ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

An International Society for the Advancement of Economic Theory in its Relation to Statistics and Mathematics

http://www.econometricsociety.org/

Econometrica, Vol. 78, No. 3 (May, 2010), 1143–1154

DOES EXPERIENCE TEACH? PROFESSIONALS AND MINIMAX PLAY IN THE LAB

JOHN WOODERS

Eller College of Management, University of Arizona, Tucson, AZ 85721-0108, U.S.A.

Econometrica, Vol. 78, No. 3 (May, 2010), 1143–1154

NOTES AND COMMENTS

DOES EXPERIENCE TEACH? PROFESSIONALS AND MINIMAX PLAY IN THE LAB

BY JOHN WOODERS[1]

Does expertise in strategic behavior obtained in the field transfer to the abstract setting of the laboratory? Palacios-Huerta and Volij (2008) argued that the behavior of professional soccer players in mixed-strategy games conforms closely to minimax play, while the behavior of students (who are presumably novices in strategic situations requiring unpredictability) does not. We reexamine their data, showing that the play of professionals is inconsistent with the minimax hypothesis in several respects: (i) professionals follow nonstationary mixtures, with action frequencies that are negatively correlated between the first and the second half of the experiment, (ii) professionals tend to switch between under- and overplaying an action relative to its equilibrium frequency, and (iii) the distribution of action frequencies across professionals is far from the distribution implied by minimax. In each respect, the behavior of students conforms more closely to the minimax hypothesis.

KEYWORDS: Minimax, mixed strategy, Nash equilibrium.

1. INTRODUCTION

SEVERAL RECENT PAPERS have established that the on-the-field behavior of sports professionals in strategic situations requiring unpredictability is consistent with the minimax hypothesis and its generalization to the theory of mixed-strategy Nash equilibrium. (See Walker and Wooders (2001) and Hsu, Huang, and Tang (2007) for tennis, and Chiappori, Levitt, and Groseclose (2002) and Palacios-Huerta (2003) for soccer.) This raises an important question: Does expertise in strategic behavior obtained in a familiar setting, for example, the tennis court or the soccer field, transfer to an unfamiliar one? If it does, a key implication is that the nature of the subject pool is a critical ingredient of whether results obtained in the laboratory are useful for predicting behavior in the field.

Following a literature using professionals as subjects in experiments, Palacios-Huerta and Volij (2008) recruited professional soccer players and students to play two mixed-strategy games—a "penalty kick" game they introduced and the well known O'Neill (1987) game.[2] They reported extraordi-

[1]I have benefited from many discussions with Mark Walker about mixed-strategy play, and my thoughts on this topic have been influenced by those discussions. I am grateful to Diego Moreno, Amnon Rapoport, Jason Shachat, Matt Van Essen, Mark Walker, and Myrna Wooders for helpful comments and to Matt for outstanding research assistance. I am also grateful to Ron Oaxaca, Jonah Gelbach, and Keisuke Hirano for comments on statistical issues. A co-editor and four anonymous referees provided valuable comments.

[2]See, for example, Cooper, Kagel, Lo, and Gu (1999), Garratt, Walker, and Wooders (2004), List (2003), and Alevy, Haigh, and List (2007).

 DOI: 10.3982/ECTA7970

nary results: the play of soccer professionals conformed remarkably closely to the behavior predicted by the theory, whereas the play of students did not. The finding for professionals is surprising since a long line of experimental research has demonstrated that student subjects do not play equilibrium in mixed-strategy games.[3]

The present paper reexamines the Palacios-Huerta and Volij (henceforth PH-V) data and argues that the behavior of soccer professionals is inconsistent with the minimax hypothesis in several respects: (i) professionals follow nonstationary mixtures, with action frequencies that are negatively correlated across the first and second half of the data, (ii) they tend to switch between halves between under- and overplaying an action relative to its equilibrium frequency, and (iii) the distribution of action frequencies across players is far from the distribution implied by minimax. In each respect, the behavior of students conforms more closely to the minimax hypothesis.

Paradoxically, this reexamination is motivated by the fact that in one respect, the actual play of professionals is *too close* to the theoretically expected play. The minimax hypothesis applied to the O'Neill game calls for each player to choose cards (*1*, *2*, *3*, or *J*(oker)) according to an independent and identically distributed (iid) mixture that assigns probability 0.4 to the *J* card, and probability 0.2 to each of the non-*J* cards.[4] In PH-V's experiment, while 80 *J*'s are expected after 200 plays, the probability a player chooses between 79 and 81 *J*'s is only 0.171.[5] Hence among the 40 subjects who were professionals, we expect to find only 6.85 for whom actual play is this close to expected play. Remarkably, we find 16. Such an outcome is extremely unlikely under minimax play.[6] The professionals' empirical card frequencies exhibit the same excessive closeness to the theoretically expected frequencies for each of the non-*J* cards as well.

PH-V hypothesized that professionals do not literally follow the iid minimax mixture, but rather they "... try to 'match' some probabilities they have, consciously or unconsciously, in mind." In other words, professionals choose cards as though a law of small numbers applies, keeping their empirical card frequencies close to the minimax expected frequencies.

To investigate this hypothesis, we partition the data in the simplest way, into the first and the second half. If professionals act to match the minimax frequencies, then their empirical card frequencies should exhibit a similar conformity to the minimax frequencies in each half as they do in the data overall. More-

[3]Important experimental studies of mixed-strategy play include O'Neill (1987), Rapoport and Boebel (1992), Mookherjee and Sopher (1994), Ochs (1994), Shachat (2002), and Rosenthal, Shachat, and Walker (2003). See Camerer (2003) for an in-depth survey.

[4]See Wooders and Shachat (2001) for an analysis of repeated binary outcome games.

[5]To avoid the ace bias noted in Brown and Rosenthal (1990), PH-V followed Shachat (2002) in labeling the strategies red, brown, purple, and green in their experiment, but used O'Neill's original labeling in describing their results.

[6]The odds that 16 or more players choose between 79 and 81 *J*'s is about 1 in 1900.

over, the striking difference found between professionals and students should also hold in each half.

We find no evidence that professionals match minimax frequencies. Even though tests based on only half the data have less power, the hypotheses that professionals choose cards individually (or jointly) according to the minimax frequencies are rejected at too high a rate to be consistent with the minimax hypothesis. They are rejected at similar rates for professionals and students.

The essential difference between professionals and students is how their play changes between halves. The minimax hypothesis requires that each player choose cards according to the same iid mixture at each round; hence, the frequencies with which a card is chosen in each half are statistically independent. The hypothesis of independence is rejected for professionals, for each card in the O'Neill game, in favor of negative correlation. For students, in contrast, this hypothesis is only rejected for the J card. In this respect, the behavior of students conforms more closely to the theory.

A related implication is that if a subject happens to play a card with, say, less than its equilibrium frequency in the first half (i.e., he "underplays" it), this has no bearing on the likelihood he will underplay it in the second half. We show that professionals who underplay a card in the first half tend to overplay it in the second half (and vice versa). Students, in conformity with the theory, exhibit no such tendency.

The tendency of professionals to change their play across halves has a powerful effect: it causes their overall card frequencies to be too closely clustered around the expected card frequencies. The Kolmogorov–Smirnov (KS) goodness of fit test applied to the professionals' overall card choices resoundingly rejects minimax play for each card individually and for all four cards jointly; the distribution of card frequencies across professionals is far from that implied by the minimax hypothesis. The same test yields a rejection only for the J card for students.

2. THE PH-V EXPERIMENT

The penalty kick (P-K) game is a stylized representation of a penalty kick in soccer, where the kicker (row) chooses whether to kick left (A) or right (B), and the goalie (column) simultaneously chooses whether to cover left (A) or right (B). In both games of PH-V's experiment, the winner receives 1 Euro and the payoff numbers are the probability that the row player wins.

	A	*B*
A	0.60	0.95
B	0.90	0.70

P-K Game

	1	*2*	*3*	*J*
1	0	1	1	0
2	1	0	1	0
3	1	1	0	0
J	0	0	0	1

The O'Neill Game

Eighty professional soccer players—40 kickers and 40 goalies—participated. For the P-K game, 40 professionals in 20 fixed pairs, with a kicker in the row role and a goalie in the column row, played 150 rounds. For the O'Neill game, another 40 professionals were paired in the same fashion, but played 200 rounds. A total of 160 college students participated, half with soccer experience and half without. To sharpen the contrast, we focus on the 80 students without soccer experience. Like the professionals, 20 fixed pairs of students played 150 rounds of the P-K game and 20 fixed pairs played 200 rounds of the O'Neill game. In both games, the subjects played 15 practice rounds. Subjects were not told the number of rounds to be played.

Comparing Professionals and Students in the O'Neill Game

Tables I and II show the empirical card frequencies of professionals and students, respectively, in the first half and the second half of the O'Neill game. Considering the overall data, PH-V found that the null hypothesis that a subject chooses the J card with probability 0.4 is rejected at the 5% level for only 3 professionals, but for 9 students (2 rejections are expected).[7] For the non-J cards, the minimax binomial model (MBM) is rejected at the 5% level in only 4 instances for professionals, but in 13 instances for students (6 rejections are expected). Applying the Pearson goodness of fit test to each subject's choice of all four cards, the null hypothesis that the subject chose cards according to the minimax mixture is rejected at the 5% level for only 2 professionals, while it is rejected for 7 students. The joint hypothesis that all 40 players choose their cards according to the equilibrium mixture is not rejected for professionals (p-value 0.988), while it is rejected at the 1% level for students (p-value 0.006).

To examine whether the PH-V data are consistent with professionals choosing cards to "match" the minimax frequencies, we partition the data into the first and last 100 rounds.[8] If professionals are indeed matching frequencies, then the close conformity of their overall empirical card frequencies to the expected frequencies should also be found when each half of the data is considered in isolation.[9] The tables show this is not the case. In particular, we note the following observations:

- The null hypothesis that a player chooses J according to the MBM is rejected at roughly three times the expected rate for both professionals and students.

[7]These results are reported in Tables X and XIV, respectively, in PH-V and on the right hand side of Tables I and II here.

[8]Besides being the obvious one, this split has several other advantages: (i) It gives the PH-V matching hypothesis its best chance since professionals might not match frequencies over very short intervals. (ii) The statistical tests have the same power on each half and hence are comparable.

[9]Subjects did not know the number of rounds to be played and hence could not match the minimax J frequency by choosing, say, 30% J's in the first half and 50% J's in the second half.

TABLE I
PROFESSIONALS PLAYING O'NEILL[a]

Pair	Player[b]	First Half 1	First Half 2	First Half 3	First Half J		Second Half 1	Second Half 2	Second Half 3	Second Half J		Overall 1	Overall 2	Overall 3	Overall J	
1	C	23	23	24	30**		16	14	18	52**	†					
	R	19	28*	24	29**	†	19	17	34**	30**	‡			**	**	‡
2	C	20	22	23	35		20	19	27*	34				*		
	R	23	22	22	33		18	21	27*	34					*	
3	C	25	21	23	31*		14	14	18	54**	‡					
	R	25	25	14	36		17	14	26	43						
4	C	18	20	21	41		11**	17	24	48	†	**				
	R	19	19	18	44		24	22	18	36						
5	C	22	21	16	41		18	18	26	38						
	R	20	17	28*	35		16	22	13*	49*						
6	C	27*	18	17	38		14	19	24	43						
	R	21	20	16	43		21	21	21	37						
7	C	23	21	19	37		18	17	22	43						
	R	20	24	13*	43		23	19	13*	45				**		
8	C	26	15	24	35		19	15	17	49*			*			
	R	22	23	19	36		17	20	20	43						
9	C	20	16	20	44		21	20	21	38						
	R	22	21	21	36		15	18	22	45						
10	C	21	25	16	38		18	14	27*	41						
	R	14	17	7**	62**	‡	21	19	27*	33					**	
11	C	27*	22	23	28**	†	23	18	18	41		*				
	R	21	21	12**	46		20	17	22	41						
12	C	17	23	20	40		22	17	21	40						
	R	16	18	23	43		24	22	16	38						
13	C	21	23	17	39		18	20	21	41						
	R	17	21	18	44		26	16	21	37						
14	C	20	31**	22	27**	‡	21	27	17	35			**		**	‡
	R	17	12**	18	53**	‡	20	25	23	32*						
15	C	21	23	24	32*		21	14	16	49*						
	R	19	19	13*	49*		24	21	21	34						
16	C	23	20	16	41		16	13*	19	52**	†				*	
	R	17	19	14	50**		24	20	25	31*						
17	C	22	27*	18	33		23	16	23	38						
	R	16	28**	21	35		25	18	17	40						
18	C	20	29**	19	32*		21	20	23	36					*	
	R	25	16	17	42		17	23	19	41						
19	C	21	21	18	40		13*	20	17	50**						
	R	19	20	18	43		22	24	29**	25**	‡				*	
20	C	22	21	21	36		15	20	15	50**						
	R	21	19	21	39		18	23	19	40						
	C	439**	442**	401	718**		362**	352**	414	872**						
	R	393	409	357**	841*	‡	411	402	433*	754**						
	Overall	832	851**	758*	1559		773	754*	847*	1626						

[a]Symbols: ** and * denote rejection of the minimax binomial model for a given card at the 5% and 10% level, respectively. ‡ and † denote rejection of the minimax multinomial model at the 5% and 10% level.

[b]C and R denote Column and Row, respectively.

TABLE II

STUDENTS PLAYING O'NEILL[a]

Pair	Player[b]	First Half 1	First Half 2	First Half 3	First Half J		Second Half 1	Second Half 2	Second Half 3	Second Half J		Overall 1	Overall 2	Overall 3	Overall J	
1	C	16	19	23	42		12**	22	29**	37	†	**		**		†
	R	19	28**	22	31*		26	26	17	31*			**		**	‡
2	C	15	19	17	49*		22	22	25	31*						
	R	19	16	14	51**		22	20	18	40						
3	C	20	21	20	39		26	17	25	32						
	R	16	24	21	39		24	17	22	37						
4	C	17	18	20	45		18	18	25	39						
	R	13*	16	15	56**	‡	16	27*	16	41		**		*	**	‡
5	C	17	20	17	46		22	15	22	41						
	R	11**	19	22	48*		16	19	25	40		**				†
6	C	25	24	21	30**		18	22	25	35					**	
	R	16	19	27*	38		21	24	20	35						
7	C	20	12**	19	49*		20	18	14	48*			*		**	†
	R	21	20	22	37		25	17	21	37						
8	C	20	20	15	45		17	25	22	36						
	R	22	21	17	40		17	24	16	43						
9	C	15	22	19	44		25	21	20	34						
	R	14	19	24	43		16	24	16	44		*				
10	C	27	15	25	33		23	22	17	38		*				
	R	33**	27	20	20**	‡	23	25	20	32		**	**		**	‡
11	C	15	28*	21	36		30**	24	20	26**	‡		**		**	‡
	R	19	14*	19	48		20	21	17	42						
12	C	20	19	13*	48		23	25	22	30**						
	R	25	21	17	37		31**	21	19	29**	‡	**			**	‡
13	C	16	18	17	49*		24	22	25	29**						
	R	17	26	18	39		18	13*	21	48*						
14	C	19	19	16	46		20	20	25	35						
	R	18	20	23	39		16	26	29**	29**	‡			**	*	†
15	C	17	21	17	45		22	20	23	35						
	R	14	19	17	50**		14	23	23	40		**				
16	C	23	21	17	39		22	22	15	41						
	R	21	22	22	35		28*	17	16	39						
17	C	19	21	18	42		23	20	22	35						
	R	22	14	20	44		17	18	20	45						
18	C	25	21	17	37		18	20	19	43						
	R	26	21	20	33		27*	21	17	35		**			*	†
19	C	28*	25	19	28**	‡	23	23	20	34		*			**	‡
	R	25	19	24	32*		35**	18	14	33	‡	**			**	‡
20	C	24	21	17	38		19	18	24	39						
	R	24	16	14	46		15	25	19	41						
	C	398	404	368*	830		427	416	439**	718**					*	
	R	395	401	398	806	‡	427	426	386	761*	†					‡
	Overall	793	805	766	1636	†	854**	842*	825	1479**					*	‡

[a]Symbols: ** and * denote rejection of the minimax binomial model for a given card at the 5% and 10% level, respectively. ‡ and † denote rejection of the minimax multinomial model at the 5% and 10% level.

[b]C and R denote Column and Row, respectively.

It is rejected at the 5% level for 7 professionals and 6 students in the first half, and 7 professionals and 5 students in the second half. Only 2 rejections in each half are expected under the null.

- For the non-*J* cards, the MBM is rejected at the expected rate for professionals and students in both halves. There are 6 and 4 rejections for professionals and students, respectively, in the first half; there are 3 and 6 rejections for professionals and students, respectively, in the second half. A total of 6 rejections (2 rejections per non-*J* card) are expected at the 5% level in each half.
- The null hypothesis that a player chooses all four cards according to the minimax multinomial model (MMM) is rejected at the 5% level (by the Pearson goodness of fit test) for 3 professionals and 3 students in the first half, and for 3 professionals and for 4 students in the second half.
- The joint hypothesis that all the players in a given role choose a given card according to the MBM is rejected more frequently for professionals than for students. In the first half, it is rejected at the 5% level for professionals in the column role for the *1*, *2*, and *J* cards and in the row role for the *3* card. The same null is not rejected for students, for either role or any card. In the second half, the joint MBM is rejected for professionals in the column role for the *1*, *2*, and *J* cards and for the row role for the *J* card. For students, it is rejected only for column playing the *3* and *J* cards.
- The joint hypothesis that all the players in a given role choose all four cards according to the MMM is rejected at the 5% level for the row role, for both professionals and students in the first half; it is not rejected for either professionals or students for either role or when the roles are combined in the second half.

These results show that in each half there is little difference between professionals and students in terms of the number of 5% rejections of either the MBM or the MMM, in contrast to the overall data.[10] For both groups, we obtain more than the expected number of rejections of the MBM for the *J* card. The joint MBM is rejected more frequently than expected, especially for professionals.

Also inconsistent with the PH-V matching hypothesis, there is no tendency in either half for the actual play of professionals to be too close to expected play. For example, in each half we expect 9.6 of 40 players to choose either 39, 40, or 41 jokers. In the first and second half there were, respectively, 7 and 8 professionals whose actual play was this close to expected play, while the analogous numbers for students are 6 and 10.

[10]See Wooders (2008) for the *p*-values for these tests. The numbers of rejections of the MBM and MMM continue to be remarkably similar for professionals and students if the data are split by thirds or quarters.

3. NEGATIVE CORRELATION AND SWITCHING

If professionals are not matching frequencies, then why are their overall empirical card frequencies too close to the minimax frequencies? We will show that for each card, the frequencies with which it is played by professionals in each half are negatively correlated. Indeed, professionals tend to switch between halves between under- and overplaying each card relative to its minimax mixture.

In the equilibrium of the repeated O'Neill game, each player chooses cards according to the same iid mixture at each round. An implication is that the frequency with which a card is played in the first half is statistically independent of the frequency with which it is played in the second half. Table III shows that this is not the case. It reports, for each card, the value of the Spearman rank correlation coefficient R between these frequencies.[11]

The choice frequencies of professionals are negatively correlated between halves for each card, with the null hypothesis of independence rejected for each card at the 5% significance level. For students, the same null is not rejected for any of the non-*J* cards. For the *J* card, however, independence is rejected at the 5% significance level, with the first- and second-half frequencies positively correlated.

A related implication of the minimax hypothesis is that if a subject plays a card with, say, a frequency below its equilibrium frequency in the first half, this has no bearing on the likelihood he will underplay the card in the second half; that is, a player is equally likely to switch as not. There were 36 professionals who either strictly under- or overplayed the *J* card in both halves. Of these, 25 switched between underplaying the *J* card and overplaying it, and hence the hypothesis that a player is equally likely to switch as not is rejected (p-value 0.029). For the *1*, *2*, and *3* cards there are, respectively, 32, 29, and 37 professionals who strictly underplayed or overplayed the card in both halves. With 26,

TABLE III

SPEARMAN RANK CORRELATION COEFFICIENTS

	Professionals			Students		
	R	t	p-Value	R	t	p-Value
J	−0.3195	−2.08	0.04432	0.3159	2.05	0.04731
1	−0.5804	−4.39	0.00009	0.2882	1.86	0.07064
2	−0.3688	−2.45	0.01900	−0.0160	−0.10	0.92087
3	−0.3463	−2.28	0.02831	−0.0427	−0.26	0.79627

[11]Under the null hypothesis of independence, the distribution of R is known and hence R can be used to obtain a nonparametric test of independence. (See Gibbons and Chakraborti (2003, pp. 422–431).) The Spearman R is computed using the web page http://faculty.vassar.edu/lowry/corr_rank.html, authored by Richard Lowry, and it corrects for ties in ranks.

24, and 27 switches, respectively, the null hypothesis that the switching probability is 0.5 is rejected for each of these cards as well.[12] In each case, the null is rejected as there are too many switches.

The behavior of students in contrast, is consistent with this implication of the minimax hypothesis. For the *J* card, there were 36 students who strictly under- or overplayed the card in both halves and, of these, only 13 switched. For the *1*, *2*, and *3* cards, respectively, 17 of 34, 16 of 31, and 16 of 31 students switched. The null hypothesis is not rejected for any card.

Switching and the Consequences for Overall Play

If a subject overplays a card in the first half, but underplays it in the second half, then the sample variance will tend to be too small relative to the sample variance under equilibrium play. Consider, for example, the 200-time repeated matching pennies game. If a subject chooses H with probability $0.5 + \gamma$ in the first half and $0.5 - \gamma$ in the second half, with $0 \leq \gamma \leq 0.5$, then 100 heads are expected, but the variance is only $50(1 - 4\gamma^2)$. Given a collection of subjects whose play varies in this fashion, there will tend to be too many having chosen approximately the expected number.[13] Under the null hypothesis of a fair coin, applying the binomial test there will tend to be too many subjects with large p-values (and too few with small p-values).

Table IV shows the results of applying the KS test to the empirical distributions of 40 p-values obtained when the randomized binomial test is applied to each card individually and the Pearson goodness of fit test is applied all four cards jointly. For each card, the KS test resoundingly rejects the joint null hy-

TABLE IV

OVERALL KS TESTS OF CONFORMITY TO $U[0, 1]$

	Professionals		Students	
	KS	p-Value	KS	p-Value
J	1.477	0.025429	1.484	0.024464
1	2.332	0.000038	1.239	0.092692
2	1.917	0.001285	1.160	0.135789
3	1.693	0.006456	1.071	0.201243
1–2–3–J	2.434	0.000014	1.292	0.071098

[12]Although each test is meaningful on its own, the tests are not independent. A subject who switches from under- to overplaying *J* must switch from over- to underplaying at least one of the non-*J* cards.

[13]This also introduces a bias in favor of too few runs and away from the negative serial correlation (see Brown and Rosenthal (1990) and O'Neill (1991)) exhibited by students in the O'Neill game.

pothesis that professionals play according to the MBM; each empirical cumulative distribution function (cdf) has too many large p-values.[14] In contrast, the same null is not rejected for any one of the non-J cards for students. The hypothesis that all four cards are jointly chosen according to the minimax mixture is rejected for professionals, but it is not rejected for students. The empirical distribution of card frequencies for professionals is far from that implied by the minimax hypothesis.

4. DISCUSSION

The differences between professionals and students found here in the O'Neill game are exhibited in the P-K game as well. Table V reports the value of the Spearman R between the first and second half (i.e., 75 of 150 plays) frequencies of right (B). Since there are only 20 subjects in each role, this test has lower power for the P-K game than for the O'Neill game, where, since the equilibrium mixture was the same for both roles, we could pool the data for both row and column. Nonetheless, for professionals, the null of independence is rejected at the 10% level and just fails to be rejected at the 5% level, while independence is not rejected for students.

Professionals in the column role tend to switch between under- and overplaying right, with 14 players switching. Under the null that professionals in the column role are equally likely to switch as not, the probability of 14 or more switches is 0.0577; the null just fails to rejected at the 10% level. Only 7 professionals in the row role switch, which is less than the number expected. This is the only instance in which professionals exhibit no tendency to switch; it is also the only instance in which PH-V rejected the joint minimax model for professionals.

Table VI shows the results of applying the KS test to the distributions of p-values obtained when the randomized binomial test is applied to the left–right choices in the P-K game. For professionals, minimax play is rejected for both roles, with the empirical cdf for the column role exhibiting too many large

TABLE V

SPEARMAN RANK CORRELATION COEFFICIENTS

	Professionals			Students		
	R	t	p-Value	R	t	p-Value
Row	−0.4153	−1.94	0.06821	−0.0982	−0.41	0.68053
Column	−0.4354	−2.05	0.05522	−0.2155	−0.94	0.36153

[14] See Wooders (2008) for a description of the randomized binomial test and for figures showing the empirical cdf's.

TABLE VI

OVERALL KS TESTS OF CONFORMITY TO $U[0, 1]$

	Professionals		Students	
	KS	p-Value	KS	p-Value
Row	1.403	0.039058	2.732	0.000001
Column	1.508	0.021187	2.509	0.000007

p-values. Minimax play is resoundingly rejected for students, with too many small p-values for both roles.

Since there is no obvious reason why professionals (but not students) would follow nonstationary mixtures, efforts to replicate the PH-V results seem especially important. Our results suggest that it may be useful to focus on the whether professionals follow nonstationary mixtures. Although not an exact replication, Levitt, List, and Reiley (2010) found no evidence that the behavior of American Major League Soccer players conforms more closely to minimax than does the behavior of students in the O'Neill game.

REFERENCES

ALEVY, J., M. HAIGH, AND J. LIST (2007): "Information Cascades: Evidence From a Field Experiment With Financial Market Professionals," *Journal of Finance*, 62, 151–180. [1143]

BROWN, D., AND R. ROSENTHAL (1990): "Testing the Minimax Hypothesis: A Reexamination of O'Neill's Experiment," *Econometrica*, 58, 1065–1081. [1144,1151]

CAMERER, C. (2003): *Behavioral Game Theory, Experiments in Strategic Interaction*. Princeton: Princeton University Press. [1144]

CHIAPPORI, P., S. LEVITT, AND T. GROSECLOSE (2002): "Testing Mixed Strategy Equilibria When Players Are Heterogeneous: The Case of Penalty Kicks in Soccer," *American Economic Review*, 92, 1138–1151. [1143]

COOPER, D., J. KAGEL, W. LO, AND Q. L. GU (1999): "Gaming Against Managers in Incentive Systems: Experimental Results With Chinese Students and Chinese Managers," *American Economic Review*, 89, 781–804. [1143]

GARRATT, R., M. WALKER, AND J. WOODERS (2004): "Behavior in Second-Price Auctions by Highly Experienced eBox Buyers and Sellers," Working Paper 02-04, University of California at Santa Barbara. [1143]

GIBBONS, J., AND S. CHAKRABORTI (2003): *Nonparametric Statistical Inference*. New York: Dekker. [1150]

HSU, S., C. HUANG, AND C. TANG (2007): "Minimax Play at Wimbledon: Comment," *American Economic Review*, 97, 517–523. [1143]

LEVITT, S., J. LIST, AND D. REILEY (2010): "What Happens in the Field Stays in the Field: Professionals Do Not Play Minimax in Laboratory Experiments," Working Paper 07-11, University of Arizona. *Econometrica* (forthcoming). [1153]

LIST, J. (2003): "Does Market Experience Eliminate Market Anomalies," *Quarterly Journal of Economics*, 118, 41–71. [1143]

MOOKHERJEE, D., AND B. SOPHER (1994): "Learning Behavior in an Experimental Matching Pennies Game," *Games and Economic Behavior*, 7, 62–91. [1144]

OCHS, J. (1994): "Games With a Unique Mixed Strategy Equilibria: An Experimental Study," *Games and Economic Behavior*, 10, 202–217. [1144]

O'NEILL, B. (1987): "Nonmetric Test of the Minimax Theory of Two-Person Zero-Sum Games," *Proceedings of the National Academy of Sciences*, 84, 2106–2109. [1143,1144]

——— (1991): "Comments on Brown and Rosenthal's Reexamination," *Econometrica*, 59, 503–507. [1151]

PALACIOS-HUERTA, I. (2003): "Professionals Play Minimax," *Review of Economic Studies*, 70, 395–415. [1143]

PALACIOS-HUERTA, I., AND O. VOLIJ (2008): "Experientia Docent: Professionals Play Minimax in Laboratory Experiments," *Econometrica*, 76, 71–115. [1143]

RAPOPORT, A., AND R. BOEBEL (1992): "Mixed Strategies in Strictly Competitive Games: A Further Test of the Minimax Hypothesis," *Games and Economic Behavior*, 4, 261–283. [1144]

ROSENTHAL, R., J. SHACHAT, AND M. WALKER (2003): "Hide and Seek in Arizona," *International Journal of Game Theory*, 32, 273–293. [1144]

SHACHAT, J. (2002): "Mixed Strategy Play and the Minimax Hypothesis," *Journal of Economic Theory*, 104, 189–226. [1144]

WALKER, M., AND J. WOODERS (2001): "Minimax Play at Wimbledon," *American Economic Review*, 91, 1521–1538. [1143]

WOODERS, J. (2008): "Does Experience Teach? Professionals and Minimax Play in the Lab," Working Paper 08-04, University of Arizona. [1149,1152]

WOODERS, J., AND J. SHACHAT (2001): "On the Irrelevance of Risk Attitudes in Repeated Two-Outcome Games," *Games and Economic Behavior*, 34, 342–363. [1144]

Dept. of Economics, Eller College of Management, University of Arizona, 401 McClelland Hall, Tucson, AZ 85721-0108, U.S.A.; jwooders@eller.arizona.edu.

Manuscript received June, 2008; final revision received December, 2009.