The Meaning of Mixing: Individual strategic randomization in games

by

Timothy O'Neill Đặng

University of Arizona

April 28, 2008

Abstract

Mixed-strategy play is an essential part of game theory, whether it comes in the form of minimax, Nash equilibrium, quantal response equilibrium, or non-equilibrium solution concepts such as level-k reasoning. Still, there remains debate about the nature of the mixing. Do individuals randomize, or is there simply heterogeneity in individuals playing pure strategies? Is a strategy's randomness purely in the subjective beliefs of other players? I conduct laboratory experiments in which players can choose to pay a fee to use a randomization device, applied to O'Neill's zero-sum game. If subjects did so, it would show a strict preference for randomization over any available pure strategy. In fact, very few subjects chose to use the randomization device. Subjects' descriptions of their decision process were consistent with the notion of purification.[1][2]

1 Introduction

Mixed strategies are at the core of much game theory–most particularly mixed-strategy Nash equilibrium (MSNE), but also non-Nash solution concepts such as quantal response equilibrium (McKelvey and Palfrey 1995) and level-k reasoning (Stahl and Wilson 1995) (Nagel 1995). However, there remains argument about what "mixed-strategy" means in practice. In most informal analyses, it is at least implicit that individual players are randomizing their actions. However, there is a paradox to this story. If a player is indifferent between strategies, then why should they choose according to particular probabilities? If they aren't indifferent, then why should they put any probability weight on (subjectively) inferior strategies? Most formal work which addresses the paradox interprets mixing as something other than individual randomization.

[1] The experiments were programmed and conducted with the software z-Tree (Fischbacher 2007).

[2] Thanks to Martin Dufwenberg for much helpful discussion, and to Barrio Devo for software testing.

Empirical research on mixed strategies has concentrated on the *how* of mixing: What model best predicts the observed probabilities of actions? Are observed actions statistically independent? A few of these experiments have provided tantalizing hints that game players desire to mix individually. However, they were not designed with this question in mind, so we can't be sure that's what we're observing. Also, since we're not sure whether players are truly motivated to mix, we can't be certain what the results say about particular mixing distributions.

I ran experiments aimed at collecting stronger data on individual mixing. I eliminated the confounding possibility of indifference by charging a fee to use an external randomization device. If subjects are willing to pay the fee to randomize, this shows a strict preference for random play. In that case, their chosen distributions truly show preferred mixes for someone who desires mixing.

There are intuitive–if not always formal–reasons to expect more randomization in repeated games, in which my opponent may otherwise identify and exploit my patterns. I look at this by running treatments with both one-shot and repeated games.

Some theories of mixed strategies are presented in Section 2. Previous experimental work is described in Section 3. Research questions are in Section 4. My experimental design and procedures are explained in Sections 5 and 6.

Section 7 presents the results. A small–but not vanishingly small–number of moves had subjects paying to use the randomization device. Overall, it appeared individual randomization was not important, although it was very important to those who used the device. Subjects did not mix according to maximin, but did improve their mixtures over the course of the repeated games. There was no evidence that subjects were more interested in randomization in a repeated game. Section 8 discusses the results.

2 Theory of Mixing

There are several explanations for what actually occurs in a MSNE, and alternative solution concepts provide their own stories of randomization[3]. Also, some theories motivate a mixed-strategy solution without actual randomization by the players.

For zero-sum games, the MSNE is also the maximin solution. The reasoning behind the maximin strategy is that a player can improve their security level–their minimum possible expected payoff–via hedging. In the equilib-

[3]Evolutionary game theory can lead to either population or individual mixing, but doesn't provide a psychology in between evolutionary pressure and behavior.

rium of a zero-sum game, maximin has a particularly strong attraction since the player is indifferent between the pure strategies. In non-zero-sum games maximin can be applied as very conservative play even though it does not maximize expected payoff.

Although maximin may be the historic origin for mixed strategies, experience in recreational games and teaching students has convinced the author that the intuitive origin for mixed strategies is "keep'em guessing". This intuition has long been part of the discussion of mixed strategies:

> The essence of randomization in a two-person zero-sum game is to preclude the adversary's gaining intelligence about one's own mode of play–to prevent his deductive anticipation of how one may make up one's own mind, and to protect oneself from tell-tale regularities of behavior that an adversary might anticipate. (Schelling 1980) p.175

The majority of recreational games and sports as well as such serious conflicts such as warfare are played with the intent to be unpredictable to one's opponent.[4] As (Oechssler 1997) points out, "Most people understand that in rock scissors paper it would be foolish to become predictable. Thus in order to avoid being outguessed players may intentionally randomize... This argument, of course, applies only if the game is played repeatedly and players choices are observable." The argument doesn't apply to one-shot games because while randomness implies unpredictability, unpredictability need not imply randomness–a player may assume that they start off unpredictable, and only become predictable through observation. Poker writer and World Series of Poker winner Dan Harrington suggests occasionally making an incorrect bet "to avoid giving opponents a read on your style" ((Harrington and Robertie 2004) p.52–53). As a piece of the advice, he suggests choosing odds for each strategy and using the second hand of your watch to randomly select the action.[5] This advice is to engage in true randomization for unpredictability, with *reputation* in mind. A "read on your style" would be harmless in a one-shot game.

In spite of possible motivations of security level or unpredictability, theorists have questioned whether players randomize at all. (Luce and Raiffa 1957) (p. 75) raise the possibility that, "the concept of a mixed strategy is a convenient mathematical tool but it completely fails to be realistic."[6]

Rubinstein ((Osborne and Rubinstein 1994) p.37, in a paragraph marked

[4](Sun Tzu (Giles 1910):"36. He must be able to mystify his officers and men ... and thus keep them in total ignorance. 37. By altering his arrangements and changing his plans, he keeps the enemy without definite knowledge. By shifting his camp and taking circuitous routes, he prevents the enemy from anticipating his purpose.")

[5] (Levitt, List, and Reiley 2007) also note this advice in their experiments involving professioal poker players.

[6]Luce & Raiffa do not advocate this view, but bring this up as one argument, along with the arguments of mixing for unpredictability and mixing for hedging.

as disagreeing with his co-author) says, "[*Matching Pennies*] is classically used to motivate the notion of mixed strategy equilibrium, but randomization is a bizarre description of a player's deliberate strategy in the game. A player's action is a response to his guess about the other player's choice; guessing is a psychological operation that is very much deliberate and not random."

(Aumann and Brandenburger 1995) describe Nash equilibrium as an equilibrium in *beliefs* about others' strategies, rather than in actual strategies. Since the beliefs need not be correct, actual play may be in pure strategies.[7] Another common "solution" to the quandry of mixed strategies is to make mixing a population property rather than an individual strategy.[8] Harsanyi's (Harsanyi 1973) purification theorem is related this idea, that individuals may have very small variations in their perceived values from strategies, so the pure strategy they play (in a Bayesian game) from their point of view is randomized from another player's point of view. The variation may be between individuals or for one individual over time–if they missed their morning coffee, for instance.

This population-mixing idea is also part of the description of other solution concepts such as QRE, where an individual makes mistakes about their payoffs according to some distribution, and level-k reasoning, where an individual best-responds to their perception of others' play, but parts of the population think more deeply about game play than others.[9]

2.1 Rationales for costly mixing

There is no room for costly randomization in MSNE, nor in most other mixed-strategy solution concepts.[10] If a player is indifferent between two strategies, then choosing either pure strategy is strictly better than paying to mix between them. If a player is not indifferent, they should not mix in any case. Any theory which allows for costly mixing must somehow allow for the mixed strategy to be *strictly* preferred to any of the pure strategies.

There are two broad classes of reasons a game player might wish to truly randomize–ambiguity aversion and observability. Various mixing behaviors (such as maximin) could be motivated by one or both.

Ambiguity-averse decision makers are uncomfortable with unknowns, and can behave in ways contrary to expected utility theory. Just as an individ-

[7] (Binmore, Swierzbinski, and Proulx 2001) also talk about "equilibrium in beliefs" but actually mean polymorphic equilibria in population-level probability distributions over actions.

[8] (Nash 1996), in part of his dissertation not included in (Nash 1951), gives average population behavior as a possible interpretation, though still allowing for individual mixing.

[9] Level-k reasoning and other non-equilibrium solution concepts are sometimes defined to include uniform mixing at level-0, or when a player is indifferent over strategies. This is useful to complete the model, but is silent about whether or not the randomization is individual.

[10] This doesn't necessarily mean that a theory which does allow for costly mixing wouldn't result in the same strategies.

ual uncertain of how many red and black balls are in an urn might avoid betting on either (Ellsberg 1961), a game player who is uncertain of their opponent's strategy might not want to commit to a single responding strategy. By randomizing their move, a player can hedge against making a poor choice.

Maximin behavior is a strong form of ambiguity aversion. Weaker forms of ambiguity aversion can lead to a variety of behavior in non-zero-sum games. However, any player who is purely averse to ambiguity will play the maximin strategy in a zero-sum game.[11] (Eichberger and Kelsey 1996) point out that different models of ambiguity as non-additive beliefs can result in a strict preference for randomization, or not.

I am considering simultaneous-move games. However, a player may still have some concern that an opponent can "read minds" and so respond to the actual strategy chosen. This could change the game–in the player's mind–into a dynamic game in which they select a strategy first. Even if a player believes there is only a small chance of mind-reading, that may be enough to make a randomizing strategy optimal.

(Reny and Robson 2004) create a model unifying the purification idea with unpredictability motivations. In their model, player i expects that, with probability t_i player j gets to observe i's strategy before moving. Bayesian heterogeneity (as in purification) comes in by having different players expect observation with different probabilities, while the motivation to be unpredictable comes from the observation itself. This model converges to MSNE as the chance of observation gets small. Unlike purification, however, most players choose to deliberately mix.[12]

In computer-mediated laboratory experiments, it is easy to label such beliefs irrational, but in naturalistic situations there may often be cause to suspect mind-reading. Small "tells" could give away planned strategies, athletes watch each other carefully, and armies employ spies. Rationally or magically, such reasoning could cause players to behave as if unpredictability does require randomness.

MSNE, ambiguity aversion, and timing uncertainty all predict maximin play in a zero-sum game. Maximin can be motivated by strongly pessimistic ambiguity aversion ("Whatever can go wrong, will.") or paranoia concerned with observability ("Stop listening to my thoughts!"). (Luce and Raiffa 1957) draw attention to the dual explanations of mixed strategies as maximin behavior (which makes sense even in games against nature), or the secrecy gained by randomizing:

[11] If a player places value on both expected value and security level, both are maximized only at the maximin equilibrium of a zero-sum game. A player who is partially ambiguity-loving may prefer a different strategy.

[12] This is a simplification for the current context. An appealing feature of their model is that it leads to deliberate mixing in pure competition zero-sum games, but pure strategy Bayesian play à la Harsanyi in coordination games.

> ...it may be sensible to use a mixed strategy as a hedge against extremely unfavorable situations and against the possibility that one's opponent has more insight into one's behavior than anticipated. (Luce and Raiffa 1957) p.73

Because of this common prediction, my experiments will use a zero-sum game (described in Section 5).

3 Experimental studies

Very few experimenters have allowed explicit mixing in games. The six experiments described below took highly varied approaches to mixing. Five of the six (excluding (Camerer and Karjalainen 1994)) have a common theme of investigating the importance of learning. The same experiments were intended to test the way in which players chose to randomize–according to what probability distributions, whether i.i.d., whether stationary, etc.–rather than whether they chose to randomize at all. The exception of (Camerer and Karjalainen 1994) did explicitly look at mixing versus pure strategy, with a focus on timing and ambiguity aversion. As I'll discuss in Section 4, none of the experiment designs is able to show whether players consider randomization strictly better than playing a pure strategy.

In a very early experiment, (Kaufman and Becker 1961) elicited mixed strategies. However, they were concerned with naive subjects learning to play minimax, and took strong steps to encourage such learning. A degree of paranoia was appropriate for subjects as they were playing against simulated opponents who observed the subjects' selected distributions and responded. The simulated opponent did not respond perfectly, but well enough that a subject could never earn more than the minimax value.

(Rapoport and Budescu 1992) had subjects enter a move for each of 100 *Matching Pennies* games. After all 100 moves were chosen by each player, all were played in the given order. Subjects "mixed" in the sense of overall choosing each strategy about half the time, but still displayed serial correlation similar to those who chose a strategy for one game at a time in a repeated-play treatment. In this study, it appears that subjects playing a one-shot game (setting the strategy for 100 "games" in advance) were just as interested in mixing as those playing repeatedly with the same opponent.

(Bloomfield 1994) elicited a sort of "mixed strategies" for a nonzero-sum game with a unique MSNE, but the payoffs were computed as the expected values given a mixture, rather than from specific realized draws from the mixtures. As Bloomfield notes, this actually transformed the original 2x2 game into a 51x51 game with a multitude of mixed-strategy equilibria and a unique pure-strategy equilibrium. Subjects' actions were closer to MSNE predictions when more feedback was provided on their opponent.[13]

(Camerer and Karjalainen 1994) used a variant on *Matching Pennies* in which player 1 wished to mismatch, and player 2 wished to match. Player 1 had the standard **H** or **T** strategies. Player 2 had three strategies: **H**, **T** or randomize ($\frac{1}{2}$**H**, $\frac{1}{2}$**T**)–so had the option to explicitly mix according to experimenter-determined probabilities. In each of two treatments, more than half of player 2's chose to randomize.[14]

(Ochs 1995) had subjects play *Matching Pennies* variants with many periods of random subject rematching–one "straight" zero-sum game and two *Asymmetric Matching Pennies* games. Ten games were played each period, and subjects defined at the beginning of the period how many of the games should be played with each strategy. The results were inconsistent with both MSNE and and maximin play.

The clearest experimental evidence of deliberate randomization comes from (Shachat 2002), who allowed for explicit mixing in repeated play of the O'Neill game (O'Neill 1987), using a computer interface with a "shoe of cards" metaphor. He found that subjects made regular use of of the mixing mechanism, but mostly didn't play minimax strategies and still exhibited serial correlation.[15]

4 Research Questions

We're left with some questions that previous experiments don't address. In each of the studies described in Section 3, subjects chose to explicitly mix, but the designs make any interpretation of that decision uncertain. We want to find out if people value individual mixing in itself. If so, then we'd like to begin to find out *why*. Particularly, does interest in mixing depend on reputation effects? And is mixing consistent with maximin, or do we need to search for new explanations?

I conducted experiments in the spirit of Shachat, but with modifications to sharpen the focus on mixing in general, rather than mixing according to a particular solution concept. The critical difference is that I require the subjects to pay to randomize. Earlier experiments show that players will often choose to mix when the option is available, but since mixing has always been costless it leaves open the possibility that the players were indifferent between mixing and pure strategies. Requiring subjects to pay a fee to use

[13]Bloomfield elicited mixed strategies by allowing subjects to decide how many "times" out of 50 they would play each strategy.

[14]I am reformulating the game slightly from Camerer & Karjalainen's presentation to make it easier to describe briefly.

[15]In the two treatments which allowed mixed strategies, 26.6% of strategies were pure strategies. This was in spite of the possibility that, "while this treatment makes playing mixed strategies cognitively simple, playing a pure strategy is less 'expensive,' as subjects can fill the shoe completely with one type of card via a single key stroke." (p.12, note 8)

the computerized randomization device clarifies whether they place a positive value on random action.

My treatments are designed to answer these three questions:

Question 1: Do players value the ability to randomize actions?

Under most theories of game-play, the players should not pay for the randomization device. (Rubinstein 1991) points out, "[T]he use of mixed strategies is particularly problematic in any situation in which their execution is costly in terms of devoting attention or time. If implementing a mixed strategy is costly for the player, then he will strictly prefer to use any of the pure strategies which appear in the support of the mixed strategy rather than to waste the resources associated with implementing the lottery." While Rubinstein was speaking of MSNE, the argument applies more generally: If a player is indifferent between two pure strategies, then either one is better than paying to randomize. If they are not indifferent, then they should choose the preferred pure strategy.

Question 2: If players do pay to randomize, do they play maximin strategies?

Maximizing security level is the only motivation commonly used in game theory which can justify paying to randomize. Maximin strategies can be justified by either ambiguity aversion or concern with observability. The O'Neill game (or *Four-Card Barry*) is a zero-sum game with binary payoffs.[16] This will help to interpret any results which disagree with maximin predictions. Because of the binary payoffs, risk attitude should be irrelevant. Since it's zero-sum, MSNE and maximin lead to the same prediction. Possible reasoning of ambiguity aversion or countering prediction also predict maximin frequencies.

Question 3: Are players more concerned with randomization in repeated games?

Repeated games have been used regularly when looking at mixed strategies in the laboratory.[17] Unpredictability may have more force as a motivation when play is repeated. Backward-looking, randomization can counter anything an opponent has learned previously. Forward-looking, randomization helps develop a reputation as unpredictable.

[16]The name *Four-Card Barry* was applied to O'Neill's game by Mark Walker and John Wooders.

[17]Field studies of sports behavior have also been of repeated games, perhaps inevitably since if the researcher can observe the history of play, so can opponents.

5 Experiment Design

There were two treatments. The first treatment had One-Shot games. The second treatment had Repeated games with constant subject matching. The core game for both was the same: *Four-Card Barry* shown in Table 1. *Four-Card Barry* has been well-studied experimentally. As a zero-sum, binary payoff game, standard theory gives strong equilibrium predictions independent of risk preferences. However, unlike *Rock-Paper-Scissors* or *Matching Pennies*, the equilibrium strategies are not of uniform probability.

Table 1: Four-Card Barry (O'Neill Game)

	♡		♠		♣		♢	
♡	$20	$0	$0	$20	$0	$20	$20	$0
♠	$0	$20	$20	$0	$0	$20	$20	$0
♣	$0	$20	$0	$20	$20	$0	$20	$0
♢	$20	$0	$20	$0	$20	$0	$0	$20

The game was described to the subjects as a card game.[18] Each player plays a card. The row player wins if the cards match, and loses if they mismatch. However, if either player plays a diamond card the rules reverse, and the row player wins with a mismatch.

The unique Nash equilibrium has each player playing hearts, spades and clubs each 20% of the time, and diamonds 40% of the time. The expected value in equilibrium is $12 for the row player and $8 for the column player.

Players received a showup payment of $6. When they played the game, they had the option of spending $1 of that payment to use the computer randomization device. This was presented as shuffling a deck to draw the card If they paid the $1 to use the device, they were allowed to choose the probabilities assigned to each pure strategy by choosing how many of each card went into their deck. Otherwise they picked any pure strategy for free. After the computer chose a strategy, the opponent saw only the revealed action, not the probability distribution or the decision whether to use the device.

In the Repeated game treatment, players played with a constantly matched opponent for 12 periods.[19] At the end of the experimental session, they were paid for a single, randomly selected period. They decided each period independently whether to use the randomization device, and what pure strategy or deck composition to use.

[18] This description of O'Neill's game was paraphrased from the instructions for (Levitt, List, and Reiley 2007). My complete instructions are in the appendix.

[19] In the first Repeated game session, subjects were told there would be 20 periods. The experiment took longer than expected, so the experiment was stopped at the end of the 14th period. In the following two Repeated game sessions, subjects were told there would be 12 periods, and 12 periods were completed.

This experiment design is inherently one-sided. If we believe people have some ability to randomize in their heads, we can't rule out that possibility in this experiment. If we don't see paying to shuffle, the data are essentially silent. Players may place a value on randomization but achieve it without paying to use an external device. (Throughout the paper, I will make this distinction by using "shuffle" to indicate using the computer randomization device, while "mix" and "randomize" will apply to any method–in one's head, using the computer, or otherwise.) Alternatively, players could truly be playing heterogeneous pure strategies. Unfortunately, in such a case the data won't provide evidence for any theory. Since this is an experiment aiming to identify motivation as well as behavior, I supplement the behavioral data with a post-experiment questionnaire on how subjects made their decisions.[20]

On the other hand, if players do pay to shuffle, we'll know that a mixed strategy can be strictly preferred to pure strategies. Such a result runs counter to the bulk of game theory, and then we can start to reconsider the motives behind mixing, and how they might affect behavior.

6 Experiment Procedures

The experiments were conducted in February and March 2008 in the Economic Science Laboratory (ESL) at the University of Arizona. Subjects were undergraduates, recruited both from Economics classes and ESL's subject database. The One-Shot games were held in two sessions of 24 subjects each, and took approximately a half hour. The Repeated game treatments were held in 3 sessions of 16, 22, and 24 subjects, which took a little less than one hour.

Experiments were conducted using z-Tree, with computerized instructions using the same program.[21] These instructions were followed by a few verbal reminders and an opportunity to ask questions. The end-of-experiment questionnaire was also conducted with z-Tree.

At the end of the Repeated game treatment, a 12-sided die was rolled on an overhead projector so all subjects could see the result. This determined the single period for which subjects were paid.[22]

7 Results

[20] Subjects were also asked their sex, and whether they had had experience with similar games.

[21] Screenshots of the instructions are in the appendix. The z-Tree treatments are available from the author.

[22] For the first session of the Repeated game treatment subjects were told there would be 20 periods, but the experiment was stopped due to time after the 14th period. A 20-sided die was rolled, and re-rolled until a valid number between 1 and 14 came up.

Table 2: **F**requency of play

One-Shot	♡	♠	♣	♢
Row Players (n=24)	29.2% (46.4%)	8.3% (28.2%)	12.5% (33.8%)	50% (51.1%)
Column Players (n=24)	37.3% (46.3%)	25.1% (40.2%)	28.2% (43.1%)	9.4% (28.0%)
Repeated	♡	♠	♣	♢
Row Players (n=388)	21.4% (38.2%)	16.2% (33.7%)	18.6% (36.1%)	43.8% (47.4%)
Column Players (n=388)	21.5% (39.0%)	19.3% (37.3%)	22.4% (39.7%)	36.8% (46.7%)

includes all single-card moves and deck-shuffling moves as fractional moves
standard errors in parentheses

Table 2 shows the overall play frequencies of the 4 cards, treating one card in a deck of a shuffling player as 1/100 of a card played. For the Repeated game, considering the relatively few periods played and limited feedback (subjects see results only from their own game), the frequencies are quite close to the maximin solution. A chi-squared goodness-of-fit test with 90% confidence rejects the maximin frequencies for only 2 row and 2 column players.

With so few periods, a test of serial correlation will have low power. Nonetheless, I ran a run test similar to (Walker and Wooders 2001) and others. With 90% confidence, we can reject 19.4% of subjects choosing their strategy independently.[23]

Table 3: **H**ow many paid to shuffle?

	Row Players	Column Players	Pooled
One-Shot (n=48 pooled)	0%	16.7%	8.3%
90% CI	-	$5.9\% \leq \rho \leq 34.2\%$	$4.2\% \leq \rho \leq 20.8\%$
90% One-Sided	-	$7.5\% \leq \rho$	$6.3\% \leq \rho$
Repeated–overall shuffles (n=776 pooled)	12.1%	9.5%	10.8%
90% CI	$5.7\% \leq \rho \leq 22.4\%$	$3.8\% \leq \rho \leq 18.3\%$	$6.7\% \leq \rho \leq 17.7\%$
90% One-Sided	$7.3\% \leq \rho$	$5.6\% \leq \rho$	$7.3\% \leq \rho$
Repeated–first period (n=62 pooled)	16.1%	19.4%	17.7%
90% CI	$6.6\% \leq \rho \leq 31.0\%$	$8.8\% \leq \rho \leq 34.7\%$	$11.3\% \leq \rho \leq 29.0\%$
90% One-Sided	$8.1\% \leq \rho$	$10.5\% \leq \rho$	$14.5\% \leq \rho$
Repeated–at least once (n=62 pooled)	22.6%	25.8%	24.2%
90% CI	$11.1\% \leq \rho \leq 38.3\%$	$13.5\% \leq \rho \leq 41.8\%$	$17.7\% \leq \rho \leq 33.9\%$
90% One-Sided	$13.0\% \leq \rho$	$15.7\% \leq \rho$	$19.4\% \leq \rho$

ρ is the probability of shuffling.
CI's for "One-Shot" and "Repeated–at least once" Row & Column are exact binomial CI's.
CI's for "Repeated–overall shuffles", and all Pooled are bias-corrected accelerated CI's from bootstrapping.

7.1 Research questions

Question 1: Do players value the ability to randomize actions?

[23] Like the other authors, I treat choice as binary for these tests–either Diamond or not-Diamond.

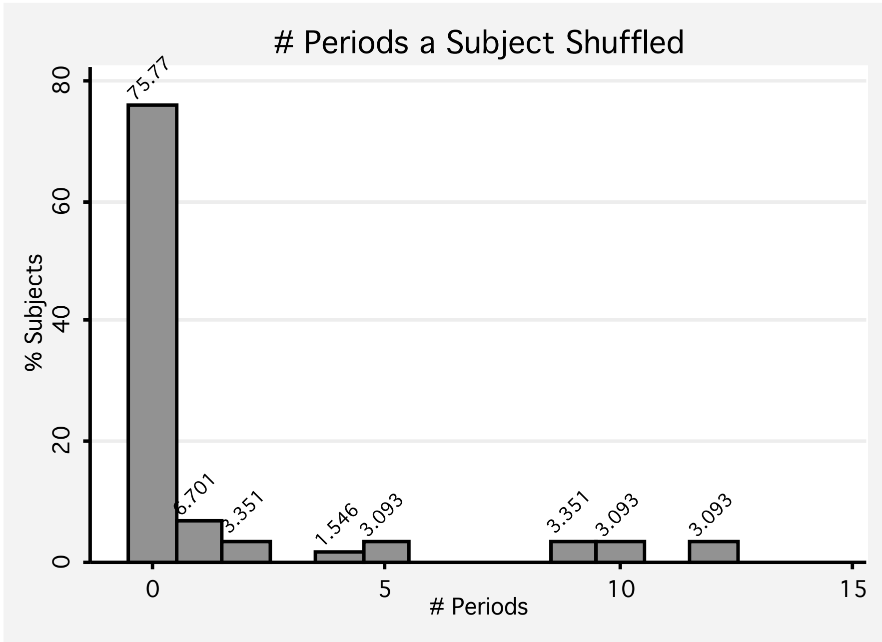

Figure 1: Number of periods a subject shuffled–Repeated game

Most theory makes a strong point prediction–no players should pay to shuffle. Clearly some players chose to do so. Table 3 shows that 8.3% of subjects in the One-Shot treatment shuffled. 10.8% of overall moves in the Repeated-Game treatment were shuffles. In the Repeated games 24.2% shuffled in at least one period, but that is natural given more periods. With 90% confidence, we would expect more than 5% of subjects to shuffle at any given time (6.3% lower bound for the One-Shot game [24] , and 7.3% for any choice in the Repeated game), and nearly 20% (19.4% lower bound [25]) of subjects to shuffle at least once over 12 periods.

Equally clearly, choosing to pay and shuffle was rare. Figure 1 shows how many times different subjects chose to shuffle in the Repeated game. Most subjects chose to never shuffle.[26] These low rates of shuffling are inconclusive. It could be that subjects who chose not to shuffle valued randomization, but were unwilling to pay $1 for it. I will look into why in the questionnaire results.

Question 2: If players do pay to randomize, do they play maximin strategies?

[24] Two of the four subjects who shuffled in the One-Shot game were also coded in the questionnaire as misunderstanding the game. If the subjects who made mistakes were excluded, then 4.8% of subjects shuffled with a 90% CI of 2.4% $\leq \rho \leq$ 16.7%. The lower bound for a one-sided test is 4.8%

[25] There is a spike in shuffling for the first period of the Repeated games. If the first period is omitted, then 22.6% of subjects shuffle at least once with a 90% CI of 16.1% $\leq \rho \leq$ 33.9%. The lower bound for a one-sided test is 17.7%

[26] There is weak evidence in the One-Shot game that row and column players shuffled with different probabilities, but I will pool row & column players for this part of the analysis. Fisher exact test p-values are 10.9% two-sided and 5.5% one-sided for the row players shuffling less often. That is marginally significant, particularly given no *a priori* reason to expect row players to shuffle less. Note also that there is little difference between row & column player frequency of shuffling in the Repeated game.

Table 4 shows the probabilities (deck composition) chosen by players who paid to shuffle. On the whole it is clear they did not shuffle according to maximin strategies. It was the norm for decks to have full support (at least one of each kind of card). In the One-Shot game, every deck composed for shuffling had full support. In the Repeated game, 94% of decks had full support.

If players are shuffling out of ambiguity aversion, the decks should be composed with security level in mind. If they are shuffling to be unpredictable, the decks should be composed to be unexploitable–which also means a high security level. Figure 2 shows the empirical CDF for security level (relative to maximum security level) of shufflers in the Repeated game. From this figure, it doesn't appear that players approached shuffling according to maximin. While there was little difference between row and column players in deciding whether to shuffle, those who did shuffle did so differently. Row players composed decks with higher security levels than column players. In a sense, security level is "easier" for row players, since their security level from a uniform strategy is \$10 (83.3% of the maximum \$12), while it's \$5 for the column players (62.5% of the maximum \$8).

Figure 3 shows security levels by period. While the plot is noisy, it shows a trend toward increasing security level. I ran a Kolmogorov-Smirnov test of the distribution of security levels in periods 1–6 vs. periods 7–12. The test won't be statistically valid, but can serve as a useful descriptive statistic. The test indicates that early period shuffles have a lower security level with p=0.000. While they didn't perfect their deck composition, subjects who chose to shuffle did learn to increase the security of their mixture.

Table 4: **C**omposition of decks for shufflers

One-Shot				
	♡	♠	♣	♢
Row Players (n=0)	-	-	-	-
Column Players (n=4)	24.0% (13.3%)	50.8% (25.5%)	19.0% (11.4%)	6.3% (3.0%)
Repeated				
	♡	♠	♣	♢
Row Players (n=47)	25.3% (7.7%)	25.5% (5.5%)	24.0% (6.2%)	25.3% (7.8%)
Column Players (n=37)	25.0% (9.1%)	24.1% (9.1%)	24.5% (8.1%)	26.4% (18.2%)

Repeated game deck compositions are not independent, including different periods for shuffling subjects. standard errors in parentheses

Question 3: Are players more concerned with randomization in Repeated games?

Table 3 shows that overall, subjects did not shuffle significantly more often in the Repeated games (10.8% of moves) than in the One-Shot games (8.3%

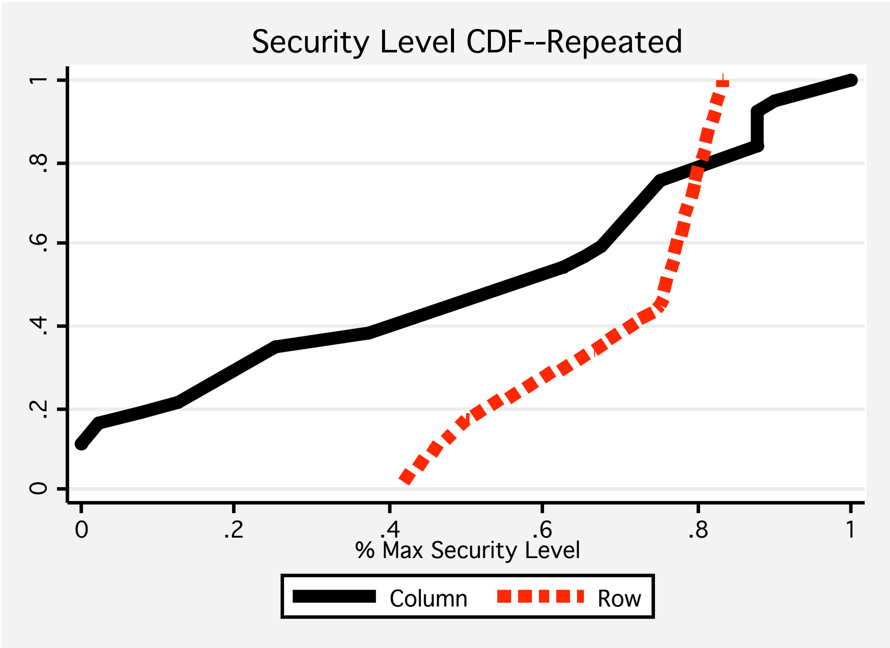

Figure 2: Security levels for shuffling subjects

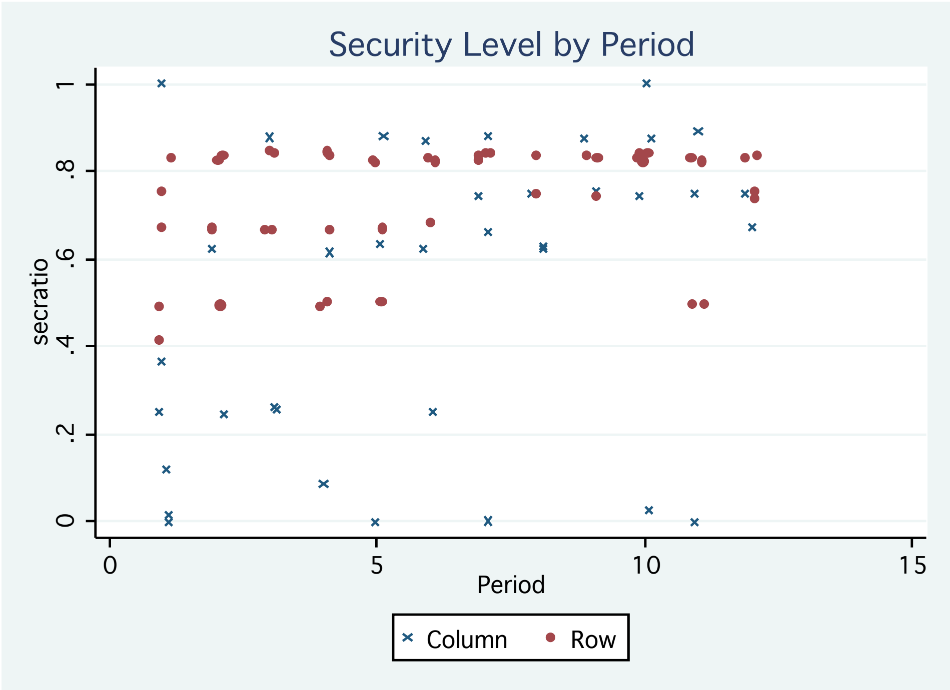

Figure 3: Security levels period-by-period

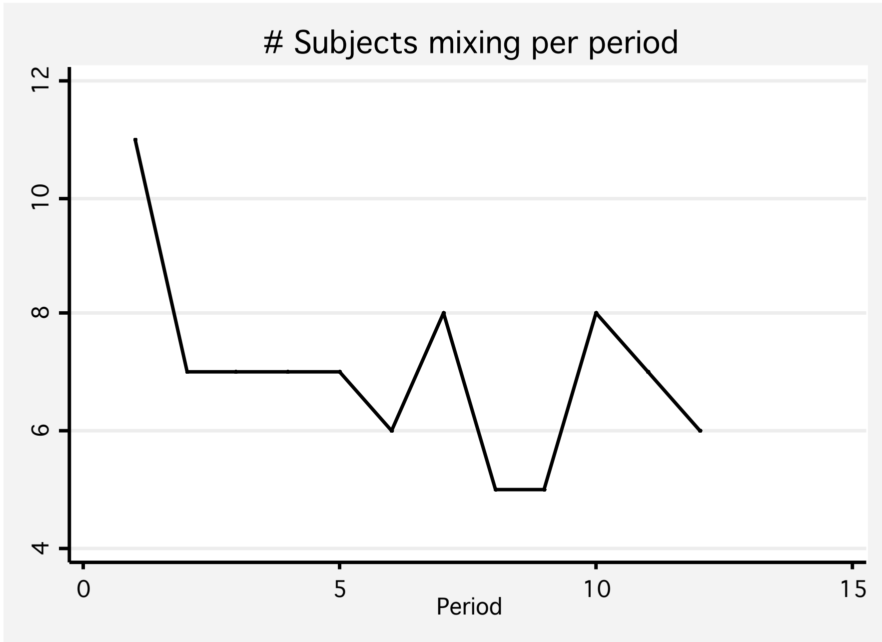

Figure 4: Number of subjects mixing per period–Repeated game

of moves). There is significantly more shuffling in the first period of the Repeated games (17.7%) than in the One-Shot games (p=0.077). This was due to a large spike in shuffling in the first period.

Figure 4 shows this spike in a timeline of the number of subjects who paid to shuffle in each period of the Repeated games. After the first period, there is no particular time trend. There are two rationales for valuing random action in a repeated game. In early periods, a player may wish to mix to establish a reputation as unpredictable; in late periods, a player may mix because they believe they've revealed their strategies. Figure 4 shows no trend in either direction, and a test on shuffling in periods 1–6 vs. periods 7–12 fails to reject a similarity (p=0.487).

The high rate of shuffling in the first period of the Repeated games is something of a puzzle. Given that there was little shuffling in other early periods, it is hard to attribute the shuffling to a concern with developing a reputation. An alternative explanation is that subjects in the Repeated games felt the cost of experimenting with shuffling was relatively low, and so tried it but rejected it after the first period.

7.2 Questionnaire Results

After the experiment was completed, subjects were asked to answer the open-ended question, "Please take a moment and explain how you decided whether to use the computer for shuffling or not." Responses can help clarify "why?"

questions which are challenging in the raw behavioral data. Responses were coded according to five categories.[27]

First, a player may strictly prefer randomization, but decide not to shuffle because they feel they can randomize adequately adequately without paying the fee to shuffle (coded as **Can mix myself**). One subject said:[28]

> **Can mix myself:** "Didn't want to risk losing a dollar. I can choose randomly without the computer's help."

Alternatively, a player might explicitly indicate they don't want to randomize–they believe they know the correct pure strategy (coded as **Want to play pure**). For example:

> **Want to play pure:** "I would rather choose diamond than have the computer shuffle because there was more probablity of our cards not matching than matching."

A player who does value randomization (by shuffling or otherwise) might indicate why. Do they perceive randomization as reducing risk (coded as **Hedging**), or thwarting the opponent's anticipation (coded as **Counter prediction**)? For example, subjects said:

> **Hedging:** "I decided to use the computer for shuffling in the case that I didn't believe my card would beat my opponent's card."
>
> **Counter prediction:** "Well I wanted to shuffle the cards because I wanted very few diamonds, because my odds for winning with diamonds were lower vs. any other random suited card. However, under the chance that the other player choose to play out the cards, then I put 3 diamonds in the deck just maybe to throw them off. Same with the other cards I put a very few cards in each then I choose a bulk card. If I played again I would have probably chosen 100 of 1 kind of card."

Finally, a player might indicate that they consider the one-shot or repeated nature of the game as a factor in deciding to either play a pure strategy or a mixed strategy (coded as **Repeated / One-Shot**). In the Repeated games, most of these responses were part of a subject indicating they wished to play a pure strategy, such as "I did not choose to use shuffling in any round, because I felt I could judge my opponents actions based on if they won the round before."

> **Repeated / One-Shot:** "My selection of a card could be just as random as the computer, since it was only one time. "

[27] Responses were also coded for whether they showed a fundamental misunderstanding of the game, such as believing it was repeated when it was in fact one-shot. 6 subjects (12.8% of nonblank responses) had some kind of misunderstanding in the One-Shot treatment, and none in the Repeated treatment.

[28] Questionnaire responses corrected for spelling. The complete list of responses and how they were coded is in the appendix.

Table 5 shows the percentage of subjects indicating each motivation in their decisions.[29]

Table 5: **C**oded questionnaire responses

	One-Shot (n=47)		Repeated (n=61)	
	Count	Percent	Count	Percent
Can mix myself	2	4.3%	1	1.6%
Want to play pure	19	40.4%	27	44.3%
Hedging	1	2.1%	3	4.9%
Counter prediction	1	2.1%	0	0.0%
Repeated / One-Shot	2	4.3 %	15	24.6%

8 Discussion & Conclusions

Some subjects were willing to pay for the ability to use a computer randomization device, approximately 8% in the One-Shot game, and 24% at some point of the Repeated game. While these numbers are small, they are economically significant given the $1 fee charged to use the device.[30]

Those subjects who did choose to shuffle did not do so according to maximin, but learned to play closer to maximin over the 12 periods of the Repeated games. This rapid learning shows that those who chose to shuffle were very attentive to the quality of their mixture. The requirement that subjects pay to use the randomization device selects for those who take randomization seriously, focuses the attention of those who randomize, or both.

It does not appear that playing a Repeated game makes randomization more appealing. Subjects did not shuffle significantly more in the Repeated games. Players who indicated on the questionnaire that reputation and patterns of play were important, did so to explain why they could predict their opponents' play well enough to pick an optimal pure strategy.

The bulk of the evidence from the experiment comes from the subjects' self-reported motivations. 40% of subjects described motivations inconsistent with an individual desire to randomize. Subjects' descriptions were consistent with purification–they believed they knew the *best* pure strategy to play. Likewise consistent with purification, and also consistent with previous experiments, the overall frequencies of play in the Repeated games were close to the maximin solution.

The results of these experiments are not surprising when taken alone. They do present a puzzle when compared with Shachat's experiments. The bulk of

[29] The percentages shown are based on non-blank responses, and categories are not exclusive. 36.2% of non-blank responses for the One-Shot treatment, and 44.3% of non-blank responses for the Repeated treatment were not coded into any of the categories.

[30] For comparison, the expected value of the game for the column player was $8.

moves in his experiments were shuffled, while the subjects in this experiment indicate they do not wish to shuffle. It may be that, rather than the $1 fee being more costly than the value of shuffling, the fee induced subjects to think differently about the game.

On balance, the results of these experiments show that individual randomization plays a small role. However, this is due to heterogeneity among the players. Faced with a cost to use a randomization device, many subjects realized they would rather play a pure strategy. Those who did choose to shuffle considered it important enough to carefully refine their mixtures over time.

There are empirical implications to whether or not individuals randomize. If players don't value randomness itself, then experience at randomization shouldn't help a player very much. As (Levitt, List, and Reiley 2007) notice, professional poker players (experienced in randomization) appear no more "random" than professional bridge players (a game which requires little randomization).

If players aren't randomizing for unpredictability, then the insights of (Reny and Robson 2004) aren't generally applicable.[31] Qualities of a game which make predictability good or bad–whether a zero-sum or coordination game, for example–should not affect the amount or kind of mixing.

Many studies have shown that people have difficulty being random in individual tasks. By appearance, this difficulty carries over to randomization in games. As (Walker and Wooders 2001) notice, game players in repeated games, "who are attempting to behave truly randomly tend to 'switch too often.'". However, if players are playing their preferred pure strategy in each period, the negative serial correlation is a correlation in beliefs, not in deliberately randomized actions.

References

Aumann, R. and A. Brandenburger (1995). Epistemic conditions for nash equilibrium. *Econometrica 63*(5), 1161–1180.

Binmore, K., J. Swierzbinski, and C. Proulx (2001). Does minimax work? an experimental study. *The Economic Journal 111*, 473.

Bloomfield, R. (1994). Learning a mixed strategy equilibrium in the laboratory. *Journal of Economic Behavior and Organization 25*, 411–436.

Camerer, C. F. and R. Karjalainen (1994). Ambiguity-aversion and non-additive beliefs in non-cooperative games: Experimental evidence. In

[31] Their model could still be very useful in circumstances where there is a real chance of observation.

B. Munier and M. J. Machina (Eds.), *Models and Experiments in Risk and Rationality*. Boston: Kluwer Academic Publishers.

Eichberger, J. and D. Kelsey (1996). Uncertainty aversion and preference for randomisation. *Journal of Economic Theory 7*(1), 31–43.

Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *The Quarterly Journal of Economics 75*(4), 643–669.

Fischbacher, U. (2007). z-tree - zurich toolbox for readymade economic experiments. *Experimental Economics 10*(2), 171–178.

Giles, L. (1910). Sun tzu on the art of war: The oldest military treatise in the world.

Harrington, D. and B. Robertie (2004). *Harrington on Hold'em*. Henderson, NV: Two Plus Two.

Harsanyi, J. (1973). Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points. *International Journal of Game Theory 2*(1), 1–23.

Kaufman, H. and G. Becker (1961). The empirical determination of game-theoretical strategies. *Journal of Experimental Psychology: General 61*, 462–468.

Levitt, S. D., J. A. List, and D. H. Reiley (2007, August). What happens in the field stays in the field: Professionals do not play minimax in laboratory experiments.

Luce, R. D. and H. Raiffa (1957). *Games and decisions. Introduction and critical survey*. New York: John Wiley & Sons, Inc.

McKelvey, R. and T. Palfrey (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior 10*, 6–38.

Nagel, R. (1995). Unraveling in guessing games: An experimental study. *The American Economic Review 85*(5), 1313–1326.

Nash, J. (1951). Non-cooperative games. *The Annals of Mathematics 54*(2), 286–295.

Nash, J. (1996). *Essays on game theory*. Cheltenham, UK: Edward Elgar.

Ochs, J. (1995). Games with unique mixed strategy equilibria: An experimental study. *Games and Economic Behavior 10*(1), 202–217.

Oechssler, J. (1997). An evolutionary interpretation of mixed-strategy equilibria. *Games and Economic Behavior 21*(1-2), 203–237.

O'Neill, B. (1987). Nonmetric test of the minimax theory of two-person zerosum games. *Proceedings of the National Academy of Sciences 84*(7), 2106–2109.

Osborne, M. and A. Rubinstein (1994). *A Course in Game Theory*. Cambridge, Massachusetts: MIT Press.

Rapoport, A. and D. Budescu (1992). Generation of random series in two-person strictly competitive games. *Journal of Experimental Psychology: General 121*(3), 352–363.

Reny, P. J. and A. J. Robson (2004). Reinterpreting mixed strategy equilibria: a unification of the classical and bayesian views. *Games and Economic Behavior 48*(2), 355–384.

Rubinstein, A. (1991). Comments on the interpretation of game theory. *Econometrica 59*(4), 909–924.

Schelling, T. C. (1980). *The Strategy of Conflict*. Cambridge: Harvard University Press.

Shachat, J. (2002). Mixed strategy play and the minimax hypothesis. *Journal of Economic Theory 104*(1), 189–226.

Stahl, D. O. and P. W. Wilson (1995). On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior 10*(1), 218–254.

Walker, M. and J. Wooders (2001). Minimax play at wimbledon. *The American Economic Review 91*(5), 1521–1538.