

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

by

E. Elisabet Rutström

and

Nathaniel T. Wilcox

## **Abstract:**

If belief elicitation during repeated game play changes game play, using stated beliefs to evaluate and compare theories of strategic action is problematic. We experimentally verify that belief elicitation can alter paths of play in a repeated asymmetric matching pennies game. Belief elicitation improves the goodness of fit of structural models of belief learning, and the prior “inferred beliefs” implied by estimates of such structural models are both stronger and more realistic when beliefs are elicited than when they are not. These effects are confined to players who see strong payoff asymmetries between strategies. We conjecture that this occurs because both automatic, unconscious, evaluative processes and more effortful, conscious, deliberative processes can conflict, and that the latter are enhanced by belief elicitation. 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: June 2007

JEL classification codes: C73, C92, D83

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

**Rutström:** Department of Economics, College of Business Administration, and the Institute for Simulation and Training, University of Central Florida, Orlando, FL 32816; [erutstrom@bus.ucf.edu](mailto:erutstrom@bus.ucf.edu). **Wilcox** (corresponding author): Department of Economics, University of Houston, Houston TX 77204-5019; [nwilcox@mail.uh.edu](mailto:nwilcox@mail.uh.edu); 713-869-6422. 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 players' beliefs about partners frequently shape players' 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 subjects 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.

As economists, we are tempted 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 estimators of any underlying true or "latent" beliefs that agents may or may not be aware of. If strategic actions are in part the product of subconscious latent beliefs, agents may have to estimate these in order to state them. In such a case, stated beliefs need not be better predictors of behavior than inferred beliefs. In fact, the experimental data we present here shows that inferred beliefs can predict

game play better than stated beliefs, contrary to the results of Nyarko and Schotter (2002). The inferred beliefs that we estimate allow for high-frequency variability, a striking aspect of stated beliefs first noticed by Nyarko and Schotter (p. 978).

If belief elicitation does result in a more deliberative consideration of the likely play of partners, then belief elicitation may also move subjects away from relatively automatic emotional or “affective” predispositions favoring specific strategies over others. In effect, game play itself may be altered by belief elicitation. 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). These games are interesting because the Nash Equilibrium often does not predict play very well, even after considerable learning has been taking place. More importantly, the predictions generated by belief-based and non-belief-based learning models differ significantly. 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. With asymmetric payoff opportunities it is likely that automatic strategic predispositions make the player attracted to choose the action with the relatively high payoff. Interrupting such predispositions, which is what appears to happen during belief elicitation, causes a more deliberative choice process that puts more emphasis on the other player’s actions. 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. Motivation, Concepts and Literature.

The potential inferential value of stated beliefs has been understood for some time (Holt 1986 and Manski 2002, 2004),<sup>1</sup> and Nyarko and Schotter (2002) showed that stated beliefs can explain game play far better than certain kinds of inferred beliefs. Contemporary interest in the use of belief elicitation procedures and stated beliefs in experimental studies of game play are therefore quite understandable. Still, it is well-known that belief 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. Thus, belief elicitation procedures are “intrusive” and this may have unintended consequences: Discovering these, and figuring out how to finesse them, if possible, is a necessary step toward confidently exploiting the full inferential potential of stated beliefs.

### 1.1 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 make three 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

---

<sup>1</sup> Both of us have used them in our own past work (McDaniel and Rutström 2002; Austin and Wilcox 2007).

compared with belief-based models.<sup>2</sup> 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 can be formulated rigorously in the following terms. Assuming that the first assumption (play is belief-based) is true, let “latent beliefs” refer to the 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.<sup>3</sup> 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).<sup>4</sup> Cheung and Friedman (1997) is a pioneering example of the inferred beliefs approach. The third assumption is that stated beliefs are better predictors of strategic actions than are inferred beliefs. Notwithstanding Nyarko and Schotter’s (2002) specific instance of support for this assumption, it may not be true in general. In the judgment and decision making literature, experimental studies of stated beliefs in nonstrategic settings document systematic biases in stated beliefs (Lichtenstein, Fischhoff and Phillips 1982). While

---

<sup>2</sup> 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.

<sup>3</sup> Costa-Gomes and Weizsäcker (2006) simply call this the “underlying belief.” We like the term “latent belief” for its continuity with both econometric and psychometric language wherein a “latent variable” is an underlying but unobserved variable which produces some observable outcome—in our case, strategic actions and stated beliefs—according to some stochastic data-generating process that involves the latent variable in question.

<sup>4</sup> Nyarko and Schotter (2002) call this the “empirical belief,” but our experience is that this terminology confuses readers, since both stated and inferred beliefs are “empirical” in the sense of being estimates and/or measurements.

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).

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 psychological processes that she cannot directly observe) to draw such inferences.<sup>5</sup> Since this position resembles that of an econometrician inferring beliefs from observed actions and theoretical identifying restrictions, such as the specification of a latent belief updating process, stated beliefs might better be viewed as estimators of latent beliefs, just as inferred beliefs are. If so, appropriate language for talking about both of them is the language of estimators and their relative statistical properties: In our particular context, we focus on the estimators' predictive content for strategic actions. This resembles the view taken by Costa-Gomes and Weizsäcker (2006),<sup>6</sup> and has some support in psychological research on the role of affect in judgments of risk (e.g. Slovic et al. 2002). In what follows, then, we distinguish between latent beliefs (thought of as the theoretical "beliefs in the

---

<sup>5</sup> Evidence of verbal reports biased by implicit (and incorrect) theoretical identifying restrictions is strong in other areas of psychology, particularly where subjects are asked to provide retrospective information or to evaluate the effect of treatments; see e.g. Nisbett and Wilson (1977) and Ross (1989). A recent economic application of these notions in the realm of program evaluation can be found in Smith, Whalley and Wilcox (2006).

<sup>6</sup> Costa-Gomes and Weizsäcker (2006) explicitly view stated beliefs as the result of a random utility model of belief statement under a proper scoring rule given underlying latent beliefs, and econometrically analyze them that way.

head” specified by some belief-based theory of strategic actions), and two classes of estimators of latent beliefs—inferred beliefs and stated beliefs.

## 1.2 Why might belief elicitation change game play?

There are several reasons for thinking that belief elicitation may influence subjects to play more in accord with belief-based models of learning, or otherwise alter decision making in ways that are traceable to changes in belief-based theoretical constructs, such as prior latent beliefs. First, in the absence of belief elicitation, some subjects may not think about partners at all: They might instead learn by reinforcement, or perhaps may play in accord with inclinations to action produced by relatively automatic and unconscious processes, such as differential affective or emotional reactions to strategies and their potential payoffs. Belief elicitation might then cause more conscious deliberation about partners amongst some of those subjects, and make their play resemble belief-based play more closely than otherwise. Second, even among subjects who would engage in belief-based thinking about play without it, belief elicitation may enhance or deepen their thinking in a manner that enhances the predictive ability of their latent beliefs. This could also enhance the descriptive success of belief-based models of learning and play.<sup>7</sup> Third, in situations where players are permanently matched, greater focus on partners could induce a more forward-looking or “sophisticated” style of thinking (Camerer, Ho and Chong 2002) that incorporates predictions about how partners, who are also learning, will respond in future periods to actions the player takes now. Fourth, Offerman, Sonnemans and Schram (1996) point out that when both stated beliefs and game play are rewarded, risk averse subjects can state

---

<sup>7</sup> Both the first and second possibilities can be motivated in other ways. From a “decision cost” or “cognitive production” viewpoint, strategy choice and beliefs are “joint products” when subjects use belief-based algorithms to choose their strategy; so rewarding the belief formation process is like subsidizing the use of belief-based choice algorithms relative to other choice algorithms. From the perspective of demand effects, asking subjects to state beliefs may make them think that experimenters wish them to think about beliefs when making decisions.

beliefs so as to insure against ex post strategic mistakes, and vice versa. This makes the action and belief statement a joint decision. Finally, Croson (1999) argues that belief elicitation enhances the perception of various reasons for preferring one strategy over another, which may then alter game play.

### 1.3 Previous studies of belief elicitation.

Previous experiments have addressed the possibility that belief elicitation procedures in games have unintended consequences, but results are decidedly mixed and reveal no obvious pattern of results. For instance, Croson (1999, 2000) finds that in certain one-shot 2x2 games with a dominant strategy or dominance-solvable equilibrium, belief elicitation procedures make subjects more likely to play dominant or iteratively dominant strategies. Yet Costa-Gomes and Weizsäcker (2006) find no evidence of this in one-shot 3x3 dominance-solvable games,<sup>8</sup> and neither did Wilcox and Feltovich (2000) in a repeated multi-player prisoner's dilemma game with a dominant strategy. While Croson found a significant decrease in public good contributions with a motivated belief elicitation mechanism, Gächter and Renner (2006) found exactly the opposite (a significant increase in contributions). Nelson (2003) reports that in dictator games, dictators give less to counterparts following the elicitation of their beliefs as to how much that counterpart will give in the role of dictator (but with a different player). In a one-shot intergroup public goods game, Erev, Bornstein and Wallsten (1993) found that with belief elicitation, subject behavior in the game did not depend on payoffs from provision of the public

---

<sup>8</sup> Costa-Gomes and Weizsäcker (2006) was designed with a similar purpose to ours, that of looking for behavioral effects from eliciting beliefs. Nevertheless, they did not investigate the effects on learning in their experiments, but rather the effects on game play in one-shot settings. Further, it is possible that the fact that they used fairly significant payoffs in the belief elicitation task may have given subjects incentives to use stated beliefs as a way to insure against the risk in the action choices (see Offerman, Sonnemans and Schram 1996), rather than to respond truthfully.

good, while it did without belief elicitation. In the realm of repeated 2x2 games with a unique mixed strategy equilibrium, Nyarko and Schotter (2002) find no effect of belief elicitation.

These existing results are obviously mixed. Many of these experiments differ in important respects from the one presented here. Poor statistical power may have been a problem in some of these experiments, so we take steps (described shortly) to design our own experiment with adequate statistical power. Aside from this, we will also argue that the mixed results frequently reflect differences between the type of games researchers choose to study. In particular, we suspect that in some games, or for some players in a game, relatively automatic affective processes will produce essentially similar strategic inclinations as more consciously deliberative judgment processes. If belief elicitation procedures primarily enhance 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.

## 2. Experimental design.

### 2.1 Overview.

In our experiment, 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 the experimental design to enhance our chance of finding significant treatment effects, using conservative nonparametric two-sample tests, when a treatment difference is present. With the data in hand, our empirical strategy first uses conservative nonparametric tests (the same tests examined in power planning) to establish significant treatment differences in our data; if 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 sources of treatment differences.

## 2.2 Power planning.

Salmon (2001) shows that power, meant to distinguish between models of learning in games, is an important and sometimes neglected aspect of experimental designs. Because of this, we use Monte Carlo simulations to choose game payoffs, the number of repetitions  $T$  played by each pair of players in each treatment, and the number of pairs  $M$  in each treatment, so as to arguably provide for adequate statistical power to detect between-treatment effects when present. In both these simulations and the experimental design, we 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,<sup>9</sup> 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. Erev and Roth's (1998) discussion of predictions of various learning models suggested that relatively large differences in predictions might emerge in asymmetric matching pennies or "AMPX" games, where "X" refers to the degree of asymmetry in the row player's payoffs.<sup>10</sup> 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

---

<sup>9</sup> 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.

<sup>10</sup> AMP games only have a mixed strategy equilibria in which players are indifferent between all strategies, and so might have poor behavioral properties in the neighborhood of such equilibria. But stochastic choice generalizations of Nash equilibrium (and stochastic learning models) do not generally predict play at (or converging to) the mixed strategy Nash equilibrium of an AMPX game, and actual play in those games is rarely near it (McKelvey, Palfrey and Weber 2000). Additionally, Bracht and Ichimura (2001) note that if observed choice probabilities are nearly equal for all strategies, there is a potentially severe identification problem in estimations of many adaptive learning models. Our simulations indicated that in the AMPX game and design we settled on, neither of these potential problems are an issue. Under the two data-generating processes we used in our design simulations, choice probabilities for players are not expected to be near either 0.5 or the mixed strategy Nash equilibrium, nor are they in our actual data.

simulations represent a large, theoretically important difference between a belief-based model and an alternative. Assuming that behavior with and without belief elicitation produces a similarly large effect (for whatever reason), the simulations tell us roughly what design gives us a good chance of detecting such a large difference. The simulations suggested that  $T = 36$  repetitions of an AMP19 game with  $M = 40$  subject pairs in each treatment would give us good power against the null of no treatment difference. For more detail on the simulations, see Rutström and Wilcox (2006). 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. Because the quadratic scoring rule is widely used

by experimenters, including Nyarko and Schotter (2002), we also use it for our SR treatment.<sup>11</sup> Couching this 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  $L_t = 0$ . The maximum payoff for each stated belief in the SR treatment (\$0.10) is deliberately low to make belief statement less interesting in expected payoff terms than strategy choice itself, which is typical of designs like this using a scoring rule.<sup>12</sup>

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

---

<sup>11</sup> In theory, proper scoring rules need to be designed exactly for the risk preferences of the judge, but almost all experiments use the quadratic scoring rule used here which, strictly speaking, is incentive compatible only for risk neutral judges. Separate elicitation of the risk attitudes in the population from which we recruited our subjects indicate that they are indeed risk averse. Harrison, Johnson, McInnes and Rutström (2005) report a moderate degree of risk aversion for this population. One good feature of the small 10-cent range of possible payoffs at stake in our quadratic scoring rule (and many other experimenters use a similarly small range) is that risk attitudes are highly unlikely to matter much over such a small payoff range. However, alternative procedures that account for differences in risk attitudes are conceivable; for an example see Andersen, Fountain, Harrison and Rutström (2006).

<sup>12</sup> This is important because a belief elicitation procedure that is "too interesting" in expected payoff terms can convert a nonzero-sum game like AMP19 with a prior belief elicitation procedure into a game where the game play and belief elicitation tasks might not be approached as separate tasks. For example, suppose the maximum earnings of the column player from the belief elicitation procedure is \$0.40 (for reports of either 100% or 0% chance of row playing Up in the event that the stated belief is correct). In such a case column could guarantee a payoff in the belief elicitation task of \$0.40 by "conditioning" the row player to always choose Up by always choosing Left herself. Since this is more than double his possible expected payoff in AMP19 as written above, such a "joint strategy" across the belief elicitation and game play tasks could tempt many column players in this instance.

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.

#### 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.<sup>13</sup> 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,<sup>14</sup> with scoring rule rewards for accuracy. The screen showed a table with eleven rows representing 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.

---

<sup>13</sup> 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>.

<sup>14</sup> Row players chose a probability that their column player partner would choose left, while column players chose a probability that their row player partner would choose up.

### 3. Statistical Tests and Econometric Models.

We first use Wilcoxon two-sample tests to establish significant treatment effects.

Nonparametric tests are ideal for doing this in a conservative manner, 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  $P_{t+1}$  that row plays up in period  $t+1$  is

$$(1) \quad P_{t+1} = (1 + \exp[-\lambda E_t(\Delta\Pi_{t+1})])^{-1}, \text{ where } E_t(\Delta\Pi_{t+1}) = 19B_{t+1} - (1 - B_{t+1}).$$

Following Camerer and Ho (1999), we call the estimable parameter  $\lambda \in \mathbb{R}^+$  the sensitivity parameter (also variously referred to as the precision, scale or “Luce” parameter by others). It represents a row player’s sensitivity to expected payoff differences between her two strategies—here,  $E_t(\Delta\Pi_{t+1}) = 19B_{t+1} - (1 - B_{t+1})$ , where  $B_{t+1}$  is row’s latent belief that her column partner will play left in period  $t+1$ .

In equation (1),  $B_{t+1}$  is 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 the specification of an estimator of latent beliefs. As

noted earlier, we regard stated beliefs  $\tilde{B}_{t+1}$  as one estimator of latent beliefs. Setting  $B_{t+1} = \tilde{B}_{t+1}$  yields the stated belief model that plays a central role in Nyarko and Schotter (2002), and later we will examine a version of this model for comparative purposes. However, the first model we estimate and discuss is instead an inferred belief model that sets  $B_{t+1} = \hat{B}_{t+1}$ , where  $\hat{B}_{t+1}$  is what we call an inferred belief. The inferred belief is, in part, estimated from some assumed specification of latent belief updating. Usually, these specifications are based solely on a player's observations of her partner's past behavior and a handful of parameters, which 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:

$$(2) \quad \hat{B}_{t+1} = \frac{\Lambda_t}{\Gamma_t}, \text{ where } \Lambda_t = \gamma\Lambda_{t-1} + L_t, \Lambda_1 = \gamma\hat{B}_1\Lambda_0 + l_1, \text{ and } \Gamma_t = \gamma\Gamma_{t-1} + 1,$$

$$\Gamma_1 = \gamma\Gamma_0 + 1, \text{ and } \Gamma_0 = \Lambda_0 \quad \forall t > 1.$$

The parameters of this model are  $\gamma \in [0,1]$ ,  $\hat{B}_1 \in [0,1]$  and  $\Lambda_0 \in [0, (1-\gamma)^{-1}]$ . The numerator accumulates past experiences of the partner playing left and the denominator accumulates all past experiences, whatever the partner played. All past experiences are discounted geometrically at rate  $\gamma$ . If players regard partners as nonstationary stochastic processes, which is sensible if partners are learning too, such discounting is reasonable (Cheung and Friedman 1997; Fudenberg and Levine 1998); it may also simply reflect decaying memory. Finally, prior beliefs and their strength at the outset of play are  $\hat{B}_1 \in [0,1]$  and  $\Lambda_0$ , respectively. Though estimable prior beliefs 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 a

separate initial strength parameter 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 fairly 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 may also partly reflect the well-known phenomenon of overconfidence in judgment (Lichtenstein, Fischhoff and Phillips 1982). Nevertheless, this finding is so striking that our belief updating process ought to be generalized to allow for high-frequency variance of beliefs. In the received GWB model, the  $\gamma$  parameter can produce high-frequency variability only at the expense of rapid discounting of history: 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. Put differently, values of  $\gamma$  close to 1 generate persistent attention to the history of partner play, but correspondingly little high-frequency variability; while values of  $\gamma$  close to 0 imply great high-frequency variability but attention to only the very recent history of partner play. Or, still differently: 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 allow for high-frequency state dependence of an otherwise gamma-weighted belief process. A simple way to do this is to split the gamma-weighted sums  $\Lambda_t$  and  $\Gamma_t$  into state-contingent parts. Let  $s \in S = \{ul, ur, dl, dr\}$  denote the four states of game play (e.g., “*ul*” means that row plays up and column plays left), let  $s_t \in S$  be the state observed in period  $t$ , and let  $I(s_t=s) = 1$  if state  $s$  is observed in period  $t$ , zero otherwise. Then, define these state-conditional gamma-weighted sums:

$$(3) \quad \Lambda_t(s) = \gamma \Lambda_{t-1}(s) + L_t I(s_{t-1} = s) \text{ and } \Gamma_t(s) = \gamma \Gamma_{t-1}(s) + I(s_{t-1} = s) \quad \forall t > 1 ;$$

With initial values in period 1 as:

$$\Lambda_1(ul) = (\gamma \hat{B}_1 \Lambda_0 + L_1) \hat{B}_1 P_1 \text{ and } \Gamma_1(ul) = (\gamma \Gamma_0 + 1) \hat{B}_1 P_1 ;$$

$$\Lambda_1(ur) = (\gamma \hat{B}_1 \Lambda_0 + L_1)(1 - \hat{B}_1) P_1 \text{ and } \Gamma_1(ur) = (\gamma \Gamma_0 + 1)(1 - \hat{B}_1) P_1 ;$$

$$\Lambda_1(dl) = (\gamma \hat{B}_1 \Lambda_0 + L_1) \hat{B}_1 (1 - P_1) \text{ and } \Gamma_1(dl) = (\gamma \Gamma_0 + 1) \hat{B}_1 (1 - P_1) ; \text{ and}$$

$$\Lambda_1(dr) = (\gamma \hat{B}_1 \Lambda_0 + L_1)(1 - \hat{B}_1)(1 - P_1) \text{ and } \Gamma_1(dr) = (\gamma \Gamma_0 + 1)(1 - \hat{B}_1)(1 - P_1) .$$

Note that since  $\sum_{s \in S} \Lambda_t(s) \equiv \Lambda_t$  and  $\sum_{s \in S} \Gamma_t(s) \equiv \Gamma_t \quad \forall t$ , these conditional sums simply split the unconditional sums  $\Lambda_t$  and  $\Gamma_t$  into four parts arising from observations following each of the four possible states of the previous play of the game. Adding a new “conditionality parameter”  $\theta \in [0,1]$ , a “gamma-theta belief process” is then given by:

$$(4) \quad \hat{B}_{t+1}(s_t) = \theta \left( \sum_{s \in S} \frac{\Lambda_t(s)}{\Gamma_t(s)} I(s_t = s) \right) + (1 - \theta) \frac{\Lambda_t}{\Gamma_t} \quad \forall t > 1.$$

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 once-lagged state of game play. From a statistical viewpoint,  $\theta \in (0,1)$  regresses conditional beliefs toward the unconditional gamma-weighted belief.

From a behavioral perspective, gamma-theta beliefs allow for a kind of “second-guessing” 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 plays: Players can learn that new actions by the partner are in part reactions to one's own recent plays. 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, the interpretation of  $\hat{B}_{t+1}(s_t)$  with  $\theta \in (0,1)$  is more subtle than the purely statistical interpretation. 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 also be a more predictable, and hence more exploitable, opponent herself in an AMP game. So while gamma-theta beliefs algebraically behave like conditional probabilities, they are perhaps better interpreted as "decision weights" since  $\theta$  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 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 a deeper introspection or changes the epistemic basis for the belief, this could result in either larger or smaller values of  $\theta$ , depending on how that introspection balances the possibility that the partner may be predictable against the possibility that the partner may exploit one's own predictability.

Another parameter  $\alpha \in \mathbb{R}$  frequently appears in models based on equation (1), usually in the form  $P_{t+1} = F[\lambda E_t(\Delta\Pi_{t+1}) + \alpha]$ . Cheung and Friedman (1997) interpret  $\alpha$  as a bias in strategy

choice not attributable to subjective expected payoff differences between strategies. Battalio, Samuelson and Van Huyck (2001) interpret  $\alpha$  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  $\alpha$  denote a row player's perception of the change in her expected payoff difference in the next period brought about by playing up in the current period rather than down.<sup>15</sup> Suppose also that the row player believes that her column player partner learns according to a GWB model with the same discount rate  $\gamma$  as her own, and believes that this implies a roughly geometrically declining impact of her

---

<sup>15</sup> We are assuming that  $\alpha$  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 knows all of the column player's adaptive learning structure, nor that  $\alpha$  is the change in one-period-ahead payoffs implied by that structure: In other words,  $\alpha$  denotes an expectation, but not necessarily a rational one.

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  $\alpha$ , we modify equation (1) to be

$$(5) \quad P_{t+1} = (1 + \exp[-\lambda[E(\Delta\Pi_{t+1}) + \alpha(1 - \gamma^{T-t})/(1 - \gamma)])^{-1}.$$

We expect  $\alpha < 0$  since players' incentives are strictly opposed in the AMP19 game.<sup>16</sup>

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

We will call equation (5), with the gamma-theta belief updating process given by equations (3) and (4), 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 simply 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 phenomenon, see, for instance, Wooldridge (2002). Although reinforcement learning models and

---

<sup>16</sup> 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.

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.<sup>17</sup>

Because of this, we employ the “mixed random estimator” recommended by Wilcox (2006) to account for heterogeneity.<sup>18</sup> This estimator assumes that  $\ln(\lambda)$  is normally distributed with mean  $\mu_\lambda$  and standard deviation  $\sigma_\lambda$  across subjects,<sup>19</sup> and also adds a normally distributed, mean zero subject-specific additive random effect  $\nu$  with standard deviation  $\sigma_\nu$  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 estimators:  $\lambda$  and  $\nu$  are numerically integrated out of the EGWB likelihood for each player, conditional on  $\mu_\lambda$ ,  $\sigma_\lambda$  and  $\sigma_\nu$ ; the log of this is summed across players in a sample; and this sum is maximized in  $\alpha$ ,  $\gamma$ ,  $B_1^l$ ,  $\Gamma_0$ ,  $\theta$ ,  $\mu_\lambda$ ,  $\sigma_\lambda$  and  $\sigma_\nu$ .

---

<sup>17</sup> In fact, there is a second, indirect conditioning on own lagged dependent variables in all fixed pairing designs whether or not the belief model is directly conditioned on them. This occurs because (for instance) the column partner’s beliefs, and hence choices, are conditioned on lagged choices of the row player. Because of this, the column partner’s choices contain trace information about the row player’s idiosyncratic learning parameters; and the row player’s beliefs are of course conditioned on the past choices of the column player. It is this conditioning on variables containing information about cross-sectional heterogeneity that causes the pooled estimation bias in the first place. This indirect conditioning is likely to be a less serious source of bias, since idiosyncratic parameter information is diluted by a pair of randomizations (in the initial choice of the row player, and then again in subsequent choices of the column player conditioned on them), whereas direct conditioning (as in the EGWB model) only passes idiosyncratic parameter information through a single randomization before subsequent choices are conditioned on own lagged choices. See Wilcox (2006) for details.

<sup>18</sup> Since it is the gamma-theta belief process that contains own lagged dependent variables, and since  $\theta$  controls the degree of conditioning on them, the expectation would be that pooled estimation will bias estimates of  $\theta$  upward. In fact, the pooled MLE of  $\theta$  is about 15-25 percent larger than the random parameters estimate of  $\theta$  in all three treatments, illustrating this expectation and demonstrating the need for random parameters estimation.

<sup>19</sup> Wilcox (2006) shows that heterogeneity of  $\lambda$  is a particularly serious source of bias in pooled MLE estimates. In the interest of robustness checking, we have also specified a 3-parameter gamma distribution for  $\lambda$  and find that this neither fits our data significantly better nor produces different parameter estimates than the lognormal specification.

The gamma-theta belief process is only one of many alternative belief updating processes that might capture high-frequency variance of beliefs. For instance, one might specify an alternative gamma-theta belief process that conditions only on once-lagged choices of partners, rather than on once-lagged states of game play.<sup>20</sup> Or, a term representing the length of recently observed unbroken “runs” of play by the partner, and an associated parameter, could represent ideas players may have about mean reversion of partner choices (or its opposite). We have examined these alternatives as well as the standard gamma-weighted belief model, and two facts are important. First, neither standard gamma-weighted beliefs, nor these alternatives, fit our data as well as the gamma-theta belief process in (3) and (4). This is true across all three treatments, and the improvement in fit is considerable in at least two of the three treatments. Second, the gamma-theta belief model generally results in larger p-values against hypotheses concerning equality of model parameters across treatments—that is, weaker evidence of treatment effects—than either standard gamma-weighted beliefs or these alternatives.<sup>21</sup> Relative to these alternatives, then, it is the least “friendly” model we have examined toward our basic hypothesis that belief elicitation changes play.

## 4. Results

### 4.1 Graphical analysis 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

---

<sup>20</sup> While this less conditional version of the gamma-theta belief model does not fit the data as well as the more conditional version in (3) and (4), it produces estimates of  $\theta$  (and other model parameters) that are very similar to those produced by the more conditional version in (3) and (4).

<sup>21</sup> Moreover, the significant treatment effects identified under all these alternatives include the ones we identify later using the gamma-theta belief process in (3).

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 (2006). Note that these are extremely conservative tests since 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 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,<sup>22</sup> 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 concerning equality of parameters across treatments. 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. These tests reveal a

---

<sup>22</sup> In the power planning simulations, we discovered that very high power for row players was achievable only with very great asymmetry. We decided we were willing to run the risk of relatively weak column player motivation, in order to keep overall experimental costs manageable, and use the high asymmetry of the AMP19 game.

strongly significant difference between the SR and NB model estimates, in agreement with the nonparametric test results in Table 1, and a weakly significant difference between the SR and EC treatments. However, there is no significant difference between the NB and EC model estimates. Thus, the EGWB models suggest that the particularly “intrusive” scoring rule procedure produces structural differences in 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  $\alpha$ ,  $\gamma$  and  $\theta$  do not appear to vary significantly across the treatments. Put differently, neither forward-looking calculation (represented by  $\alpha$ ) nor the dynamics of belief learning (represented by  $\gamma$  and  $\theta$ ) appear to be altered by belief elicitation. In all three treatments,  $\alpha$  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  $\gamma$ , 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  $\theta$  takes values from 0.2 to 0.35 and is highly significant in all treatments.<sup>23</sup> We return to the question of the practical significance of these estimates of  $\theta$  shortly. Significant parametric effects of scoring rule procedures relative to no belief elicitation are confined to three

---

<sup>23</sup> 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  $\chi^2$  with one degree of freedom—all (obviously) producing miniscule p-values.

areas. These are: (1) More realistic and stronger initial conditions of beliefs (represented by  $\hat{B}_1$  and  $\Lambda_0$ ); (2) greater sensitivity to expected payoff differences (represented by the expected value of  $\lambda$ ) 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  $\sigma_v$ ). We now discuss each of these in turn.

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.<sup>24</sup> It is clear that there is a radical difference in this fit in early periods (periods 1-12); however, by the late periods (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  $\hat{B}_1$  are 0.89, 0.62 and 0.45, with standard errors of 0.29, 0.15 and 0.09, respectively. It is evident that 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.

---

<sup>24</sup> 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.

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.<sup>25</sup> Parametrically, the estimated expected value of  $\lambda$ —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.<sup>26</sup>

Third,  $\sigma_v$  is significantly greater in the SR treatment than the NB treatment. We interpret  $v$  as representing relatively persistent but unmodeled differences between the behavior of individual row players. Under this interpretation, larger  $\sigma_v$  implies greater unexplained but systematic variance of row player behavior in the SR treatment. We interpret this as meaning that conscious belief processes encouraged by scoring rule procedures are more variable across players than the largely unconscious ones that prevail without them. If conscious attentional resources vary a great deal across individuals while unconscious ones do not, 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 2003), this difference is wholly sensible. Since attentional resources are likely generated through experience and training, their heterogeneity is consistent with findings elsewhere that cognitive capital

---

<sup>25</sup> To put this 5.62% poorer model fit in the NB treatment in 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 percent. If this is an interesting difference in fit between two different models of behavior, then a difference that is twice as large for a single model, game, and player type, across two different treatments, should be interesting too.

<sup>26</sup> While it is customary to view  $\lambda$  as a measure of sensitivity to expected payoff differences, 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  $\lambda$  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  $\lambda$  can signal greater predictive force of inferred belief updating processes.

varies greatly across individuals (Camerer and Hogarth (1999)). And, as pointed out by Wilcox (2006), heterogeneity can be a serious inferential nuisance in comparisons of learning models, so any design choice that increases unexplained heterogeneity is (other things equal) a potentially dangerous one.

#### 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 (gamma-weighted beliefs, gamma-theta beliefs and stated beliefs) of latent beliefs in the SR treatment. The gamma-theta beliefs are those implied by the estimates of  $\alpha$ ,  $\gamma$ ,  $\theta$ ,  $\hat{B}_1$  and  $\Lambda_0$  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 player's stated beliefs in the SR treatment, and therefore (obviously) includes none of the three inferred belief process parameters  $\theta$ ,  $\hat{B}_1$  or  $\Lambda_0$ . However, it does include the parameters  $\alpha$  and  $\gamma$  to represent forward-looking calculation by the players as in equation (5). 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  $\theta$  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  $\bar{B}_i$ , for each player  $i$ .<sup>27</sup> 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 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

---

<sup>27</sup> 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  $\tau$  could also be used, but it throws away potentially useful magnitude information represented by the rank of  $\bar{B}_i$  that is used by the Spearman coefficient, replacing it with purely ordinal information about  $\bar{B}_i$ .

than either gamma-theta or gamma-weighted beliefs, though the comparison is statistically convincing only for gamma-theta beliefs.<sup>28</sup> 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 either of 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. This may explain part of the difference between our findings and those of Nyarko and Schotter (2002). Nyarko and Schotter’s game results in play that is, over most of the time series, not too different from equal frequencies of play of both strategies by both players. This means that the frequency of changes in play, relative to any game where play frequencies are well away from equal mixing, must be relatively large. An i.i.d. Bernoulli sequence with  $p = 0.5$  produces high-frequency changes in outcomes 50% of the time. In our game, in the SR treatment, the overall proportion of “up” plays by row players is about 0.64, well away from 0.5. An i.i.d Bernoulli sequence with  $p = 0.64$  produces high-frequency changes in outcomes 46% of the time. While not an enormous difference in this particular comparison, the example illustrates that if stated beliefs have a comparative advantage in predicting changes in play relative to frequencies of

---

<sup>28</sup> 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), which might explain the difference between their findings and ours. We rather doubt this, however. Using Nyarko and Schotter’s data, and using their 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), we find that rounding their players’ stated beliefs to the nearest 0.1 prior to estimation has no significant effect on parameter estimates and actually improves the fit of that model (the log likelihood per subject, over the sixty periods of play, is  $-36.9285$  without rounding versus  $-36.8888$  with rounding), albeit very slightly.

play, their overall predictive advantage would tend to be maximal in games with roughly equal observed mixing, such as Nyarko and Schotter's game, and smaller in games with highly unequal mixing.

## 5. Discussion and Conclusions

Is belief elicitation a good tool for inquiry into economic behavior in general, and behavior in games in particular? It is easy to see why belief elicitation is, in principle, such an attractive mechanism: An unbiased and efficient estimator of latent beliefs, constructed without reference to (or effect on) observed actions in games and without any structural model, would be extremely valuable for empirical study of behavior in games (Manski 2002) and in many other contexts (Manski 2004). Nyarko and Schotter's (2002) results encouraged hopes for the value of stated beliefs, and suggested that belief elicitation had no effect on subsequent actions, and graphically demonstrated just how inadequate inferred beliefs can sometimes be.

Unfortunately, our results show that none of this is robust. In our experiment, 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. Finally, we find that stated beliefs can actually be poorer predictors of observed player actions than inferred beliefs. Note that this alone cannot decisively weigh against Manski's (2002, 2004) position, which is built not only upon the predictive value of stated beliefs, but also on their potential strict exogeneity from observed

actions. It is this latter hoped-for property that would give stated beliefs a decisive econometric advantage in model identification. Unfortunately, if belief statement itself 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.

As none of this matches what Nyarko and Schotter (2002) observe, what might explain the difference? First, we suspect that statistical power plays a role. Rutström and Wilcox (2006) show that, under the same assumptions used to plan power for our experiment, the expected power of Nyarko and Schotter’s design to detect belief elicitation effects on game play is very poor. Critically, our experiment was designed from the outset to create relatively high power against the null hypothesis that belief elicitation has no effect on game play. This was not true of Nyarko and Schotter’s experiment: Their comparison of otherwise identical treatments without belief elicitation was an after-the-fact add-on undertaken as part of editorial processes.<sup>29</sup> Put differently, Nyarko and Schotter originally set out to compare stated belief and inferred belief models and chose their game with that central purpose in mind—not to ask whether belief elicitation changes game play. Therefore, it should not be surprising to discover after the fact that their specific game produces no evidence that belief elicitation changes game play while we and Croson (1999, 2000), who chose games to look for this, do.

Beyond the issue of statistical power, we are willing to make one conjecture on the basis of our own findings and those of Croson (2000) and Nyarko and Schotter (2002). Recall that we only find belief elicitation effects for row players, but not column players, in our AMP19 game. 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.

---

<sup>29</sup> Personal communication with Andrew 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).

We suspect that pronounced affective asymmetries between strategies create relatively automatic emotional predispositions favoring one strategy over another, and that these predispositions tend to determine actions, especially early ones in a repeated game, in the absence of more conscious deliberative judgment. Suppose then that belief elicitation procedures enhance deliberative judgment: Then one should expect that belief elicitation changes game play only when relatively automatic predispositions conflict with the recommendations of more deliberative judgments. Because affective and emotional tags on strategies should evolve with payoff experience in a repeated game, such an account is potentially consistent with the gradual vanishing of differences in play between subjects who do and do not state beliefs that we see in our own study. Details of this affective account need to be worked out in future research. While we think it is a promising explanation for the results of Nelson (2003) and Erev, Bornstein and Wallsten (1993), the results of Costa-Gomes and Weizsäcker (2006) are not consistent with it. It may be that 3x3 games, such as those examined by Costa-Gomes and Weizsäcker, are too

complex to allow for relatively automatic affective predispositions that strongly favor any one strategy.

Our finding that stated beliefs can be worse predictors of actions than inferred beliefs needs to be replicated by others, 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. This is obviously disappointing especially since strictly exogenous measures of beliefs could play a decisive identification role in games (Manski 2002). Experimenters understand their potential role for efficient experimental designs and hypothesis testing. 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, effects play to a lesser degree than does the scoring rule.<sup>30</sup> Of course, this elicitation is a rather blunt instrument for continuous underlying true beliefs. Yet this fact gives us hope that some belief elicitation mechanism might be informative enough about latent beliefs, but unintrusive enough not to qualitatively alter play, and thus give researchers a suitable "belief instrument" for identification and theory tests.

As Holt (1986) points out, any belief measure monotonically related to true beliefs could be a suitable instrument for those beliefs—depending on the particular experimental context. One potentially useful possibility is a suitably fine-grained verbal response scale. This technique

---

<sup>30</sup> But Croson (1999, 2000) documents belief elicitation effects even with this procedure.

has been extensively studied by statisticians and psychologists for many years; Mosteller and Youtz (1990) provide a useful review. Wallsten, Budescu and Zwick (1993) offer practical advice on the construction of these scales, as well as comparisons between them and proper scoring rules based on numerical responses. Subjects prefer verbal response scales to numerical ones, suggesting that verbal responses may be less cognitively intrusive than numerical ones. And although the two methods each have their own small advantages over one another, Wallsten, Budescu and Zwick conclude that, in general, there is nothing that decisively recommends one over the other. 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.

## 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 & E. Elisabet Rutström. 2006. Eliciting beliefs in the laboratory. University of Canterbury (New Zealand) Department of Economics working paper.
- Austin, A. and N. Wilcox. 2007. Believing in economic theories: Sex, lies, evidence, trust and ideology. Economic Inquiry 45:502-518.
- 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
- Camerer, Colin and Robin Hogarth. 1999. The effects of financial incentives in experiments: a review and capital- labor-production framework, Journal of Risk and Uncertainty 19, 7-42.
- 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.
- Costa-Gomes, Miguel A. and Georg Weizsäcker. 2006. Stated beliefs and play in normal-form games. London School of Economics working paper.
- 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.
- McDaniel, T. and R. Rutström. 2001. Decision making costs and problem solving performance. Experimental Economics 4:145-61.
- 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.
- Offerman, Theo, Joep Sonnemans and Arthur Schram. 1996. Value orientation, expectations and voluntary contributions in public goods. Economic Journal 106:817-45.
- 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. 2006. 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.
- Slovic, Paul, M. Finucane, E. Peters and D. G. MacGregor. 2002. The affect heuristic. In T. D. Gilovich, D. W. Griffin and D. Kahneman, eds., Heuristics and Biases: The Psychology of Intuitive Judgment. Cambridge, U.K.: Cambridge University Press.
- Smith, J., A. Whalley and N. Wilcox. 2006. Are program participants good evaluators? University of Michigan Department of Economics, working paper.
- Wallsten, Thomas S., David V. Budescu and Rami Zwick. 1993. Comparing the calibration and coherence of numerical and verbal probability judgments. Management Science 39:176-90.

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.

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		
	<b>NB versus SR</b>	NB versus EC	SR versus EC	<b>NB versus SR</b>	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>

Table 2.

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

Likelihood Ratio Tests of Equivalence of Parameter Vectors Across Treatments						
Treatments Compared	Extended GWB Model w/ lognormal $\lambda$ 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
$\alpha$	-0.296 (0.090)	-0.477 (0.146)	-0.358 (0.0701)	equal $\alpha$	$\chi^2_1 = 0.86$ $p = 0.35$	$\chi^2_1 = 0.33$ $p = 0.56$
$\gamma$	0.950 (0.0261)	0.897 (0.0352)	0.898 (0.0261)	equal $\gamma$	$\chi^2_1 = 0.0005$ $p = 0.98$	$\chi^2_1 = 2.53$ $p = 0.11$
$\hat{B}_1$	0.892 (0.285)	0.615 (0.152)	0.446 (0.0901)	equal $\hat{B}_1$	$\chi^2_1 = 1.54$ $p = 0.21$	$\chi^2_1 = 5.21$ $p = 0.022$
$\Lambda_0$	0.329 (0.445)	1.24 (0.661)	5.38 (3.82)	equal $\Lambda_0$	$\chi^2_1 = 3.77$ $p = 0.052$	$\chi^2_1 = 6.49$ $p = 0.011$
$\theta$	0.352 (0.112)	0.214 (0.0692)	0.297 (0.0691)	equal $\theta$	$\chi^2_1 = 0.72$ $p = 0.40$	$\chi^2_1 = 0.24$ $p = 0.62$
$\mu_\lambda$	-1.77 (0.174)	-1.74 (0.169)	-1.28 (0.137)	equal $\mu_\lambda$	$\chi^2_1 = 4.41$ $p = 0.036$	$\chi^2_1 = 4.32$ $p = 0.038$
$\sigma_\lambda$	0.414 (0.184)	0.640 (0.227)	0.496 (0.182)	equal $\sigma_\lambda$	$\chi^2_1 = 0.42$ $p = 0.51$	$\chi^2_1 = 0.15$ $p = 0.69$
$\sigma_v$	0.459 (0.115)	0.577 (0.190)	0.844 (0.158)	equal $\sigma_v$	$\chi^2_1 = 2.48$ $p = 0.12$	$\chi^2_1 = 5.67$ $p = 0.017$
$E(\lambda)^a$	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 $\mu_\lambda$ and equal $\sigma_\lambda$	$\chi^2_2 = 3.69$ $p = 0.16$	$\chi^2_2 = 6.55$ $p = 0.038$
Log Likelihood per 40 subjects	-866.78	-867.49	-820.68	equal $\hat{B}_1$ and $\Lambda_0$	$\chi^2_2 = 3.80$ $p = 0.15$	$\chi^2_2 = 6.91$ $p = 0.032$

<sup>a</sup>Because  $\mu_\lambda$  and  $\sigma_\lambda$  are the mean and variance of  $\ln(\lambda)$ , the actual mean and variance of  $\lambda$  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).