know: Verbal reports on mental processes. Psychological Review 84(3): 231-259.
- Nyarko, Yaw and Andrew Schotter. 2002. An experimental study of belief learning using elicited beliefs. Econometrica 70:971-1005.
- Ochs, Jack. 1995. Games with unique, mixed strategy equilibria: An experimental study. Games and Economic Behavior 10:202-17.
- Palfrey, Thomas R. and Stephanie W. Wang. 2007. On eliciting beliefs in strategic games. California Institute of Technology, Social Science Working Paper 1271, July.
- Ross, Michael. 1989. Relation of implicit theories to the construction of personal histories. Psychological Review 96(2):341-357.
- Rutström, E. Elisabet and Nathaniel T. Wilcox. 2008. Power planning appendix to “Stated beliefs versus inferred beliefs: A methodological inquiry and experimental test.” University of Houston Department of Economics working paper. Available in ExLab Digital Library, <http://exlab.bus.ucf.edu>.
- Salmon, Timothy C. 2001. An evaluation of econometric models of adaptive learning. Econometrica 69(6): 1597-1628.
- Sarin, R. and F. Vahid. 2001. Predicting how people play games: A simple dynamic model of choice. Games and Economic Behavior 34:104-122.
- Stahl, Dale O. and Paul O. Wilson. 1995. “On players models of other players: Theory and experimental evidence”. Games and Economic Behavior, July, pp. 218-254.
- Wilcox, Nathaniel T. 2006. Theories of learning in games and heterogeneity bias. Econometrica 74:1271-1292.

Wilcox, Nathaniel T. and Nick Feltovich. 2000. Thinking like a game theorist: Comment. University of Houston Department of Economics working paper.

Wooldridge, Jeffrey. *Econometric Analysis of Cross Section and Panel Data*. Cambridge, MA: MIT Press, 2002.

Wright, W. and M. Aboul-Ezz. 1989. Effects of extrinsic incentives on the quality of frequency assessments. Organizational Behavior and Human Decision Processes 41:143-52.

Table 1

Simple Tests for Differences Between the Three Treatments

block (periods)	Proportion of row "up" choices			Proportion of column "left" choices		
	NB versus SR	NB versus EC	SR versus EC	NB versus SR	NB versus EC	SR versus EC
All Periods (1 to 36)	<i>p=0.67</i> <i>p=0.62</i>	<i>p=0.63</i> <i>p=0.46</i>	<i>p=0.38</i> <i>p=0.27</i>	<i>p=0.35</i> <i>p=0.34</i>	<i>p=0.94</i> <i>p=0.78</i>	<i>p=0.27</i> <i>p=0.23</i>
Early Periods (1 to 12)	<i>p=0.04</i> <i>p=0.055</i>	<i>p=0.92</i> <i>p=0.94</i>	<i>p=0.023</i> <i>p=0.026</i>	<i>p=0.52</i> <i>p=0.40</i>	<i>p=0.54</i> <i>p=0.72</i>	<i>p=0.21</i> <i>p=0.25</i>
Late Periods (25 to 36)	<i>p=0.35</i> <i>p=0.47</i>	<i>p=0.16</i> <i>p=0.19</i>	<i>p=0.66</i> <i>p=0.59</i>	<i>p=0.89</i> <i>p=0.78</i>	<i>p=0.92</i> <i>p=0.55</i>	<i>p=0.79</i> <i>p=0.81</i>
Change Between Early and Late Periods	<i>p=0.027</i> <i>p=0.027</i>	<i>p=0.24</i> <i>p=0.28</i>	<i>p=0.25</i> <i>p=0.25</i>	<i>p=0.63</i> <i>p=0.58</i>	<i>p=0.61</i> <i>p=0.89</i>	<i>p=0.32</i> <i>p=0.41</i>

The shaded cells indicate where the power planning predicts low power.

Table 2.

Maximum likelihood analysis of treatment differences in row player behavior, using the 6-parameter EGWB Model with random coefficients λ and random effects v .

Likelihood Ratio Tests of Equivalence of Parameter Vectors Across Treatments						
Treatments Compared	Extended GWB Model w/ lognormal λ and normal random effects v					
	Constrained Log L		Unconstrained Log L		Likelihood Ratio Test against equivalent parameter vectors	
NB versus SR	-1823.52		-1810.56		$\chi^2_8 = 25.92, p = 0.0011$	
NB versus EC	-1868.69		-1864.41		$\chi^2_8 = 8.56, p = 0.38$	
SR versus EC	-1948.25		-1941.40		$\chi^2_8 = 13.70, p = 0.090$	
All Treatments	-2824.56		-2808.18		$\chi^2_{16} = 32.77, p = 0.0079$	
Estimated Parameters (Asymptotic Standard Errors in Parentheses)						
Extended GWB Model				Likelihood Ratio Test of Various Parameter Equality Restrictions Between the SR Treatment and the other Treatments		
Parameter	NB Treat.	EC Treat.	SR Treat.	Restriction	LR Test: SR vs. EC	LR Test: SR vs. NB
α	-0.296 (0.090)	-0.477 (0.146)	-0.358 (0.0701)	equal α	$\chi^2_1 = 0.86$ $p = 0.35$	$\chi^2_1 = 0.33$ $p = 0.56$
γ	0.950 (0.0261)	0.897 (0.0352)	0.898 (0.0261)	equal γ	$\chi^2_1 = 0.0005$ $p = 0.98$	$\chi^2_1 = 2.53$ $p = 0.11$
\hat{B}'_1	0.892 (0.285)	0.615 (0.152)	0.446 (0.0901)	equal \hat{B}'_1	$\chi^2_1 = 1.54$ $p = 0.21$	$\chi^2_1 = 5.21$ $p = 0.022$
Γ_0	0.329 (0.445)	1.24 (0.661)	5.38 (3.82)	equal Γ_0	$\chi^2_1 = 3.77$ $p = 0.052$	$\chi^2_1 = 6.49$ $p = 0.011$
θ	0.352 (0.112)	0.214 (0.0692)	0.297 (0.0691)	equal θ	$\chi^2_1 = 0.72$ $p = 0.40$	$\chi^2_1 = 0.24$ $p = 0.62$
μ_λ	-1.77 (0.174)	-1.74 (0.169)	-1.28 (0.137)	equal μ_λ	$\chi^2_1 = 4.41$ $p = 0.036$	$\chi^2_1 = 4.32$ $p = 0.038$
σ_λ	0.414 (0.184)	0.640 (0.227)	0.496 (0.182)	equal σ_λ	$\chi^2_1 = 0.42$ $p = 0.51$	$\chi^2_1 = 0.15$ $p = 0.69$
σ_v	0.459 (0.115)	0.577 (0.190)	0.844 (0.158)	equal σ_v	$\chi^2_1 = 2.48$ $p = 0.12$	$\chi^2_1 = 5.67$ $p = 0.017$
$E(\lambda)^a$	0.186 (0.0301)	0.214 (0.0400)	0.314 (0.0357)	equal $E(\lambda)$	$\chi^2_1 = 2.94$ $p = 0.087$	$\chi^2_1 = 6.43$ $p = 0.011$
$\text{Var}(\lambda)^a$	0.00647 (0.00674)	0.0233 (0.0256)	0.0276 (0.0243)	equal $\text{Var}(\lambda)$	$\chi^2_1 = 0.026$ $p = 0.87$	$\chi^2_1 = 1.75$ $p = 0.19$
Log Likelihood	-866.78	-997.62	-943.78	equal μ_λ and equal σ_λ	$\chi^2_2 = 3.69$ $p = 0.16$	$\chi^2_2 = 6.55$ $p = 0.038$
Log Likelihood per 40 subjects	-866.78	-867.49	-820.68	equal \hat{B}'_1 and Γ_0	$\chi^2_2 = 3.80$ $p = 0.15$	$\chi^2_2 = 6.91$ $p = 0.032$

^aBecause μ_λ and σ_λ are the mean and variance of $\ln(\lambda)$, the actual mean and variance of λ are given by $E(\lambda) = \exp(\mu_\lambda + \sigma_\lambda^2/2)$ and $\text{Var}(\lambda) = \exp(2\mu_\lambda + \sigma_\lambda^2) \cdot [\exp(\sigma_\lambda^2) - 1]$.

Table 3

Comparison of Characteristics of Stated Beliefs, Gamma-Weighted Beliefs and Gamma-Theta Beliefs in the SR Treatment

Type of within-player statistic	Mean across row players of each type of within-player statistic, using each belief estimator			Two-tailed p-values of one-sample t-test, sign test and signed rank test against no difference across pairs of belief measures		
	Gamma-weighted beliefs	Gamma-theta beliefs	Stated beliefs	Stated vs gamma-weighted	Stated vs gamma-theta	gamma-weighted vs gamma-theta
$\hat{\sigma}_{i\Delta B} / \hat{\sigma}_{iB}$, ratio of standard deviation of high frequency changes in belief estimator to standard deviation of estimated beliefs, for row player i .	0.553	0.780	1.19	$p < 0.0001$ $p < 0.0001$ $p < 0.0001$	$p < 0.0001$ $p < 0.0001$ $p < 0.0001$	$p < 0.0001$ $p < 0.0001$ $p < 0.0001$
$\hat{\rho}_{iBu}$, Spearman correlation between belief estimator and choices of “up”, for row player i .	0.266	0.293	0.214	$p = 0.25$ $p = 0.10$ $p = 0.13$	$p = 0.097$ $p = 0.026$ $p = 0.050$	$p = 0.093$ $p = 0.044$ $p = 0.059$
LL_i , log likelihood of the estimated model for row player i .	-20.77	-20.52	-20.92	$p = 0.86$ $p = 0.46$ $p = 0.36$	$p = 0.50$ $p = 0.026$ $p = 0.072$	$p = 0.71$ $p = 0.30$ $p = 0.37$
$\hat{\rho}_{i\Delta B\Delta u}$, Spearman correlation between changes in estimated beliefs and changes in choices of “up”, for row player i .	0.087	0.230	0.184	$p = 0.077$ $p = 0.18$ $p = 0.096$	$p = 0.44$ $p = 0.30$ $p = 0.40$	$p < 0.0001$ $p < 0.0001$ $p < 0.0001$

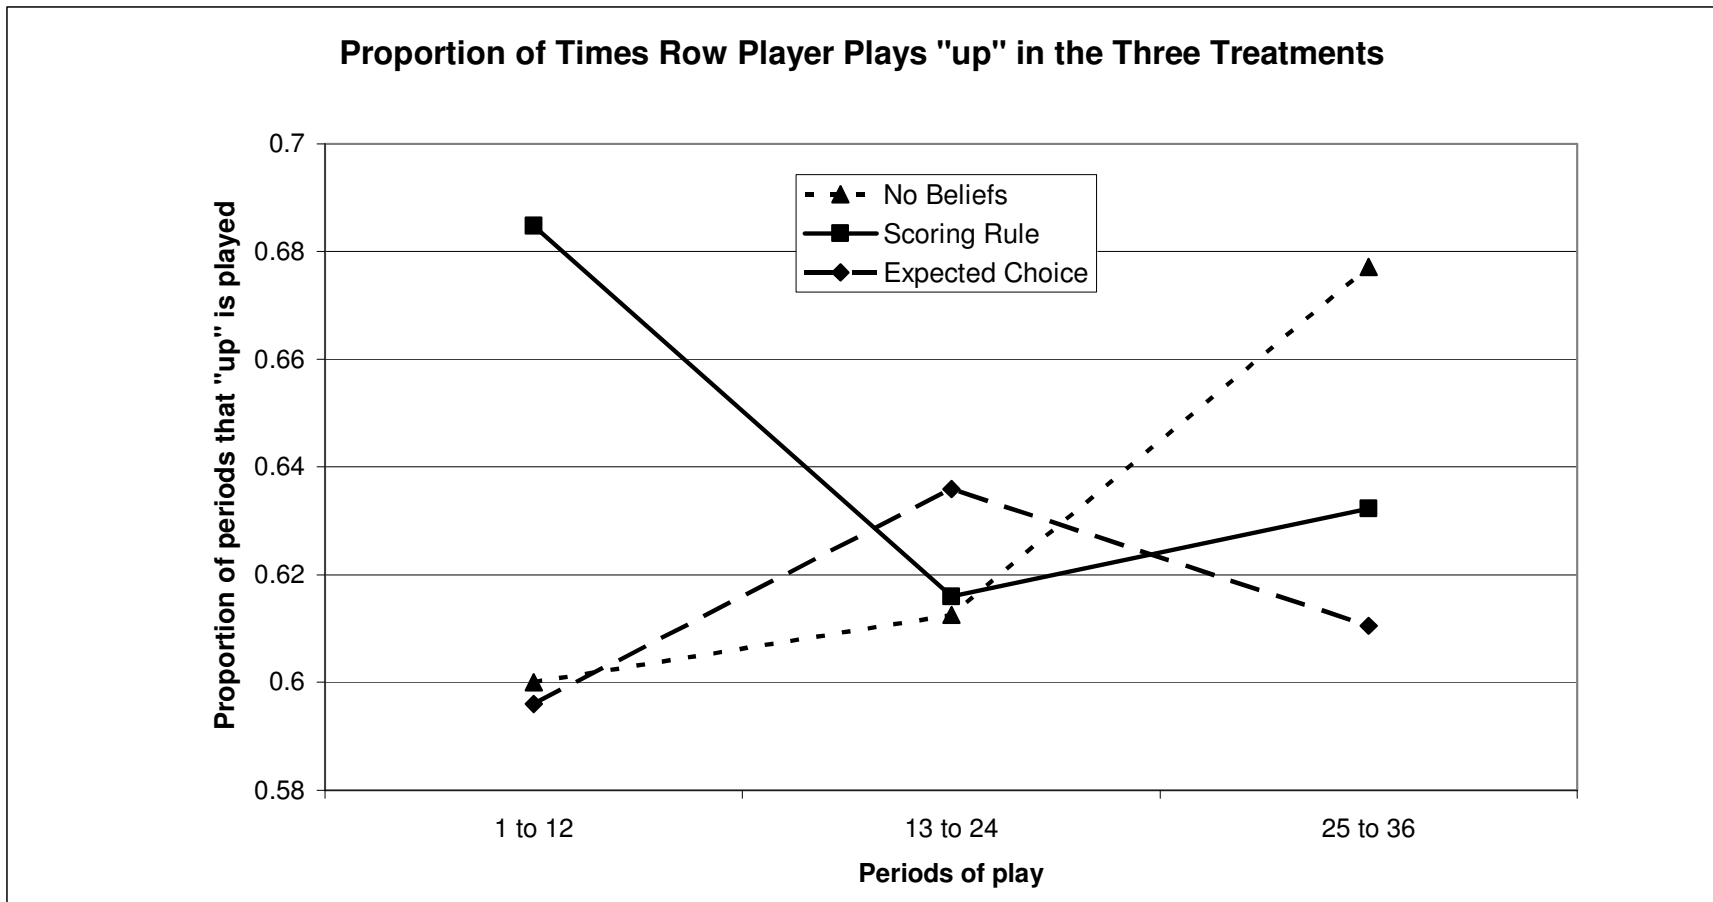


Figure 1. Proportion of times row player plays "up" in the three treatments.

Proportion of Times Column Player Plays "left" in NB and SR Treatments

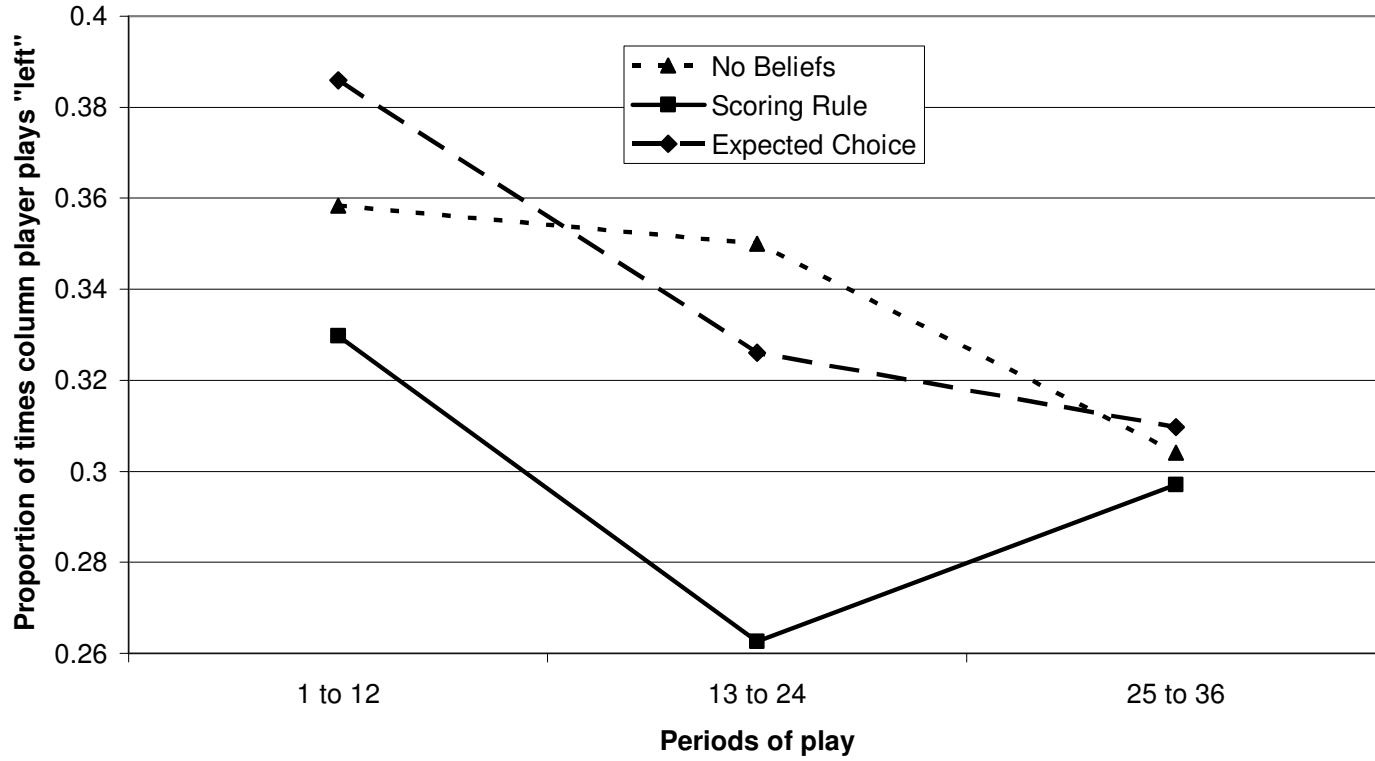


Figure 2. Proportion of times column player plays "left" in the three treatments.

Fit of row players' estimated gamma-theta beliefs to actual play of column players, by treatment and period blocks

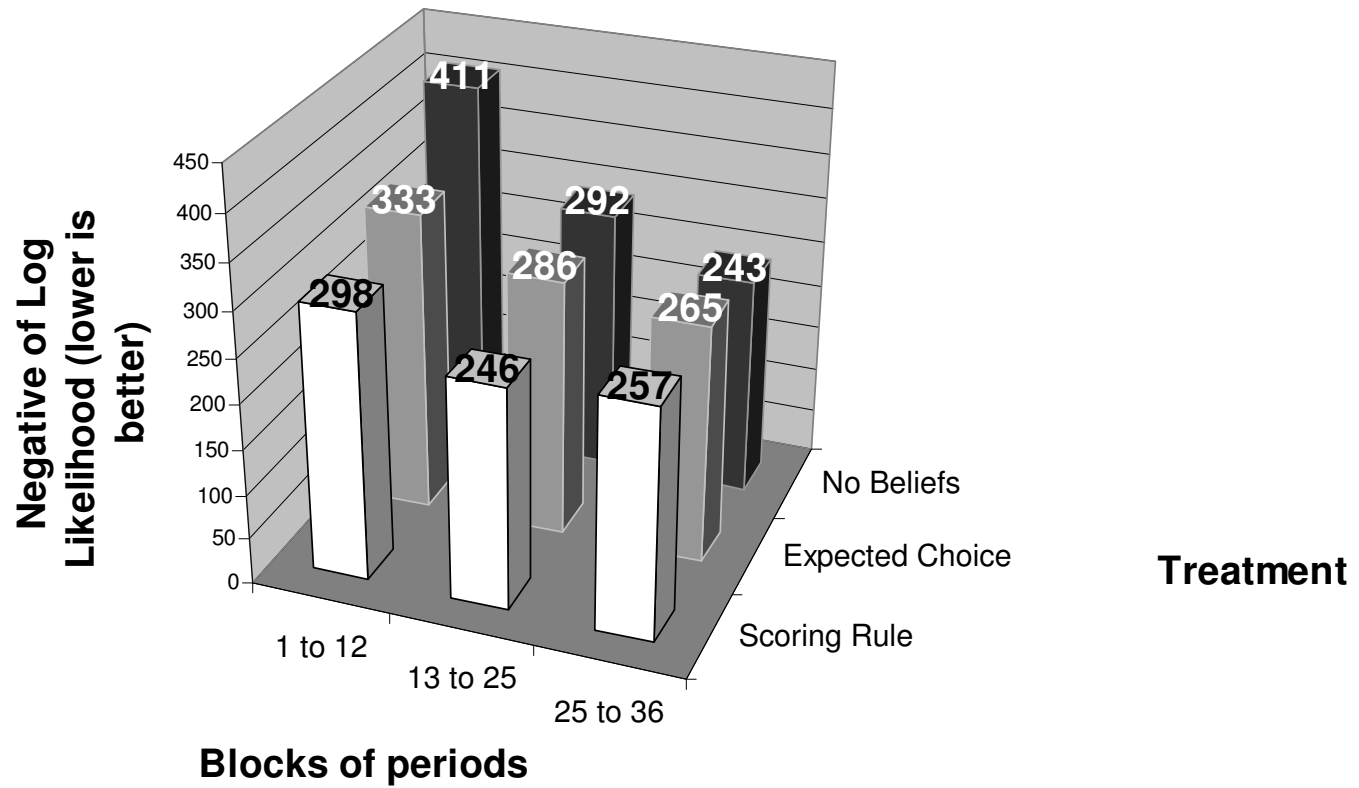


Figure 3. Fit of row players' estimated gamma-theta beliefs to actual play of column partners, by treatments and period blocks.

This study

	<i>l</i> (left)	<i>r</i> (right)
<i>u</i> (up)	(19,0)	(0,1)
<i>d</i> (down)	(0,1)	(1,0)

Nyarko and Schotter (2002)

	<i>l</i> (left)	<i>r</i> (right)
<i>u</i> (up)	(6,2)	(3,5)
<i>d</i> (down)	(3,5)	(5,3)

Croson (1999, 2000)

	<i>l</i> (left)	<i>r</i> (right)
<i>u</i> (up)	(75,75)	(25,85)
<i>d</i> (down)	(85,25)	(30,30)

Figure 4. Comparison of the games used here and in Nyarko and Schotter (2002) and Croson (1999, 2000).

Appendix: Statistical power and payoff asymmetry

Table A1 reports an analysis (based on Monte Carlo simulation) of the power of two-sample tests for differences in AMPX game behavior: X is row's payoff in the {up,left} cell of the game, whereas all other payoffs are zero or one as in symmetric matching pennies or SMP (e.g. AMP1 is identical to SMP, and AMP19 is our game). The analysis considers two samples, each with 40 row and column players in fixed pairings for 36 periods (as in our experiment). Table entries are probabilities that a null hypothesis (no difference between two such samples) is rejected by a Wilcoxon test at an $\alpha = .05$ significance level (entries are $1-\beta$, where β is the probability of type II error). As described in the text, the simulations use a 3-parameter weighted fictitious play model as the true data-generating process or DGP in one sample, and a 3-parameter reinforcement learning model as the true DGP in the other sample (see Rutström and Wilcox 2008 for more details). The first row of the table below reports power in our experimental game, AMP19. Below that, we also report power in AMPX games with degrees of asymmetry $X=1, 2, 3, 4,$ and 9 .

The table suggests that power is very good during early periods for all $X>2$. But learning models make relatively ad hoc assumptions about initial conditions—after all, initial conditions are not their primary focus. So the fact that the two assumed DGPs (learning models) produce detectable differences in early periods is relatively uninteresting—probably (at least partially) an artifact of the different assumptions the DGPs make about initial conditions. We wanted a design capable of detecting differences in later behavior too. The table shows that we only get to conventionally accepted power levels (β less than or equal to 4α , or power in excess of 0.80; see e.g. Cohen 1988) during late period play (last third of periods) when the asymmetry is strong, $X=19$, and even then only for the row player.

Notice that SMP (AMP1) produces no power at all. Ichimura and Bracht (2001) first described why this is so. Games with equal mixing over strategies in equilibrium will result in poor identification of learning model parameters in equilibrium. Both the MSNE and QRE (quantal response equilibrium) of SMP is equal mixing, and equal mixing is commonly observed in SMP games (Goeree and Holt 2001). Under these circumstances, any learning model with a precision parameter λ (EWA and GWB for instance) will be poorly identified, since $\lambda = 0$ will fit such data well. (Because λ multiplies a function of all other parameters in these models, this implies that these other model parameters are also poorly identified in SMP games.) But if no EWA parameter is well-identified in SMP equilibria, then its δ parameter, which distinguishes between reinforcement- and belief-learning processes, is poorly identified too. Therefore, it should be unsurprising that the nonparametric Wilcoxon test cannot detect the difference between a reinforcement learning and belief learning model in an SMP game. After all, such a test operates with “less structural knowledge” than an EWA estimation would: It cannot be that it is able to detect something that a (correctly specified) EWA estimation cannot. For distinguishing between learning models with precision parameters λ , QRE equilibria away from equal mixing over strategies is a necessary design component.

TABLE A1

POWER OF WILCOXON TESTS AGAINST THE HYPOTHESIS OF NO TREATMENT DIFFERENCE (SIMULATIONS—SEE NOTES BELOW): AMP GAMES

Game	Power of test (at 5% significance level) to detect a treatment difference in the proportion of times each player plays the indicated strategy.						Power of test to detect change in proportion (between first third and last third of periods)	
	All periods		First third of all periods		Last third of all periods			
	Row plays Up	Column plays Left	Row plays Up	Column plays Left	Row plays Up	Column plays Left	Row plays Up	Column plays Left
AMP19	0.18	0.99	0.84	0.99	0.85	0.11	0.98	0.79
AMP9	0.12	0.98	0.78	0.99	0.70	0.10	0.94	0.71
AMP4	0.07	0.92	0.59	0.96	0.36	0.08	0.71	0.52
AMP3	0.06	0.83	0.47	0.89	0.24	0.07	0.54	0.41
AMP2	0.05	0.56	0.26	0.63	0.11	0.06	0.26	0.21
SMP (AMP1)	0.06	0.06	0.05	0.05	0.05	0.05	0.04	0.05