Learning-by-Doing in an Ambiguous Environment

Jim Engle-Warnick*
McGill University

Sonia Laszlo
McGill University

December 2008

Abstract: We experimentally test whether risk aversion or ambiguity aversion can explain decisions in a learning-by-doing game. We first measure subjects' preferences toward risk and ambiguity, and then use these measures to predict behavior in the game. We find that ambiguity averse subjects pay more often to resolve ambiguity, and we find that less risk averse subjects earn more in the game. Our results, in light of a previous field study of farmers in a developing economy, provide further evidence of a link between ambiguity aversion and technology choice, as well as a link between risk aversion and farm profitability.

Keywords: Learning-by-doing; Technology choice; Risk preferences; Risk measurement instruments; Ambiguity Aversion; Experimental economics.

JEL classification nos. C91, D80.

Contact Information: Engle-Warnick: Department of Economics, 855 Sherbrooke St. W., McGill University, Montreal, QC, H3A 2T7, Canada, e-mail: jim.engle-warnick@mcgill.ca, tel.: (514) 398-1559; fax: (514) 398-4938; Laszlo: Department of Economics, 855 Sherbrooke St. W., McGill University, Montreal, QC, H3A 2T7, Canada, e-mail: sonia.laszlo@mcgill.ca, tel.: (514) 398-1924; fax: (514) 398-4938.

*Corresponding author.

Acknowledgements: We thank Claude Montmarquette, Peter Moffatt, Daniel Parent, and participants at the 2006 Economics Science Association European Meetings in Nottingham. We thank Duncan Whitehead for research assistance. We gratefully acknowledge the Social Science and Humanities Research Council of Canada for funding, and the Centre for Interuniversity Research and Analysis on Organizations (CIRANO) and the Bell University Laboratory in Electronic Commerce and Experimental Economy.

1 Introduction

The adoption of new technologies is a fundamentally important issue for economies. From farmers' decisions to switch to new seeds, to the choice of computer technology, to contraceptive choice, to the diffusion of technology throughout an economic sector, new technology adoption is a widely studied and important phenomenon.

For example, it has long been accepted that farmers in developing countries are slow to adopt new technologies. Among many competing hypotheses, risk aversion is viewed as an important determinant of technology adoption (Feder, Just, and Zilberman (1985) and subsequent literature). Because they are poor and thus have little recourse to credit or insurance markets, subsistence farmers tend to be relatively risk averse. Important evidence for this idea came from Binswanger (1980) who, in his well-known experiment with choices between lotteries, established the existence of a group of farmers in India whose measured preferences were indeed risk averse.

Despite the theory and evidence, it is difficult to make the connection between risk preference and technology choice. This is because risk preferences are neither observable nor typically formally revealed in survey data. Many field studies have compensated by using a combination of survey questions on risk attitudes and modeling assumptions to correlate measures of risk with technology choice (e.g., Knight, Weir, and Woldehanna (2003), Antle and Crissman (1990)). Results from these studies normally find a negative correlation between risk aversion and new technology adoption.

A recent field study took a further step by combining a Binswanger-like laboratory experiment in the field with a socioeconomic survey to correlate an incentivized behavioral measure of risk with technology decisions reported in the survey. With this methodology, Engle-Warnick, Escobal and Laszlo (2008) provided evidence from subsistence farmers that it is not actually risk aversion but aversion toward ambiguity that predicts seed technology decisions on the farm.[1] Furthermore, among competing hypotheses typically explored in the context of farming in developing countries, learning-by-doing was found to be an important mechanism through which farmers made new technology decisions.[2] The inference is that farmers learn about other technologies as they gain experience with one of them, and that they view unknown technologies as having unknown probability distributions over possible outcomes. Farmers apparently learn by doing in an ambiguous environment.

This result is surprising given the reliance on the hypothesis that risk aversion drives such results. This result is new because it relies on both the field laboratory experiment and the survey results from the same subject to correlate the measure with the decision. However, evidence from this study is indirect in the sense that we do not really know that the farmers see the technology choice as ambiguous in the way that we define it as ambiguous, nor do we have direct measurable data regarding their technology choice.

By contrast, in the traditional experimental laboratory we can present subjects with a truly ambiguous learning-by-doing environment, and in this paper this is precisely what we do. The idea is to explore further the connection between ambiguity preferences and new technology choice. The advantage of this study compared with others is that, while it is informed by field results, all the data are collected under the control of the experimental laboratory.

In our experiments subjects first report for a session during which they respond to individual choice problems under uncertainty to measure their risk and ambiguity preferences. The instruments are standard: to measure risk preferences, subjects are asked to choose between risky and safe lotteries, and to measure ambiguity preferences, subjects are given the option to buy their way out of playing ambiguous lotteries, i.e., lotteries with unknown probability distributions over outcomes.

[1] Using incentivized laboratory measures of risk and time preferences are becoming increasingly used both to predict decisions in laboratory games in the field (e.g., Barr (2003) for behavior in games, and Eckel, Johnson, and Montmarquette (2007) to predict decisions in the field).

[2] See Foster and Rosenzweig (1995) for other examples of such a finding.

The subjects are then recalled a month later to play a technology choice game. In the game, modeled after Jovanovic and Nyarko (1996), as the subjects use a technology, they learn in a noisy way about a more efficient technology. The sole decision is when to switch from the first technology into the more efficient one.[3] When to switch is ambiguous: the subjects do not know the probability distribution of earnings for the possible switch times, thus they must gain experience with the game to resolve the ambiguity. Resolving ambiguity comes at a nominal cost.

We report two main findings. First, our measure of ambiguity aversion is correlated to the degree to which subjects are willing to incur a cost to reduce ambiguity in the learning-by-doing game. This is important, because it suggests that the ambiguity aversion measure is measuring what is intended. Second, risk aversion is negatively correlated with performance, measured by net earnings, in the learning-by-doing game. This is surprising, and suggests both further exploration into the mechanism that causes this result, and further exploration into whether this is also true in the field.

Our paper makes the following contributions. We present the first laboratory ambiguity instrument with a scale formally derived from an ambiguity model (Klibanoff, Marinacci, and Mukerji (2005)). We validate the use of the instrument in learning-by-doing contexts by showing that behavior is robust to the change in framing and context from the instrument to the learning-by-doing game. And we present two new conjectures to test in the field: (1) that risk averse subjects do worse under learning-by-doing, and (2) that this result is driven by the extremely risk averse.

Our paper sharpens knowledge regarding findings in the field, and suggests future field investigation.[4] The main result implies that reducing ambiguity may be an avenue for

[3] Experimental tests of learning-by-doing have typically focused on fitting dynamic learning models to choice data: Camerer (2003) provides a survey. Merlo and Schotter (2003) provide an experimental test, using a model similar in spirit to Jovanovic and Nyarko (1996), of whether observational learning of an agent who is learning-by-doing can be an efficient form of learning.

[4] Other studies provide links between the laboratory and the field. For example, Kagel and Roth (2000) report a study in which a laboratory experiment is conducted to isolate the effect of an institution on

improving real technology choice decisions. The relatively poor performance by the extremely risk averse could have implications for the complexity of the instrument needed to measure risk aversion: identifying extreme types in the field is likely to be less difficult than classifying all types.

We also provide a methodological contribution with a demonstration of a method to handle inference with a highly skewed distribution of payoffs. Because subject payoffs are determined by a quadratic loss function in the learning-by-doing model, the distribution of payoffs is not normal. We illustrate a method to transform the variables to correct for this, and to estimate marginal effects with the transformation.

The next section describes the experiment for measuring risk and ambiguity preferences. The following section describes the learning-by-doing experiment, which was conducted approximately one month later. We then present the experimental results and conclude.

2 Measuring Preferences

The first part of our study consisted of a laboratory experiment conducted to measure risk and ambiguity preferences. The same subjects were recalled one month later to play the learning-by-doing game. The goal is to experimentally test the validity of the ambiguity preference instrument in predicting behavior in learning-by-doing. This section details the risk and ambiguity preference instruments, the experimental procedures, and the models that generate the explanatory variables.

2.1 Risk Instrument

Our risk preference measure is based on the well-known instrument used by Binswanger (1980), and Eckel and Grossman (2003). This instrument, shown in Figure 1, is simply a

decisions in a labor market. The study provided support for implementing a new institution in the field.

choice of a most preferred lottery from a collection of five lotteries. Beginning with the top lottery and moving clockwise around the remaining lotteries, both the expected value and the variance of the lotteries increase. This trade-off permits inference regarding the subject's attitude toward risk from her choice.

We decomposed this instrument into a series of binary choices, shown in Figure 2, where each row is a separate decision making problem, and where the two lotteries involved in each choice are contiguous in Figure 1. This decomposition, similar to Holt and Laury (2002) and used in Engle-Warnick, Escobal, and Laszlo (2009), will allow the construction of an ambiguity instrument that closely resembles the risk instrument, without altering the theoretical basis of the instrument in Figure 1. Specifically, as a subject moves down the rows of Figure 2, once she chooses the riskier of the two lotteries, she should choose the safer of the two in all subsequent rows. For example, if a subject chooses the riskier lottery in row 1, then this lottery becomes the safer lottery in row 2. Since she has already revealed her maximum acceptable risk level in row 2, all subsequent row choices should be safe. This implies that the risk preference is revealed by the number of times the safe lottery is chosen.

2.2 Ambiguity Instrument

Our ambiguity preference measurement instrument, shown in Figure 3, presents the subjects with five binary choices between a lottery with unknown probabilities and a lottery with the same outcomes but with known 50/50 probabilities. Each binary choice corresponds to each one of the gambles in Figure 1 along with its ambiguous counterpart. Choosing the lottery with the known probability distribution over outcomes comes with a small cost, shown just below the lottery itself. Thus the decision problem for the subject is whether or not to pay a small cost to eliminate ambiguity, where ambiguity is uncertainty regarding the probability distribution over outcomes. This problem is similar to the one posed in the well-known Ellsberg Paradox.

2.3 Experimental Procedures

The sessions were conducted with paper and pencil. Subjects were given a book with one decision to make on each of forty-four pages.[5] The pages were randomly ordered, as was the left to right presentation of the gambles, and the instructions were given orally. Subjects indicated their decisions by placing a mark above their choice in their booklet, and an experimenter verified that there was exactly one choice made on each page when completed. To prevent influencing the results, the subjects were not informed in advance that their booklets would be verified. Subjects were privately paid for one randomly chosen decision. All payoffs were displayed in Canadian dollars.

Inference from choices in the ambiguity instrument requires controlling the subjects' priors over the probability distribution of outcomes. Otherwise, the prior itself is a parameter of the model and must be estimated from the data. To the extent possible, we implemented the instrument in a manner consistent with uniform priors over this subjective distribution. There were ten chips in a bag, all chips were either blue or yellow, and the subjects were not told how many chips were blue or yellow. To control for beliefs regarding colors, the subjects were asked to choose which color represented the better of the two lottery outcomes. For the risk instrument, there were ten chips in a bag, five of them blue, and five of them yellow.

We conducted six sessions, which were run at the experimental laboratory at the Centre for Interuniversity Research and Analysis on Organizations in Montreal. The subjects were recruited by e-mail from the English-speaking subject pool (the laboratory also has a French-speaking subject pool), using the Online Recruitment System for Economic Experiments (Greiner (2004)). Subjects were paid a $10 show up fee upon arrival before making their decisions, and the same experimenter conducted the sessions and read the script to the

[5] The experimental design consists of an additional set of questions that study the effect of additional choices. In addition to the risk and ambiguity measures, there were decisions to reveal the effect of additional alternatives on choice, and to reveal preferences for payoff dominated alternatives. The effect for this experimental study was to randomly scatter the nine questions we are interested in here among thirty-five other questions.

subjects in all the sessions. One-hundred and six subjects participated in this experiment, with session sizes of fifteen to twenty. Subjects earned an average of $20 in addition to the $10 show up fee. The experiments lasted approximately one hour.

2.4 Risk Preference Model

We infer risk preferences from choices using standard expected utility theory. Risk is characterized by a probability distribution over payoffs. Risk preferences are characterized by a standard utility function over outcomes.

All lotteries in our instruments are composed of a high and low outcome, x_l and x_h, and all outcomes occur with equal probability. There is always a left and right lottery to choose from, which we will indicate with superscripts L and R. A subject chooses the left lottery if

$$\frac{1}{2}u(x_l^L) + \frac{1}{2}u(x_h^L) > \frac{1}{2}u(x_l^R) + \frac{1}{2}u(x_h^R).$$

We use this equation and the CARA utility function $u(x) = -\frac{1}{r}\exp -rx$ to compute an interval estimate of the parameter r for each of the five possible number of safe lottery choices.[6] We compute r numerically, setting the two sides of the equation equal to each other. The interpretation of r is as follows: $r > 0$ represents risk aversion, $r < 0$ represents risk preferring, and linearity, i.e., $u(x) = x$, represents risk neutrality.

The first column of Table 1, labeled "Five-Circles", shows the midpoint of the interval estimate of the parameter r from the CARA utility function for each of the 5-Circle choices (Figure 1), where risk aversion is increasing in r. We compute the interval by numerically computing the value of r that sets the expected utility of two neighboring gambles equal. The second column of Table 1, labeled "# Safe Choices", maps our decomposed instrument

[6] The CARA utility function is restrictive because the degree of risk aversion is not constant over stakes. It is acceptable for our application because the stakes for both the instrument and the learning-by-doing game were similar, and the stakes were relevant because the subjects earned enough money to cover their opportunity costs. It is convenient to use CARA for both risk and ambiguity preferences in numerical solutions later.

into the Five-Circle instrument.

Notice that the middle three rows of the table produce a midpoint for the risk parameter r. That is, if a subject makes one, two, or three safe choices, then we can identify her risk preference up to an interval. However, at the two ends of the scale, if the subject makes zero or four safe choices, then we can only identify a maximum or a minimum value of the parameter r. This is common with risk preference instruments, and will require attention when constructing the ambiguity table as well.

2.5 Ambiguity Preference Model

We infer ambiguity preferences from choices using the "Smooth Model of Decision Making Under Ambiguity" in Klibanoff, Marinacci, and Mukerji (2005). In this model, ambiguity is characterized by uncertainty about the probabilities of the lottery outcomes. Ambiguity preferences are characterized by two elements: (1) a prior over the probability distribution of outcomes, and (2) a subjective utility function V that operates on the lotteries.

Assuming a uniform prior over the distributions of outcomes, and noting that there could have been from zero to ten chips representing the higher of the two outcomes, the subject chooses the ambiguous lottery if

$$\frac{1}{11}\sum_{i=0}^{10} V(\frac{i}{10}u(x_l) + \frac{10-i}{10}u(x_h)) > V(\frac{1}{2}u(x_l - 0.50) + \frac{1}{2}u(x_h - 0.50)).$$

This inequality illustrates two important facts regarding revealed ambiguity preference. First, it is necessary to know the subject's risk preference before one can draw conclusions regarding her ambiguity preferences. Second, the assumption of the uniform prior over the probabilities allows us to identify a parameter of her subjective utility function over *gambles*, V, which characterizes the attitude toward ambiguity. Without this assumption, the form of the prior and the parameter of the utility function for ambiguity would have to be jointly estimated.

For convenience, we use the functional form $V(x) = -\frac{1}{\alpha}\exp -\alpha x$ for our analysis, which is the same CARA function we use for modeling risk preferences, and α is the ambiguity preference parameter. The interpretation of the subjective utility function is similar to that of the expected utility function $u(x)$: concavity represents aversion to ambiguity, linearity represents ambiguity neutrality, and convexity represents a preference for ambiguity. Notice that, as is usually explained in the Ellsberg Paradox, an ambiguity neutral subject would never pay to avoid ambiguity: any payment for avoidance results in a strictly lower utility level when V is linear.

Since inference regarding ambiguity preferences is conditional on risk preferences, we construct a two-dimensional table to present the predictions of our ambiguity preference instrument. We use the midpoints (and endpoints) of the interval estimate of the risk parameter r in Table 1 to compute the interval estimate of the ambiguity parameter α. We then find the midpoints and endpoints of α for each combination of risk and ambiguity preference. The results are displayed in Table 2.

In Table 2, the rows represent the decision made in the risk preference instrument, and the columns represent the decision made in the ambiguity preference instrument. For example, if a subject makes 2 safe choices in the risk instrument, and pays 3 times to avoid ambiguity in the ambiguity instrument, then the midpoint α is 0.103. Notice the extreme left and right columns: if a subject never pays to avoid ambiguity, then we call her ambiguity neutral, and if the subject always pays to avoid ambiguity (even in the case of $13 for sure), then we call her infinitely averse to ambiguity.

2.6 Constructing the Risk and Ambiguity Indexes: The Explanatory Variables

Tables 1 and 2 present the interval estimate of the parameter representing the preference, risk and ambiguity, where the interval may or may not be bounded on one or both sides.

The top and bottom rows in Table 1 contain unbounded intervals, as does the border of Table 2. For empirical work we need the rank ordering of risk and ambiguity preferences, which is presented in Tables 3 and 4.

Table 3 is straightforward: the index for risk preference is simply the number of safe choices made in our risk instrument.[7] Thus the instrument is increasing in risk aversion. Table 4 is a bit more complicated because of the border cells that represent unbounded intervals. The left and right column borders represent never or always averse to ambiguity, while the top and bottom rows are analogous to the top and bottom rows of Table 3, with no bound to the interval either above or below.

To construct an index for ambiguity preference we first label the left column (non-aversion to ambiguity) 0, and then index the interior cells of the table in order of increasing aversion to ambiguity. For the first row of the table (four safe choices in the risk instrument), we find the largest entry a non-border cell of the table that is smaller than the entry in this row. For example, in the cell representing four safe choices and one choice to avoid ambiguity, the endpoint for α is 0.0774. The largest entry that is smaller than 0.0774 in the interior of the table is at one safe choice and three choices to avoid ambiguity, 0.0625, and its rank is 6. We thus label the rank of the cell in the one, four position of the table > 6. We perform the analogous search in the bottom row of the table, i.e., finding the smallest entry that is greater then the entry in the bottom row of the table.

The last task remaining is to deal with the top and bottom rows of Table 4: we take the midpoint of the interval that contains the rank of the cell. For example, in the cell representing four safe choices and one choice to avoid ambiguity, the rank is at least 7 (> 6), and our highest rank in the table is 13; thus we label this cell as the midpoint between 7 and 13, which is 10. We do the same for the remaining undetermined cells, and present the final rank ordering for the ambiguity choices in Table 5.

[7] This is also the measure used by Holt and Laury (2002).

3 Learning-by-Doing Experimental Design

Approximately one month after completing the preference measure experiment, subjects were recalled to play the learning-by-doing game. The subjects were not informed that the second experiment was related to the first experiment.

3.1 Learning-by-Doing Model and Game

We use the learning-by-doing model of Jovanovic and Nyarko (1996) as the basis of our game. In this model, a firm learns about a parameter of a technology by using it. At the same time the firm learns in a noisy way about a parameter of a more efficient technology. Think of a farmer planting seeds at the beginning of a growing season, with more modern and efficient varieties available. Learning how to plant the traditional seed assists with learning some aspects about how to plant the modern seed. But the learning is noisy, because choices such as type of irrigation and type of fertilization might be different with the modern seed. This same model has been used in the economic development literature concerned with technology adoption (e.g., Foster and Rosenzweig (1995) and Rosenzweig (1995)).

The game is played repeatedly, where the firm chooses to continue with the least efficient technology (technology 1), or to permanently switch to the more efficient one (technology 2). Whichever technology the firm chooses, it must also choose an intensity of use. Switching from the first technology to the more efficient technology results in an immediate loss in profits, because learning about the more efficient technology is noisy, and the firm's prior for the optimal intensity of use is thus inaccurate. However, switching also results in the opportunity to earn higher profits in the long-term, because learning will be faster, and because of the efficiency gain.

Formally, the payoff, q, to the firm is determined by a quadratic loss function, which measures the time t difference between the firm's selected intensity of technology use, x, and an optimal intensity of use, y_t, which is randomly determined:

$$q = \gamma^n[a - (y_t - x)^2], \gamma > 1 \tag{1}$$

The parameter γ determines the increase in efficiency from a new technology, where the available technologies are indexed by the integer n. At time t, the firm selects x, then sees q, at which time it can update its beliefs with Bayes' rule about the technology parameter θ_n by inferring y_t.

The optimal choice for technology intensity is y_t, and this optimal level is determined by the technology specific parameter θ_n and a random variable:

$$y_t = \theta_n + w_t \tag{2}$$

where w_t is normally distributed i.i.d. with zero mean. The technologies are linked through θ:

$$\theta_{n+1} = \sqrt{\alpha}\theta_n + \epsilon_{n+1} \tag{3}$$

The optimal behavior of the firm involves using Bayes' rule to update its belief about $x = E[y_t] = E_t[\theta_t]$ each time it observes its payoff q. At some point, the immediate cost of switching no longer exceeds the future cumulative gains from efficiency, and the firm should switch. If the firm switches too soon, it loses profits from not having learned enough. If the firm switches too late, it loses profits from efficiency gains.

3.2 Experimental Procedures

The experiment was programmed and conducted with the software z-Tree (Fischbacher 1999). The subjects played the learning-by-doing game for twenty-five rounds. Their sole decision was which period to switch from the less efficient technology (technology 1) to the more efficient technology (technology 2). In our implementation of this game, we gave the computer a prior over the optimal use of technology 1 (x), and allowed the computer to update its prior using Bayes' Rule for both technologies after the realization of the optimal use (y_t) each round. The computer played its estimate of x, the period payoff was realized, then the computer updated its new estimate of x for both technologies. Our design thus limited the subjects' strategy to finding the optimal switch-point from technology 1 to technology 2.

The subjects' computer display included the round number, the technology currently in use, the computer's estimate of x for the technology currently in use, the period realization of the optimal use, the period payoff, the total payoff for all periods played, and the computer's estimate of x for technology 2 (this last information reminded the subjects that the computer was learning about the unused technology as long as technology 1 was in use). Once the subjects switched to technology 2, they were not permitted to switch back.

One challenge in implementing this model is to find parameters that result in a steep enough surface of maximization to be behaviorally meaningful. Quadratic loss functions, which are flat at the maximum, can be poor with regard to providing economic incentives for human subjects to optimize. We chose the following parameters for the model:

$$a = 50; \gamma = 1.8; \alpha = 20; \epsilon \propto N(0, 0.25); w \propto N(0, 0.25)$$

We played our game, switching thirty times after each period of the twenty-five period game, and computed an average payoff for switching in each period. This computation, which reports actual values from our computer program that implemented the experiment, is shown in Figure 4. Figure 4 confirms that our chosen model parameters result in a fairly steep surface of maximization with a switch period that should not be easily guessed by the subjects. The theoretical optimal switch-period is $t = 8$, and the maximum expected payoff is approximately \$20. The worst thing to do is to switch right away; this is because at this point not enough has been learned about the optimal intensity of use of technology 2.

In the instructions the subjects were informed that the task was to choose whether or not (and when) to switch to technology 2 in a twenty-five period game. The subjects were shown the loss function that determined their payoffs so that in theory they were aware that the payoff function was smooth and contained a unique maximum. The subjects were told that the computer updated its information and learned about both technologies. The subjects were not given equations (2) or (3), so that they knew neither the process generating the optimal intensity of use, nor the way the technologies were linked. They were told that if they never switched technologies, they could expect to earn approximately $9.00.

Notice that in this game, there is a distribution for the payoff for each possible switch point. To the subjects, this distribution is unknown because they did not have full information about the model. Thus, this information condition is the basis for making the technology choice environment ambiguous. Subjects who pay to avoid ambiguity in the preference measurement experiment should also pay to resolve this payoff ambiguity in the learning-by-doing experiment.

After setting out the decision making problem, the instructions then informed the subjects that they could pay $0.50 to practice the game for no pay as many times as they wanted. This gave the subjects the opportunity to resolve the ambiguity regarding when to switch from technology 1 to technology 2, at a low cost. The question is whether we can use our preference measures from the first experiment to predict behavior and performance in the second experiment.

We conducted seven sessions, which were run at the same experimental laboratory as were the risk and ambiguity preference measuring experiments. The subjects were recruited by e-mail from the list of subjects who participated in the preference elicitation experiment; all previous participants were invited to participate in the new experiment. Subjects were paid a $10 show up fee. Seventy-two subjects participated in the experiments. Subjects earned an average of $15.40 in addition to the $10 show up fee. The experiments lasted approximately one hour.

4 Linking the Instruments with the Learning-by-Doing Game

In the learning-by-doing game the problem is to decide whether to pay for a chance to *reduce* ambiguity regarding the distribution of payoffs across the twenty-five strategies. In the ambiguity instrument, the problem is to decide whether to *eliminate* ambiguity regarding the distribution of lottery payoffs. The idea is to validate the use of the instrument for predicting behavior under learning-by-doing, shedding additional light on a tentative result found in the field.

In the learning-by-doing game, the choice to practice or not to practice for a fee represents a choice over strategy sets. Playing the game for pay requires selecting a strategy in the chosen set. We assume that the set of strategies available to play consists of the switch point that resulted in the best payoff during practice and the switch points not yet practiced. For every practice history, the subjects have a strategy with a current best known payoff. This payoff is relatively certain because there is a bit of randomness, and they can always throw away a strategy that results in a low payoff (recall that the subjects knew the expected payoff for never switching). Call this payoff x_c. While the actual possible payoffs are not known, it is logical to assume that paying to practice the game will uncover either a better or a worse payoff than is currently known.[8]

As with the ambiguity preference instrument, assume that subjects have a uniform prior over the distributions of possible payoffs for the set of switching periods not yet played. How

[8] While unknown payoffs are not part of the standard model of ambiguity, it was not possible to reveal the underlying payoff structure in the learning-by-doing game and still have an interesting decision problem that mimics the problem in the field. This feature of the experimental design allows the simplification of analysis of the game by assuming that practicing results simply in either a better or a worse outcome.

many such distributions are there? Whenever there are n strategies with unknown payoffs, there are $n + 1$ ways the unknown payoffs can be distributed with respect to known best payoff: n are better than x_c and none worse, $n - 1$ are better than x_c and 1 worse,..., none are better than x_c and n are worse. These distributions are analogous to the distributions in the ambiguity instrument: ten chips in the bag implied eleven possible distributions, from ten blue and no yellow, nine blue and one yellow,..., to no blue and ten yellow.

Thus paying to uncover the value of a strategy makes the distribution of unknown payoffs less ambiguous by reducing the number of ways the payoffs can be distributed. It is analogous to paying to discover the color of one chip in the bag of ten chips used in the ambiguity instrument. Imagine that you paid $0.50 to see a single chip in that bag: if the chip was blue, you could eliminate, as a possibility, the distribution in which all chips are yellow. With the ambiguity instrument, subjects revealed their willingness to pay a small amount to see all of the chips. In the learning-by-doing game, the distribution of payoffs was revealed one 'chip' at a time.[9]

Our experimental procedures in the learning-by-doing game, as in the ambiguity instrument, were designed to induce a uniform prior over payoff distributions: while we told the subjects the quadratic loss function, we did not explicitly discuss the shape of the surface of maximization, nor did any subject ever ask about it. This was important to induce a uniform prior over payoff distributions, as we did in the ambiguity preference instrument.[10]

This discussion linking the ambiguity instrument to the learning-by-doing game results in the following conjecture:

Conjecture 1 *Holding risk preference constant, the number of times a subject pays to practice is increasing in ambiguity aversion.*

[9] With the ambiguity instrument, all chips had an equal probability of being selected to determine the outcome, which corresponds to mixing over available strategies in learning-by-doing for pay (Olszewski (2006) makes a similar assumption in a theory of choice over sets). We investigate the empirical validity of this assumption in Section 5.5.

[10] The underlying payoff structure was not random, and could theoretically be exploited if the subjects had knowledge of this fact. We investigate this possibility in Section 5.5.

There is possibly a second way to interpret the game, simply viewing paying to practice as a risky lottery, and viewing the currently known best strategy as a certainty. In this case, the higher the degree of risk aversion, the better the subjective beliefs about the lottery must be to pay to practice, resulting in the second conjecture:

Conjecture 2 *Holding ambiguity preference constant, the number of times a subject pays to practice is decreasing in risk aversion.*

5 Experimental Results

5.1 Preference Measurement Experiment

In what follows, we analyze data from the 72 subjects who participated in both sets of experiments (i.e., the preference measurement and the learning-by-doing experiments).[11] Descriptive statistics of the observed socio-economic characteristics of this sample are provided in the appendix. Figure 5 shows a histogram of the number of safe decisions made by the subjects in the binary gamble. The figure reveals heterogeneity in decision-making, with subjects choosing all possible numbers of risky choices from zero to four. There is a mode at three safe choices, and the second-most chosen number of safe choices is two. Figure 6 presents a histogram of the number of times subjects paid to avoid the ambiguous gamble. There is a mode at zero, but roughly two thirds of the subjects paid to avoid the ambiguous gamble at least once. Since the ambiguity index described above is constructed from both measures, both showing a great deal of heterogeneity, our indices of risk and ambiguity aversion should have some predictive value. We now turn to the results from the learning-by-doing experiment.

[11] The subjects who did not participate in the learning-by-doing experiments are no different in their observed socio-economic conditions or in their responses to the preference experiments than those that did. We confirmed this by using t-tests for all independent variables, and in no case were we able to reject that the included sample of 72 observations is any different than the excluded sample of 34 observations.

5.2 Learning-by-Doing Experiment

Figure 7 presents a histogram of the number of times subjects paid to practice the learning-by-doing game. There is a mode at one, and the second-most number of times practicing is two. Five subjects did not practice the game at all, and nine subjects practiced three or four times. Figure 8 presents the distribution of payoffs, which is skewed to the right. This distribution is driven by the quadratic loss function (see equation (1)). To see this, recall Figure 4, which revealed a steep climb to the left of the optimal switch point of eight rounds, and a relatively flat area to its right. One could choose a switch point of six through fifteen rounds and expect to earn at least $15 in the experiment. Finding the optimal point adds approximately $5 in expectation to earnings, which is not trivial. However, a subject who experiments with moving the switch point down from later rounds (say, from round fifteen to round fourteen or thirteen), will find that reinforcement from experimentation may result in small increases to earnings, thus may stop experimenting. And many subjects may find themselves closer to the maximum with very similar earnings.

Thus without strong economic incentives to find the maximum, and with many switch points resulting in near-optimal earnings, we may expect to find the earnings of many subjects who play the game relatively well to be clustered in this range, and Figure 8 reveals that this is indeed the case. Furthermore, approximately two-thirds of our subjects earned payoffs in the range between $17 and $18, and these payoffs occur at the flattest part of the payoff function. Those subjects who do not do as well we find scattered to the left of this range. These relatively few subjects switch very early or very late, where the range of payoffs is larger. Our belief is that a non-normal distribution of payoffs may occur with a combination of this type of economic incentive and heterogeneous subjects. Our empirical analysis will take the non-normality of payoffs into account.[12]

[12] For another example of this type of result, Engle-Warnick and Turdaliev (2006) find a similar payoff distribution in a central banking game, which uses a quadratic loss function. They accounted for this by performing a regression analysis for each individual subject.

5.3 Predicting Practice Rounds

Section 4 provided a theoretical link between subjects' performance in the learning-by-doing experiment and their attitudes towards risk and ambiguity. Conjecture 1 stated that ambiguity averse subjects are more likely to pay to practice the learning-by-doing experiment, while Conjecture 2 stated that risk averse subjects are more likely to pay to practice. We show here that the decisions are correlated in the directions predicted by the conjectures.

As a first piece of evidence, the correlations between the risk and ambiguity aversion indexes and the number of times subjects paid to practice in the learning-by-doing are -0.1733 and 0.0878. While these raw correlations are small in magnitude, their signs correspond to the conjectures. Table 6 reveals additional evidence of the effects of risk and ambiguity aversion on the number of times subjects practiced in the learning-by-doing experiment. The table reports results from an ordered probit of the number of times practiced on our risk and ambiguity preference indices, including session controls. Risk aversion statistically significantly negatively predicts the number of times practiced, while the ambiguity aversion index significantly positively predicts it. Table 6 also reports the marginal effects from this exercise. More risk averse individuals are more likely to never practice or practice only once and less likely to practice more than once. Meanwhile, we find that subjects who are more ambiguity averse are less likely to never practice or practice only once, but more likely to practice twice.[13]

Given that the ambiguity index is constructed using the number of times subjects chose the safe gamble, a reader might be concerned about collinearity between the two measures leading to potential bias. We do not believe that this is a concern. First, the raw pairwise correlation between the two measures, while statistically different from zero, is only 0.4551.[14]

[13] These and all other regression results in this paper are reported with non-robust standard errors because of the small sample size. All of the results in this paper remain unchanged when considering Huber-White robust or bootstrapped standard errors.

[14] The correlations are even smaller when we use Spearman's or Kendall's rank correlation.

Second, informal tests examining tolerances, variance inflation factors and eigenvalues all rule out high collinearity problems. Third, the correlation of the estimated coefficients on the two measures, while also statistically different from zero, is -0.4816. Fourth, the predictive patterns are robust to randomly removing observations from the regression. Fifth, an F-test soundly rejects that the coefficients on both behavioral measures are jointly insignificant. Finally, the theory tells us that both measures belong in the regression, so not including either one would lead to an omitted variables bias.[15]

In summary, while the theory did not predict the magnitudes of the effects of risk and ambiguity aversion indices on the number of times practiced in the learning-by-doing experiment, the comparative statics did provide predictions about the signs of these effects. The results in Table 6 empirically corroborate our theoretical conjectures. We next use the preference measures to predict performance in the game.

5.4 Main Results

We wish to investigate the effects of risk aversion index (RM), ambiguity aversion (AM) and the number of times the subjects practiced (NP), on the payoffs (y) earned by each subject. To do so, we are interested in estimating the following regression:

$$y = \mathbf{X}'\beta + \epsilon \tag{4}$$

where $\mathbf{X} = [RM, AM, NP, \mathbf{Z}]$, $\mathbf{Z}$ a vector of control variables and ϵ a random disturbance term. In columns (1) and (2) of Table 7 we present the results from estimating (4) by ordinary least squares. The model estimated in column (1) does not contain socio-economic controls, while the model estimated in column (2) does.

We find that more risk averse individuals have higher payoffs, while subjects that practice

[15] A crude check of the robustness of this collinearity analysis is to replace the ambiguity index with simply the number of times subjects paid to avoid ambiguity. Doing so does not affect any of the patterns seen in Table 6.

more have lower payoffs. Specifically, subjects who practiced four times made lower earnings than those who never practiced (never practiced is the omitted category). The ambiguity aversion index does not affect payoffs in the game.[16] However, as seen in Figure 8, the distribution of payoffs is highly skewed to the right. Ordinary least squares may yield inconsistent estimates because such skewed distributions generate non-normal error terms. In fact, the Shapiro-Wilks test (reported in the table) resoundingly rejects that the error is normally distributed.

To rule out the possibility that the results found in Table 7 are driven by the skewness of the dependent variable and the rejection of normally distributed error terms, we transform the dependent variable using a 'zero-skewness logarithmic transformation' $\ln(\pm y+k)$, where $\text{sign}(y)$ and k are to be estimated.[17] The retransformation is shown in Figure 9, where we superimpose a normal distribution for comparison. The untransformed data clearly cannot approximate a normal distribution, while the transformed data look much more like a normally distributed variable. We thus estimate with ordinary least squares the following variant of (4):

$$\ln(\pm y + k) = \mathbf{X}'\beta + \epsilon \tag{5}$$

The results of this estimation are found in columns (3) and (4), where column (4) includes socio-economic control variables. We find that the estimated effects of the risk aversion index and the number of times practiced are statistically significant determinants of payoffs, consistent with the results from model (4). The Shapiro-Wilks test of normality of the residuals can no longer be rejected. However, the signs and magnitudes of these effects are quite different than those estimated by (4), because of the zero-skewness logarithmic transformation. To get the marginal effect of the independent variables of interest on payoffs,

[16] As for the results in Table 6, we checked for the possibility of collinearity among the behavioral and practice variables. Using the same techniques as above, we are able to rule out that collinearity is a problem.

[17] A simple log transform yields an equally skewed distribution and non-normal errors. We use the `lnskew0` command in `Stata 9.2` to transform the dependent variable and estimate $\text{sign}(y)$ and k.

we must retransform the model. We follow Duan (1983) and Abrevaya (2002) and apply Duan's smearing estimator:

$$\hat{y} = \frac{1}{n}\sum_{i=1}^{n}(\exp\{\mathbf{X}'\hat{\beta} + \hat{\epsilon}_i\} - k) \tag{6}$$

where i indexes over observations, and $\hat{\beta}$ and $\hat{\epsilon}$ are the estimated coefficients and error terms from (5). We calculate the marginal effect $m_j(x_0, \hat{\beta})$ by taking the derivative of (6) with respect to variable X_j, evaluated at a certain x_0:

$$m_j(x_0, \hat{\beta}) = \frac{\hat{\beta}_j}{n}\sum_{i=1}^{n}\exp\{\mathbf{X}'\hat{\beta} + \hat{\epsilon}_i\} \tag{7}$$

We evaluate $m_j(x_0, \hat{\beta})$ at the mean values of the X's. The standard errors are calculated by bootstrap and 500 replications. These marginal effects are presented in Table 8. The first two columns evaluate the model without socio-economic controls while the second two columns do include them. Notice that both models again tell the same story. Without the controls, at mean X's the marginal effect of the number of safe choices in the binary gamble is \$-0.860. With the controls, the marginal effect increases in magnitude to \$-1.346. There is a large and significant negative marginal effect for practicing the game four times (from a low of -\$4.891 to a high of -\$6.149). Thus the estimated marginal effects are both economically and significantly significant. The more safe choices a subject makes, i.e., the more risk averse the subject is, the lower her earnings.

5.5 Other Considerations

In this section, we provide evidence regarding assumptions that were necessary to apply the ambiguity model to the learning-by-doing game. First, we explore the validity of the assumption of the uniform prior over payoffs in the learning-by-doing game. Second, we look at evidence regarding mixing over remaining strategies in the strategy set in the same game.

We assumed a uniform prior over the possible distributions of better and worse payoffs in the learning-by-doing game in order to match the theory to the ambiguity instrument. In fact the underlying payoffs were not random, and subjects may have learned this, resulting in hill-climbing strategies. We checked to see whether we could use subjects' past uncovered information to predict their subsequent choices for evidence of this kind of behavior.

In this case, the trade-off is between paying to practice the game and finding a higher spot on the hill. Generally, if subjects understand the shape of the payoff function, the probability of stopping their search should increase as the improvement realized from practicing decreases. And the direction they move with their strategy should depend on the slope of the payoff generated by their previous two practice results.

First, for each practice round, we ran a regression where the dependent variable was whether or not subjects practiced the game in a practice round, and the independent variable was the difference between payoffs generated by practicing the previous two times. The regression was insignificant. Second, for each practice round, we ran a regression where the dependent variable was the direction in which the strategy moved, i.e., the sign of the difference between the switch period chosen in the previous two practices, and the independent variable was the difference in payoffs generated by practicing the previous two times. Again, the regression was insignificant. These regressions suggest that subjects did not know the shape of the payoff function, thus do not contradict the assumption of the uniform prior.

To shed some light as to what type of behavior we find in the data, we can investigate typical decisions made by our subjects, as depicted by Figures 10 and 11. Each figure shows the payoff function (same as in Figure 4), the practice points and the point played for pay. In Figure 10, subject #115 first practices by switching in the second period, uncovers about the same payoff as if he had never switched, which he confirms in the second practice period. In the third period, this subject switches right away, uncovers a very low payoff, stops practicing there, and plays for pay a very different strategy: he switches in the middle of the round, at the 12th period, coinciding with a relatively higher payoff.

In Figure 11, subject #111's first practice takes her to the actual maximum. In her second practice, she switches right away, uncovering a very low payoff. Given the huge difference in payoffs, a hill-climbing subject would perhaps practice again, switching sometime after period 8 (corresponding to her first practice strategy), as if she were following the gradient. However, she does not, and she chooses to play for pay in the third round, playing a strategy very close to her highest practice payoff strategy.

It was also necessary to assume mixing between available strategies in the learning-by-doing game for pay. This assumption is analogous to each chip being uniformly probable for selection in the ambiguity instrument. An obvious alternative to this is to play the strategy with the best known payoff. The patterns in Figures 10 and 11 suggest that subjects do not necessarily do this. The figures do lead us to ask whether subjects tend to play for pay by switching in a period close to the switch period with the higher practice payoff. Figure 12 presents the histogram of the absolute value difference in periods between the switch period played for pay and the switch period corresponding to the highest practice payoff. We observe how this distribution is skewed towards one period: subjects tend to play close to, but not exactly, their best practice strategy.

Thus while the subjects do not mix with uniform probability over all available strategies, they do not always play their best strategy either: they play strategies that they have not practiced as well. Our view is that this is consistent with the notion that ambiguity averse subjects pay to reduce ambiguity in the learning-by-doing game, then select among the remaining strategies when they are satisfied with the remaining degree of ambiguity.

This issue is probably the main point of difference in framing between the two experiments. Our experiment shows that ambiguity preferences nevertheless do transfer across the games. While a direct test of the theory could easily involve having nature choose the final strategy in the learning-by-doing game, the goal of this paper is to validate the instrument with a task that mimics decisions in the field, where nature does not make the technology choice. Rather, some technologies are known (risky), some are unknown (ambiguous), and either type may be chosen, as we found in our laboratory experiment. Most importantly, ambiguity averse behavior is robust to the two contexts.

6 Conclusions

In this paper we apply an instrument to measure ambiguity aversion, and we derive a scale of ambiguity aversion from a formal ambiguity model. The instrument is simple to implement, easy to understand, and parallel to existing risk preference instruments. We design a laboratory experiment to see whether subjects who measure ambiguity averse are willing to pay to reduce ambiguity in a learning-by-doing game. We find that they are. We also find that subjects who are extremely risk averse do worse in the ambiguous environment of the learning-by-doing game.

Our laboratory study is designed to sharpen knowledge regarding previous findings from the field, where we previously found that ambiguity aversion predicts technology choice. As part of these findings, we suspected, but could not observe directly, that the field decisions are influenced by learning-by-doing. In the laboratory, we validate the use of the instrument in the context of learning-by-doing. Thus the laboratory experiment helps to interpret the field result.

Our paper suggests that reducing ambiguity can improve technology choices. It predicts that risk averse subjects may do worse in this type of ambiguous environment. If, as occurred in our experiment, extremely risk averse decision makers drive this result, it should be easier to identify these types of people in the field for technical assistance. Our paper thus continues the conversation between field and laboratory observation. It answers a question that arose in the field, and it suggests further field investigation.

References

Abrevaya, J. (2002) "Computing Marginal Effects in the Box-Cox Model," *Econometric Reviews* 21:3:383-393.

Antle, J.M. and Crissman, C.C. (1990) "Risk, Efficiency, and the Adoption of Modern Crop Varieties" Evidence from the Philippines," *Economic Development and Cultural Change*, 38:3:517-537.

Barr, A. (2003), "Risk Pooling, Commitment and Information: An Experimental Test of Two Assumptions," Oxford Department of Economics Working Paper 2003-5.

Becker, G., M. DeGroot, and J. Marschak (1964), "Measuring Utility by a Single-Response Sequential Method," *Behavioral Science*, 8:4:293-315.

Binswanger, H. (1980), "Attitudes toward Risk: Experimental Measurement in Rural India", *American Journal of Agricultural Economics*, 62:3:395-407.

Camerer, C. (2003), *Behavioral Game Theory: Experiments in Strategic Interaction*, Princeton University Press, Princeton.

Duan, N. (1983), "Smearing Estimate: A Non-Parametric Retransformation Method," *Journal of the American Statistical Association* 78:383:605-610.

Eckel, C. and P. Grossman (2003), "Forecasting Risk Attitudes: An Experimental Study of Actual and Forecast Risk Attitudes of Women and Men," forthcoming in *Journal of Economic Behavior and Organization*.

Eckel, C., C. Johnson and C. Montmarquette (2007) "Saving Decisions of the Working Poor: Short- and Long-Term Horizons,"*Public Finance Review* 35:233-262.

Engle-Warnick, J., Escobal J., and S. Laszlo (2009), "How Do Additional Alternatives Affect Choice Under Uncertainty?", forthcoming *Canadian Journal of Economics* 42:1.

Engle-Warnick, J., Escobal J., and S. Laszlo (2008), "Risk Preference, Ambiguity Aversion, and Technology Choice: Experimental and Survey Evidence from Peru", McGill University Working Paper.

Engle-Warnick, J. and N. Turdaliev (2006), "An Experimental Test of Taylor-Type Rules with Inexperienced Central Bankers", McGill University Working Paper.

Feder, G. R. Just and D. Zilberman (1985), "Adoption of Agricultural Innovations in Developing Countries: A Survey," *Economic Development and Cultural Change*, January.

Fischbacher, U. (1999), "z-Tree - Zurich Toolbox for Readymade Economic Experiments - Experimenter's Manual", Working Paper Nr. 21, Institute for Empirical Research in Economics, University of Zurich.

Greiner, B. (2004), "An Online Recruitment System for Economic Experiments", in K. Kremer and V. Macho (Eds.): *Forschung und wissenschaftliches Rechnen 2003.* GWDG Bericht 63, Gttingen : Ges. fr Wiss. Datenverarbeitung, 79-93.

Foster, A. and M. Rosenzweig (1995), "Learning-by-Doing and Learning from Others: Human Capital and Technical Change in Agriculture," *Journal of Political Economy* 103:6:1176-1209.

Holt, C. and S. Laury (2002), "Risk Aversion and Incentive Effects in Lottery Choices," *American Economic Review*, 96:1644-1655.

Jovanovic, B. and Y. Nyarko (1996), "Learning-by-Doing and the Choice of Technology" *Econometrica* 64:6:1299-1310.

Kagel, J. and A. Roth (2000), "The Dynamics of Reorganization in Matching Markets: A Laboratory Experiment Motivated by a Natural Experiment", *Quarterly Journal of Economics* February, 2000: 201-235.

Klibanoff, P., M. Marinacci, and S. Mukerji (2005), "A Smooth Model of Decision Making under Ambiguity," *Econometrica*, 73(6):1849-1892.

Knight, J., S. Weir and T. Woldehanna (2003), "The Role of Education in Facilitating Risk-Taking and Innovation in Agriculture", *Journal of Development Studies*, 39:6:1-22.

Merlo, A. and A. Schotter (2003), "Learning by not doing: an experimental investigation of observational learning", *Games and Economic Behavior* 42:116-136.

Olszewski, W. (2007), "Preferences Over Sets of Lotteries", *Review of Economic Studies* 74:567-595.

Rosenzweig, M. (1995), "Why are there Returns to Education?" *American Economic Review* 85:2:153-158.

Table 1: Risk Preference Instrument CARA Parameter Midpoints and Endpoints

5-Circle Choice	# Safe Choices	r
13/13	4	>0.1096
10/17.5	3	0.07285
7/22	2	0.02885
4/26.5	1	0.01855
1/31	0	<0.0155

Table 2: Ambiguity Preference Instrument CARA Parameter Midpoints and Endpoints

# Safe	Number of Times Paid to Avoid Ambiguity					
Choices	0	1 (1/31)	2 (4/26.5)	3 (7/22)	4 (10/17.5)	5 (13/13)
4	0	>0.0774	>0.1246	>0.6974	>0.8651	inf
3	0	0.05925	0.09565	0.4171	0.68705	inf
2	0	0.0297	0.0491	0.10285	0.3904	inf
1	0	0.01635	0.0284	0.06255	0.24845	inf
0	0	<0.0144	<0.0253	<0.0562	<0.2251	inf

Table 3: Risk Preference Instrument Ranking

5-Circle Choice	# Safe Choices	Index
13/13	4	4
10/17.5	3	3
7/22	2	2
4/26.5	1	1
1/31	0	0

Table 4: Ambiguity Preference Instrument Initial Ranking

	Number of Times Paid to Avoid Ambiguity					
# Safe	0	1 (1/31)	2 (4/26.5)	3 (7/22)	4 (10/17.5)	5 (13/13)
4	0	>6	>8	>12	>12	inf
3	0	5	7	11	12	inf
2	0	3	4	8	10	inf
1	0	1	2	6	9	inf
0	0	<1	<2	<5	<9	inf

Table 5: Ambiguity Preference Instrument Ranking

# Safe	Number of Times Paid to Avoid Ambiguity 0	1 (1/31)	2 (4/26.5)	3 (7/22)	4 (10/17.5)	5 (13/13)
4	0	10	11	13	13	13
3	0	5	7	11	12	13
2	0	3	4	8	10	13
1	0	1	2	6	9	13
0	0	0	0.5	2	4	13

Table 6: Correlation between Ambiguity Aversion Measure and the Number of Times Practiced in the Learning-By-Doing Experiment

	Ordered Probit Coefficient	Marginal Effects if Practiced…				
		Never	Once	Twice	Three Times	Four Times
Number of Times Chose the Safe Gamble	-0.300	0.028	0.087	-0.063	-0.026	-0.025
	(0.118)**	(0.015)*	(0.038)**	(0.029)**	(0.015)*	(0.015)*
Ambiguity Aversion Index	0.064	-0.006	-0.019	0.014	0.006	0.005
	(0.030)**	(0.004)*	(0.010)**	(0.007)*	(0.003)	(0.003)
Pedicted probability practiced N times…		0.044	0.574	0.287	0.058	0.037
F-test of joint significance of both measures	7.45***					
Pseudo R-Squared	0.0921					
Wald Chi Squared (5)	16.24**					

Regressions include session controls. Standard errors in parentheses. * significant at 10%; ** at 5% and *** at 1%.

Table 7: Predictors of Earnings in the Learning-By-Doing Experiment

	OLS		Zero Skewness Logarithmic Transform	
Number of safe choices in binary gamble	-1.101	-1.570	0.322	0.542
	(0.409)***	(0.520)***	(0.123)**	(0.150)***
Ambiguity aversion index	-0.031	-0.084	0.023	0.016
	(0.104)	(0.124)	(0.031)	(0.036)
Practiced once	-2.317	-2.519	0.318	0.617
	(1.913)	(2.301)	(0.577)	(0.665)
Practiced twice	0.201	-0.396	-0.014	0.526
	(2.038)	(2.530)	(0.614)	(0.731)
Practiced three times	-0.425	-1.401	0.299	0.537
	(2.455)	(2.693)	(0.740)	(0.778)
Practiced four times	-7.604	-9.016	1.834	2.474
	(2.663)**	(3.111)***	(0.803)***	(0.899)***
Socio-Economic Controls	No	Yes	No	Yes
Skewness parameter (*k*)			-18.5997	
95% confidence interval for *k*			[-18.96013, -18.52517]	
R-Squared	0.3157	0.4165	0.2659	0.3935
Joint F-test of behavioural + practice variables	4.16***	4.28***	3.26***	3.91***
Shapiro-Wilks test for normality of residuals [p-value]	0.0002	0.0013	0.9757	0.9779
Observations	72	69	72	69

Standard errors in parentheses. * significant at 10%; ** at 5% and *** at 1%.

Table 8: Zero Skewness Logarithmic Transformation Marginal Effects for the Predictors of Earnings in the Learning-by-Doing Experiment

	Marginal Effects Evaluated at Mean Xs	
Number of safe choices in binary gamble	-0.860	-1.346
	(0.372)**	(0.464)***
Ambiguity aversion index	-0.062	-0.039
	(0.095)	(0.093)
Practiced once	-0.849	-1.534
	(1.590)	(1.841)
Practiced twice	0.039	-1.306
	(1.662)	(2.012)
Practiced three times	-0.796	-1.335
	(1.762)	(1.786)
Practiced four times	-4.891	-6.149
	(2.179)**	(2.808)**
Socio-Economic Controls	No	Yes

Standard errors in parentheses. * significant at 10%; ** at 5% and *** at 1%.

Figure 1: ‘Five Options’ Risk Preference Measurement Instrument

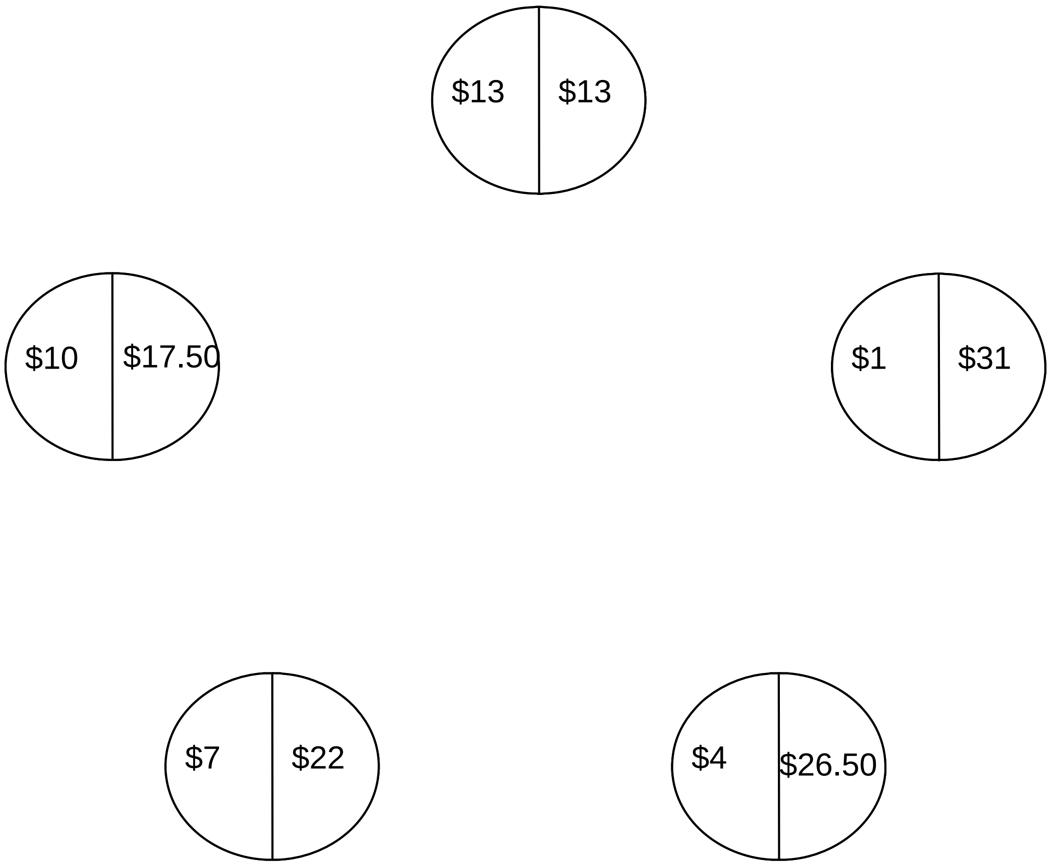

Figure 2: Decomposing the 'Five Options' Instrument into a Series of 'Binary Options' Instruments

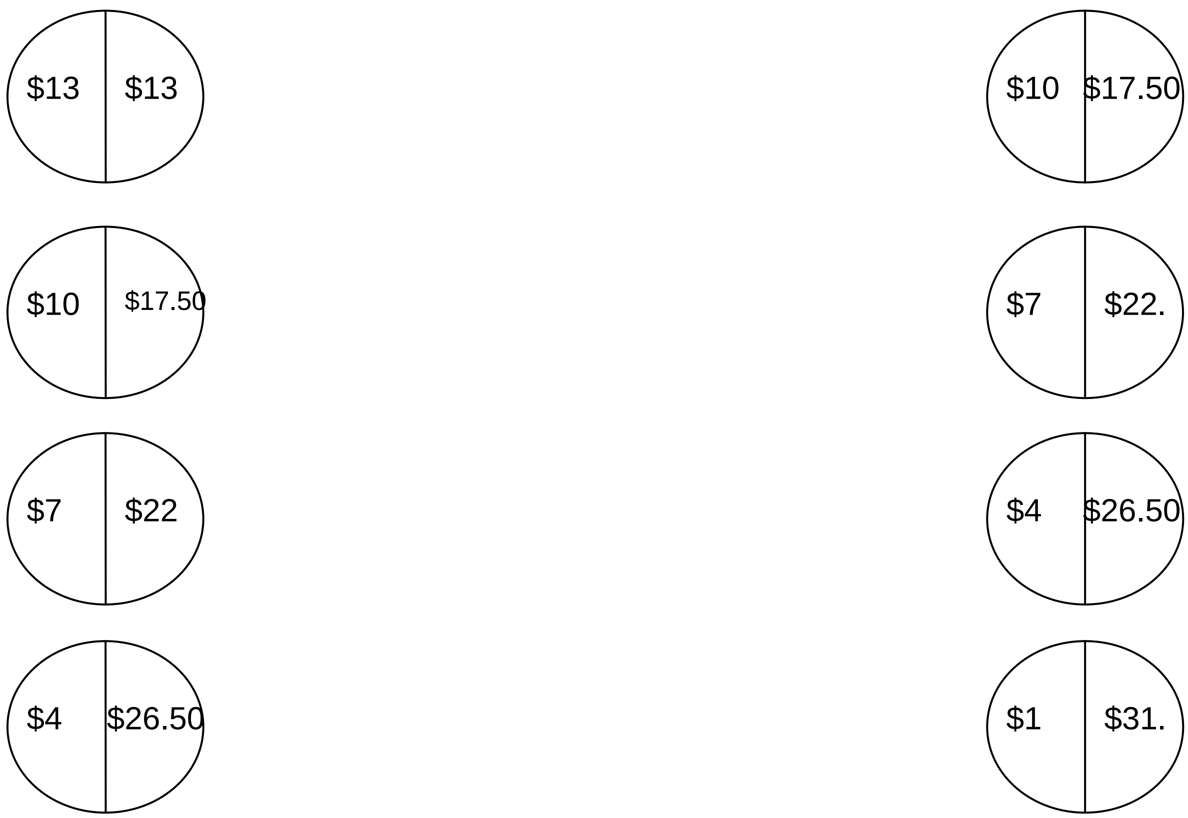

Figure 3: Binary Choices to Reveal Preferences for Ambiguity

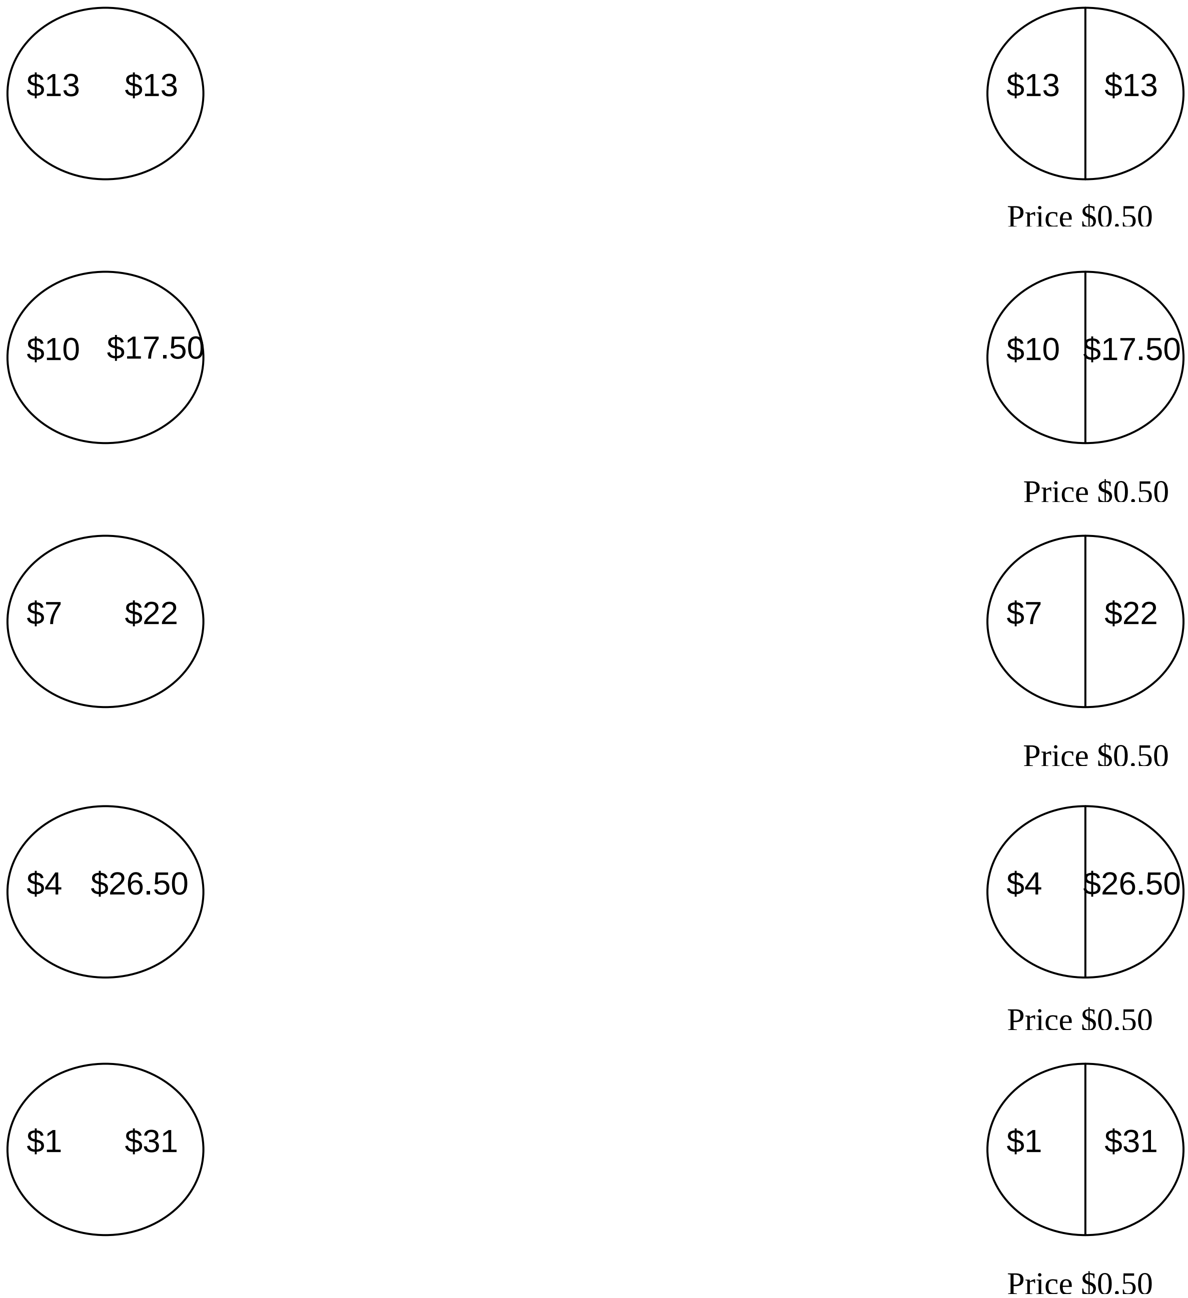

Figure 4: Average Payoff by Switchpoint in the Learning by Doing Game

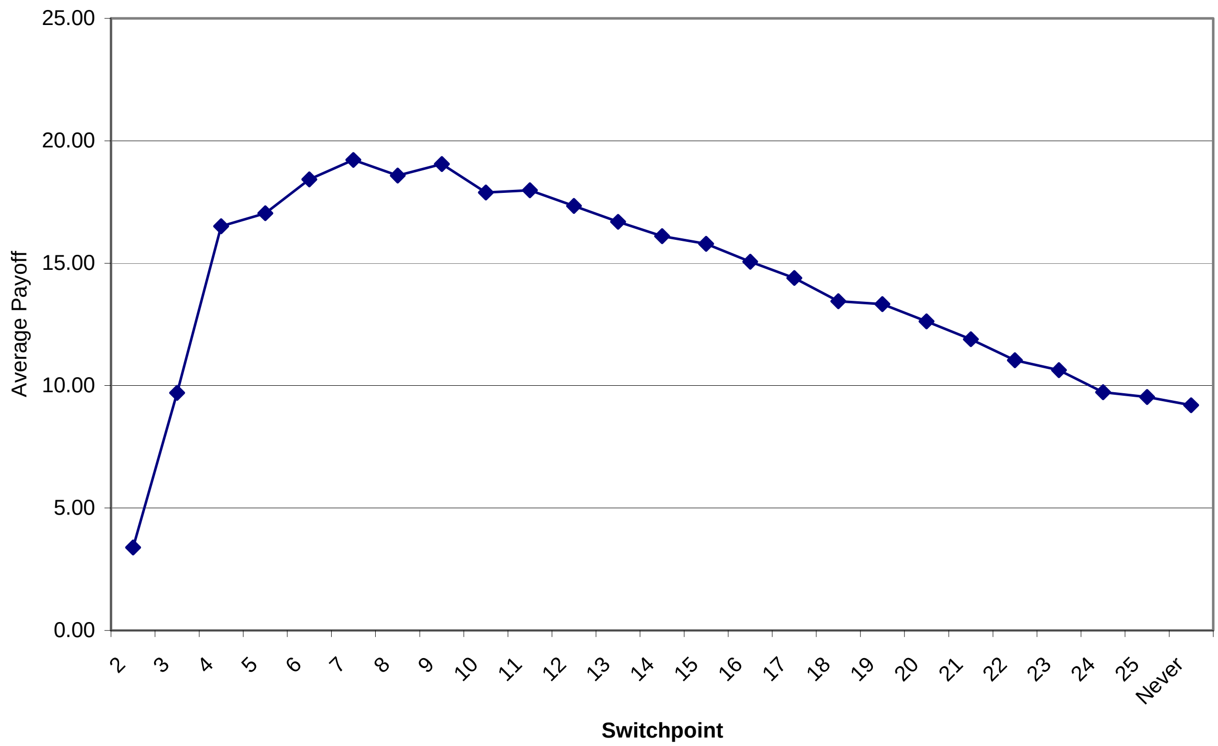

Figure 5: Distribution of Safe Choices in the Binary Game

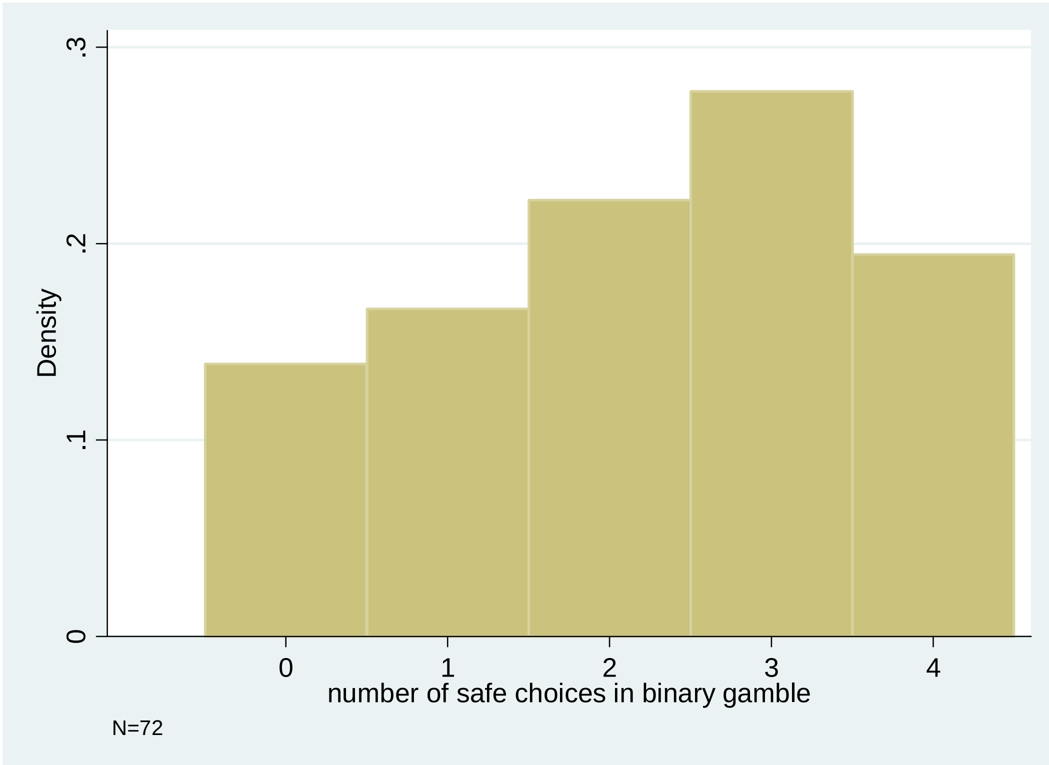

Figure 6: Distribution of the Number of Times Subjects Paid to Avoid Ambiguity

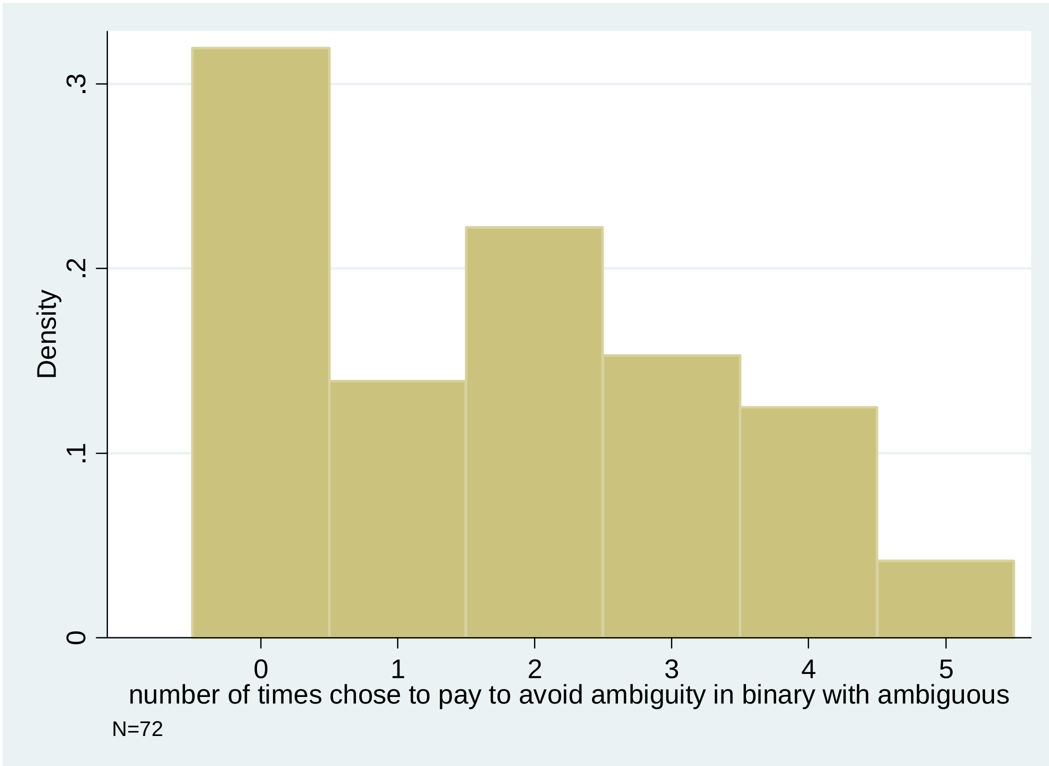

Figure 7: Distribution of the Number of Times Subjects Practiced the Learning-by-Doing Experiment

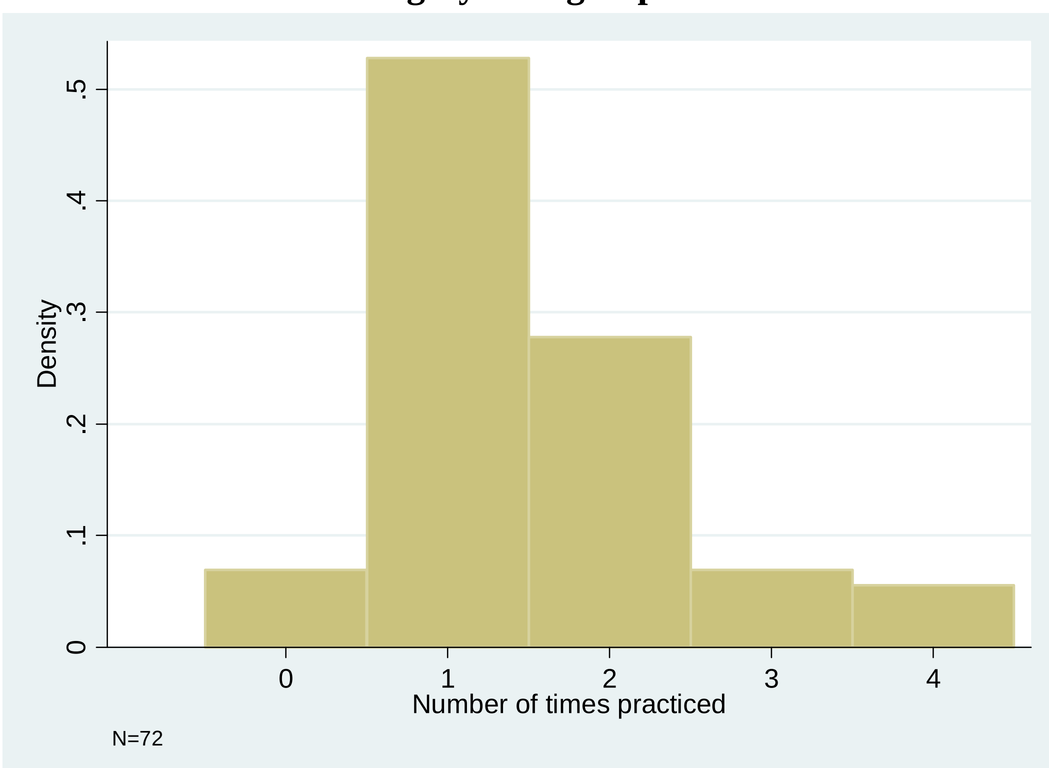

Figure 8: Distribution of Payoffs in the Learning-by-Doing Experiment

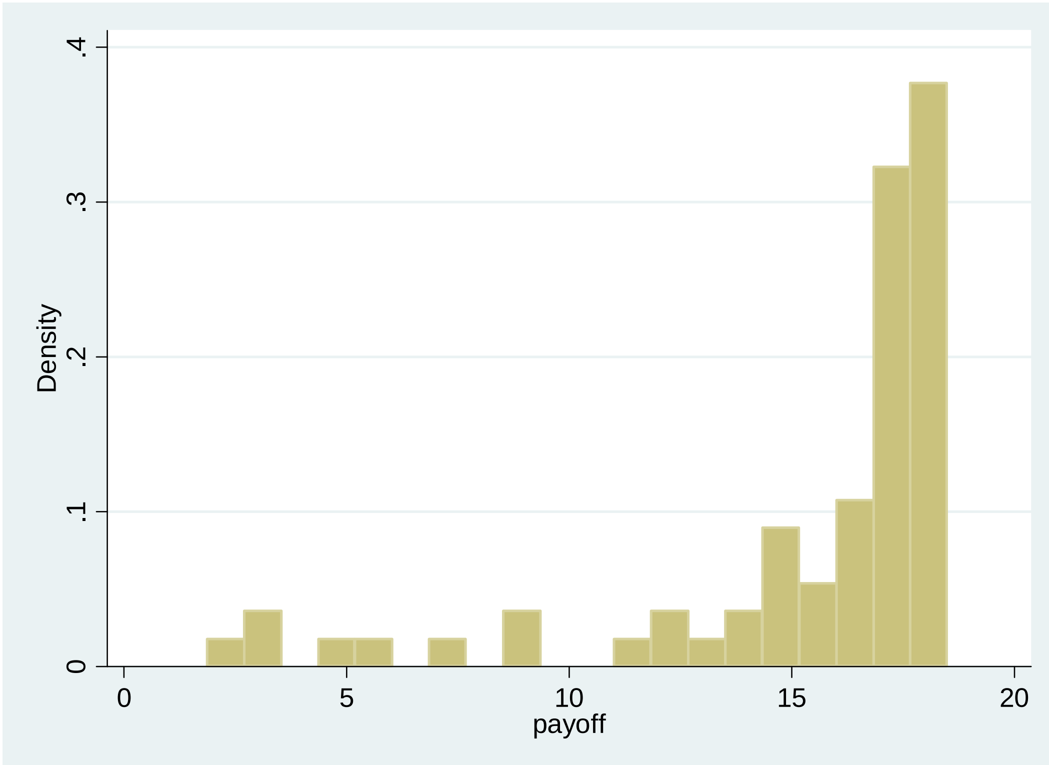

Figure 9: Zero-Skewness Logarithmic Transformation

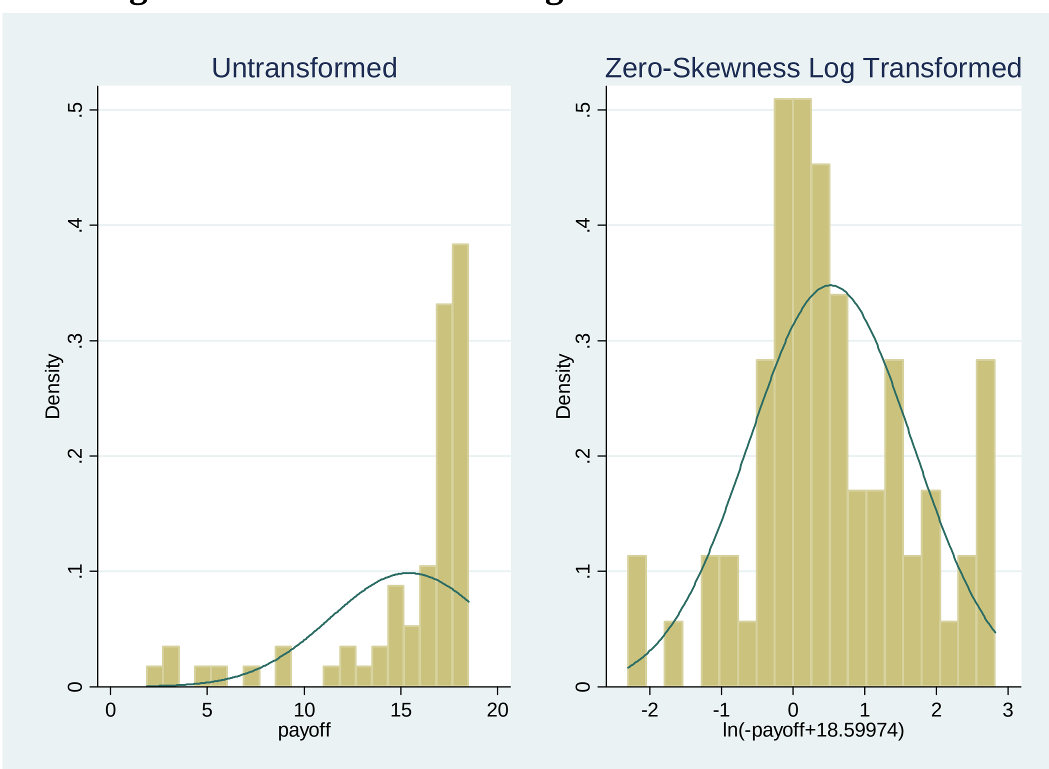

Figure 10: Example of Subjects' decisions

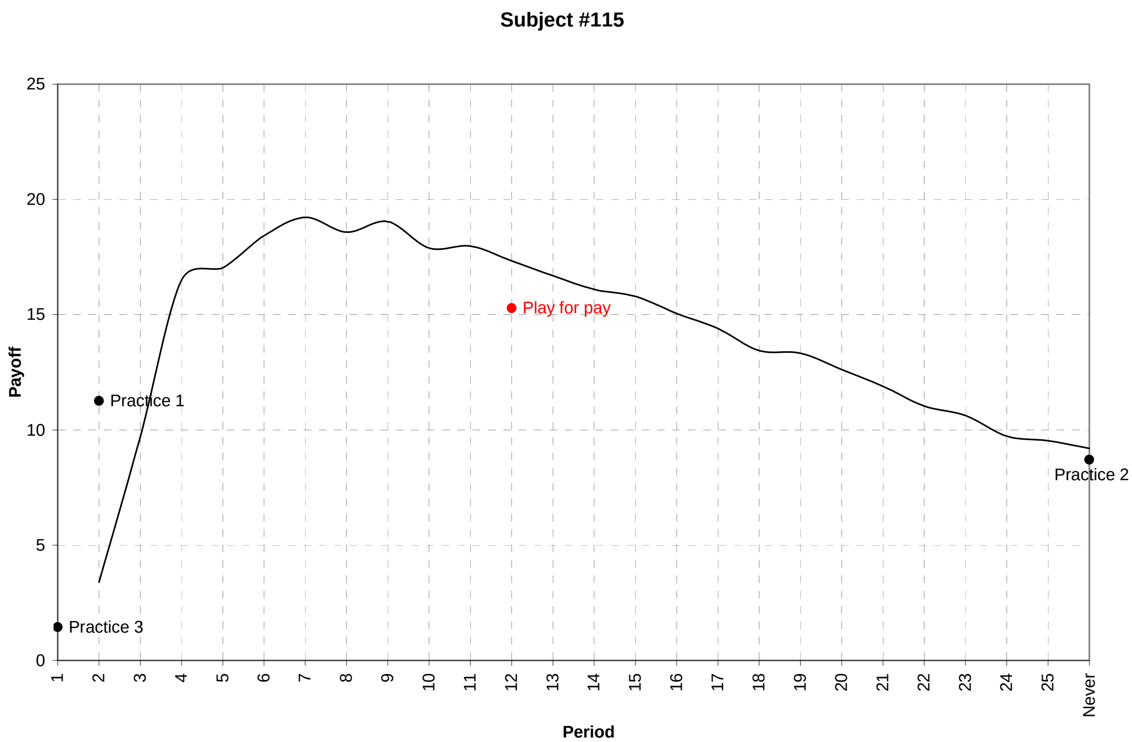

Figure 11: Example of Subjects' decisions

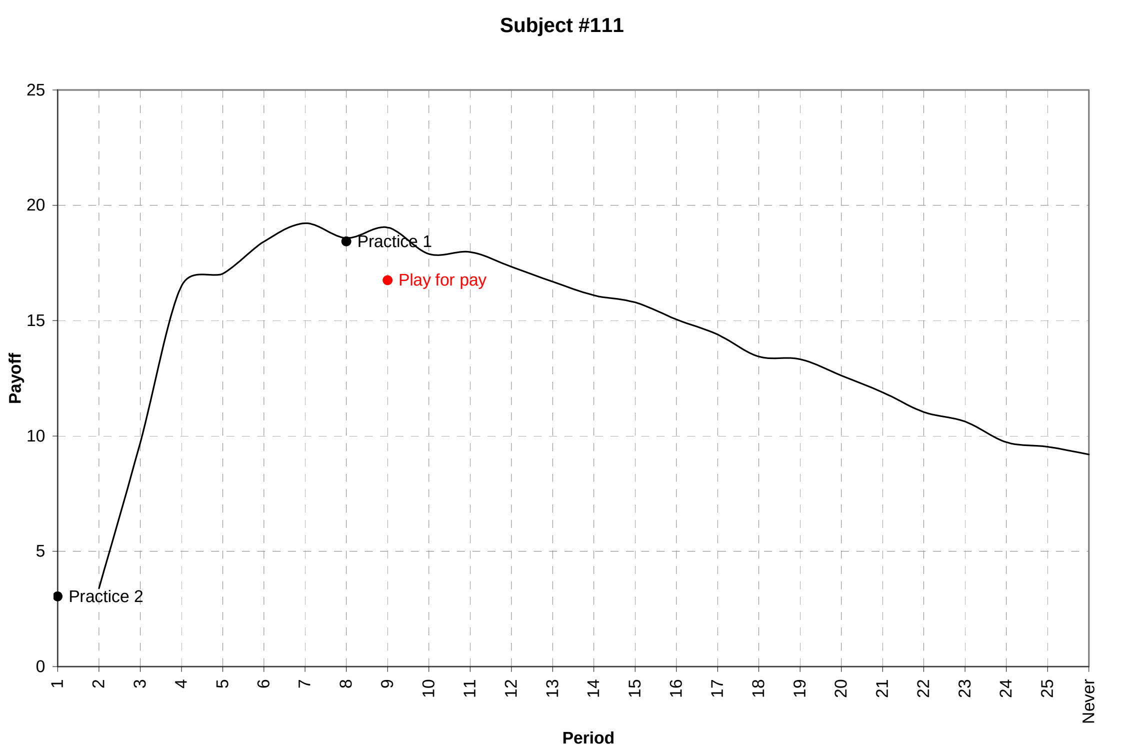

Figure 12: Difference in Best Practice Strategy and Strategy Played for Pay