CIRANO
Allier savoir et décision

2016s-46

Learning-by-Doing in an Ambiguous Environment

Jim Engle-Warnick, Sonia Laszlo

Série Scientifique/Scientific Series

2016s-46

Learning-by-Doing in an Ambiguous Environment

Jim Engle-Warnick, Sonia Laszlo

Série Scientifique
Scientific Series

Montréal
Août/August 2016

CIRANO
Allier savoir et décision

Centre interuniversitaire de recherche en analyse des organisations

CIRANO

Le CIRANO est un organisme sans but lucratif constitué en vertu de la Loi des compagnies du Québec. Le financement de son infrastructure et de ses activités de recherche provient des cotisations de ses organisations-membres, d'une subvention d'infrastructure du ministère de l'Économie, de l'Innovation et des Exportations, de même que des subventions et mandats obtenus par ses équipes de recherche.

CIRANO is a private non-profit organization incorporated under the Quebec Companies Act. Its infrastructure and research activities are funded through fees paid by member organizations, an infrastructure grant from the ministère de l'Économie, de l'Innovation et des Exportations, and grants and research mandates obtained by its research teams.

Les partenaires du CIRANO

Partenaires corporatifs
Autorité des marchés financiers
Banque de développement du Canada
Banque du Canada
Banque Laurentienne du Canada
Banque Nationale du Canada
Bell Canada
BMO Groupe financier
Caisse de dépôt et placement du Québec
Fédération des caisses Desjardins du Québec
Gaz Métro
Hydro-Québec
Industrie Canada
Intact
Investissements PSP
Ministère de l'Économie, de l'Innovation et des Exportations
Ministère des Finances du Québec
Power Corporation du Canada
Rio Tinto
Ville de Montréal

Partenaires universitaires
École Polytechnique de Montréal
École de technologie supérieure (ÉTS)
HEC Montréal
Institut national de la recherche scientifique (INRS)
McGill University
Université Concordia
Université de Montréal
Université de Sherbrooke
Université du Québec
Université du Québec à Montréal
Université Laval

Le CIRANO collabore avec de nombreux centres et chaires de recherche universitaires dont on peut consulter la liste sur son site web.

ISSN 2292-0838 (en ligne)

Learning-by-Doing in an Ambiguous Environment*

Jim Engle-Warnick†, *Sonia Laszlo*‡

Abstract

We apply an instrument to measure ambiguity preferences in an experiment and show that revealed ambiguity preferences, but not risk preferences, predict behavior in a separate game that involves exploitation vs. exploration of a maximization problem. We provide direct evidence of ambiguity preferences acting on decision making separately from risk preferences, and advance knowledge regarding how ambiguity preferences operate on decision-making.

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

* We thank Claude Montmarquette, Peter Moffatt, Daniel Parent, and participants at the Economics Science Association European Meetings in Nottingham. We are especially grateful to the editor and an anonymous referee for their useful comments. 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.

† McGill University and CIRANO.

‡ McGill University and CIRANO. Corresponding author: 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.

1. Introduction

Since Ellsberg's (1961) hypothetical findings, a large amount of work has been done to establish theoretical foundations of ambiguity preferences, to replicate and extend Ellsberg's findings, and to establish empirical regularities with regard to decision making under ambiguity. This paper is an empirical test of the ability of revealed ambiguity preferences to predict choices in a different domain.

In this paper we seek evidence for a link between ambiguity preferences and decisions in games in a laboratory experiment. Any study of ambiguity in decision-making must start with a theory. There are several, many of which are non-encompassing. They include the max-min expected utility theory of Gilboa and Schmedler (1989) and the α max-min expected utility model of Ghirardato, Maccheroni and Marinacci (2004); and the models of recursive expected utility that operate on beliefs regarding the likelihood of possible outcomes (Ahn (2008), Halevy and Feltkapm (2005), and Klibanoff, Marinacci, and Mukerji (2005)). Another approach (Abdellaoui, Baillon, Placido, and Wakker (2011)) models beliefs with a subjective probability weighting function based on Prelec (1998).

Our method is to first introduce an instrument, based on the theory of Klibanoff, Marincacci, and Mukerji (2005), and then to determine to what extent the instrument can predict behavior in a separate ambiguous environment. We chose this theory for its simplicity of application, because the instrument we develop based on it is analogous to the risk instrument of Eckel and Grossman (2008), and because its basis of unknown probability distributions translates well to many other decision-making environments, including the game we present in the next phase. What matters most in this paper is the link between the theory-based instrument and the subsequent game more than the choice of the particular theory itself.

After measuring risk and ambiguity preferences, we next invite the subjects back to the laboratory to play a game that illustrates an important aspect of decision-making under ambiguity. The learning-by-doing model of Jovanovic and Nyarko (1996) allows us, within the experimental control of the laboratory, to seek evidence whether ambiguity can operate on decision-making through a channel separate from risk. It does this by presenting an exploitation vs. exploration situation to the subjects with an ambiguous optimal choice. Although the learning in the model is Bayesian, unlike in other games such as bandit games, we are able to remove the Bayesian updating from the subjects' decision-making problem so as not to confound updating capabilities with ambiguity aversion. Our experimental design reveals whether ambiguity or risk preference can capture the preference to explore in a situation where the probability distribution of outcomes over strategies is not known. We hypothesize that it should.

We find heterogeneity in both the risk and the ambiguity preferences of our subject population. More importantly we find that the ambiguity preferences, and not risk preferences, help to explain choices in the learning-by-doing game. Ambiguity averse subjects are more likely to be willing to pay to explore the problem before selecting their strategy. Using this first laboratory ambiguity instrument with a scale formally derived from the ambiguity model of Klibanoff, Marinacci, and Mukerji (2005) in combination with the learning-by doing- model, we thus contribute to our understanding of how ambiguity operates on decision-making separate from risk.

Our paper provides an important new twist on the empirical relevance of ambiguity preference. Important work currently takes at least three approaches, and there is at times a common theme of separating ambiguity from risk. The first approach focuses on the design of better instruments as has been done for risk preference (Eckel and Grossman (2008); Holt and Laury (2002)). Krahnen, Ockenfels and Wilde (2014), for example, measure ambiguity several different ways and find it to be uncorrelated with risk. Gneezy, Imas and List (2015) also distinguish carefully between ambiguity and risk. A second approach explores the effect of the environment on the ambiguity preference. Charness, Karni, and Levin (2013), for example, examine the ability of people to be persuasive toward others in making their choices, and Engle, Engle-Warnick, and Laszlo (2011) examine the malleability of the revealed ambiguity preference to an experience in a chat room. Keller, Sarin, and Sounderpandian (2007) show that subject dyads make different choices than individual.

A third approach, found here, searches for structural evidence for choices under ambiguity, including a basis for the preference itself or a link between the preference and decision-making. Chew, Ebstein, and Zhong (2012), for example, find evidence for correlations between receptor genes and choices under ambiguity. Engle-Warnick, Escobal, and Laszlo (2011) correlate ambiguity aversion with new technology choices with Peruvian farmers. Anagol et al. (2010) find evidence for ambiguity aversion among trick-or-treaters.

This paper brings several of these elements together to add to the foundation of empirical evidence for how ambiguity operates on choices. In it we present an instrument for ambiguity and control for risk preferences with an analogous instrument. We observe behavior in a game that should reveal ambiguity averse behavior, while circumventing the difficulties with field results including the lack of control of the environment as well as possible omitted explanatory variables. Our results strengthen the field results, and give some confidence that what is being measured by the instruments is empirically relevant.

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 consists 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 (2008). This instrument, shown in Figure 1, is simply a 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 decompose 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), allows 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.[1] The pages were randomly ordered, as was the left to right presentation of the gambles, and the instructions were given orally.[2] 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

[1] 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.

[2] All instructions for the experiments reported in this paper are available in an online appendix.

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 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:

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

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. This model is advantageous because it is tractable.[3]

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:

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

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.

---

[3] The theoretical basis behind our ambiguity instrument is consistent with expected utility theory, and can be roughly thought of as a way to model aversion to compound gambles. This tractable theory is convenient for the development of field instruments of ambiguity preference, where the instrument must be quick, easy to understand, and easy to administer. There is a parallel literature on non-expected utility ambiguity theory. See Camerer and Weber (1992) for a survey of the approaches and Abdellaoui, Baillon, Placido, and Wakker (2011) for recent experimental results.

2.6 Constructing the Risk and Ambiguity Indices: The Explanatory Variables

The index for risk preference is simply the number of safe choices made in our risk instrument. This is also the measure used by Holt and Laury (2002). Thus the instrument is increasing in risk aversion. Similarly, we simply count the number of times subjects paid to avoid the ambiguous gamble in the ambiguity instrument as our measure of ambiguity aversion. Thus the index is increasing in ambiguity aversion.

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. The problem involves the choice when to switch to the better technology.

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:

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

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:

$$y_t = \theta_n + w_t \quad (2)$$

where w_t is normally distributed i.i.d. with zero mean.

The technologies are linked through θ_n:

$$\theta_{n+1} = \sqrt{\alpha}\theta_n + \varepsilon_{n+1} \quad (3)$$

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:

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

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. Hypothesis

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.

Our main behavioral conjecture is simple. In the learning-by-doing game, the probability distribution over the rank order of strategies with regard to outcomes in unknown, i.e., ambiguous. In the instrument, ambiguity aversion is revealed by paying for information about the probability distribution over outcomes. This leads in a straightforward way to our first behavioral conjecture.

Conjecture 1: *Holding risk preference constant, the number of times a subject pays to practice in the learning-by-doing game is increasing in ambiguity aversion. Paying to practice reveals information about the probability distribution over outcomes in the game.*

There is possibly a second way to interpret the game, which is not what the design intended, i.e., 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.*

The experimental design is intended to test Conjecture 1. Conjecture 2 is possible if the subjects interpret the game as risky rather than ambiguous.

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).[4] 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. We now turn to the results from the learning-by-doing experiment.

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

[4] 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.

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.[5]

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 less likely to pay to practice. We show here that the decisions are correlated in the directions predicted by the conjectures.

[5] For another example of this type of result, Engle-Warnick and Turdaliev (2010) 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.

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 0.2061. While these raw correlations are small in magnitude, their signs correspond to the conjectures. Table 1 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 1 also reports the marginal effects from this exercise. More risk averse individuals are more likely to 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 two or three times.[6]

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 1 empirically corroborate our theoretical conjectures. We next use the preference measures to predict performance in the game.

5.4 Explaining Payoffs to Learning by Doing

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:

$$y = \boldsymbol{X}'\beta + \varepsilon \qquad (4)$$

[6] 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.

where $\boldsymbol{X}$=[RM, AM, NP, $\boldsymbol{Z}$], $\boldsymbol{Z}$ a vector of control variables and ε a random disturbance term. In columns (1) and (2) of Table 2 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 more have lower payoffs. Specifically, subjects who practiced four times made lower earnings than those who never practiced (never practiced is the omitted category). Ambiguity aversion does not affect payoffs in the game.[7] 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 2 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 sign(y) and k are to be estimated.[8] 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):

$$\ln(\pm y + k) = \boldsymbol{X}'\beta + \varepsilon \qquad (5)$$

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

[7] As for the results in Table 1, 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.

[8] 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 sign(y) and k.

the zero-skewness logarithmic transformation. To get the marginal effect of the independent variables of interest on payoffs, we must retransform the model. We follow Duan (1983) and Abrevaya (2002) and apply Duan's smearing estimator:

$$\hat{y} = \frac{1}{n}\sum_{i=1}^{n}\left(e^{X'\hat{\beta}+\hat{\varepsilon}_i} - k\right) \tag{6}$$

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

$$m_j\left(x_0, \hat{\beta}\right) = \frac{\hat{\beta}_j}{n}\sum_{i=1}^{n} e^{X'\hat{\beta}+\hat{\varepsilon}_i} \tag{7}$$

We evaluate $m_j\left(x_0, \hat{\beta}\right)$ at the mean values of the ***X***'s. The standard errors are calculated by bootstrap and 500 replications. These marginal effects are presented in the last two columns of Table 2. 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 \$-0.929. With the controls, the marginal effect increases in magnitude to \$-1.411. There is a large and significant negative marginal effect for practicing the game four times (from a low of -\$4.717 to a high of -\$6.009). 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.[9]

[9] We re-ran this analysis replacing the dependent variable with a ranked index for the switch point from highest to lowest expected payoff. Our main results are robust to this alternative analysis, and provide additional evidence that subjects are optimizing their earnings. We thank an anonymous referee for suggesting this analysis.

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 presented evidence that ambiguity attitudes operate separately from risk attitudes using a laboratory experiment. We introduced a ambiguity instrument that is similar to the existing risk instrument of Eckel and Grossman (1996), and used results from the instrument to predict choices in a learning-by-doing game. We found that the ambiguity preferences, and not risk preferences, help to explain behavior in the learning-by-doing game. Specifically, in an exploration vs. exploitation situation, ambiguity preferences help to predict who will explore.

Using this first laboratory ambiguity instrument with a scale formally derived from the ambiguity model of Klibanoff, Marinacci, and Mukerji (2005) in combination with the learning-by doing- model, we contributed to our understanding of how ambiguity operates on decision-making separate from risk. Our results add to a line of results that explore the empirical effect of ambiguity preference on choices. For example, Charness, Karni and Levin (2013) find that ambiguity neutral subjects are persuasive for non-neutral subjects to become neutral. Keller, Sarin and Sounderpandian (2007) find ambiguity aversion in a group setting. Chew, Ebstein, and Zhong (2014) correlate the presence of receptor genes with ambiguity aversion.

Much is known about decision-making under uncertainty. Similar data are now being organized regarding decision-making under ambiguity. Our results show that ambiguity preferences are separate from risk, and add new evidence regarding how such preferences translate into decisions.

References

Abdellaoui, M., A. Baillon, L. Placido, and P. Wakker. (2011) "The Rich Domain of Uncertainty: Source Functions and Their Experimental Implementation." *American Economic Review* 101:695-723.

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

Ahn, D. (2008) "Ambiguity Without a State Space," *Review of Economic Studies* 75:3-28.

Anagol, S., S. Bennett, G. Bryan, T. Davenport, N. Hite, D. Karlan. (2010) "There's something about ambiguity: evidence from Halloween" Yale working paper.

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

Camerer, C., and M. Weber. (1992) "Recent Developments in Modeling Preferences: Uncertainty and Ambiguity." *Journal of Risk and Uncertainty* 5:325-370.

Charness, G., E. Karni, and D. Levin (2013), "Ambiguity Attitudes and Social Interactions: An Experimental Investigation," *Journal of Risk and Uncertainty* 46:1-25.

Chew, S., R. Ebstein, and S. Zhong (2012), "Ambiguity Aversion and Familiarity Bias: Evidence from Behavioral and Gene Association Studies," *Journal of Risk and Uncertainty* 44:1-18.

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 (2008), "Forecasting Risk Attitudes: An Experimental Study using Actual and Forecast Gamble Choices,"*Journal of Economic Behavior and Organization* 68:1:1-17

Ellsberg, D. (1961), "Risk, Ambiguity and the Savage Axioms," *Quarterly Journal of Economics* 75:643-669.

Engle, J., Engle-Warnick, J. and S. Laszlo (2011), "Social Exchange and Risk and Ambiguity Preferences," CIRANO Working Paper 2011s-46.

Engle-Warnick, J., Escobal J., and S. Laszlo (2009), "How Do Additional Alternatives Affect Choice Under Uncertainty?" *Canadian Journal of Economics* 42:1:113-140.

Engle-Warnick, J., Escobal J., and S. Laszlo (2011), "Ambiguity Aversion, and Portfolio Choice in Small Scale Peruvian Farming," *B.E. Journal of Economic Analysis and Policy* 11:1 Article 68.

Engle-Warnick, J. and N. Turdaliev (2010), "An Experimental Test of Taylor-Type Rules with Inexperienced Central Bankers," *Experimental Economics* 13:2:146-166.

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.

Ghirardato, P., F. Maccheroni, and M. Marinacci (2004), "Differentiating Ambiguity and Ambiguity Attitude," *Journal of Economic Theory* 118:133-173.

Gilboa, I., and D. Schmeidler (1989), "Maxmin Expected Utility with Non-unique Prior," *Journal of Mathematical Economics* 18: 141-153.

Gneezy, U., A. Imas and J. List (2015), "Estimating Individual Ambiguity Aversion: A Simple Approach," CESifo working paper no. 5220.

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, Göttingen: Ges. für Wiss. Datenverarbeitung, 79-93.

Halevy, Y., and V. Feltkamp (2005), "Bayesian Approach to Uncertainty Aversion," *Review of Economic Studies* 72:449-466.

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.

Keller, R., R. Sarin, and S. Sounderpadian (2009), "An Examination of Ambiguity Aversion: Are Two Heads Better than One?" *Judgement and Decision Making* 2:390-397.

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

Krahnen, J., P. Ockenfels, and C. Wilde (2014), "Measuring Ambiguity Aversion: A Systematic Approach," working paper.

Prelec, D. (1998), "The Probability Weighting Function," *Econometrica* 66:497-527.

Table 1: Correlation between Ambiguity Aversion Measure and the Number of Times Practiced in the Learning-by-Doing Experiment

	Ordered probit	Marginal effects if practiced…				
		Never	Once	Twice	Three times	Four times
# of safe choices	-0.207**	0.019	0.060*	-0.044*	-0.018	-0.0167
	(0.104)	(0.012)	(0.033)	(0.024)	(0.012)	(0.011)
# of times paid to avoid ambiguity	0.192**	-0.018*	-0.056**	0.041*	0.017*	0.0015
	(0.008)	(0.010)	(0.028)	(0.021)	(0.010)	(0.009)
Predicted Prob. practiced N times		0.043	0.574	0.288	0.057	0.037
F-test for joint significance	7.85*					
Pseudo R^2	0.0943					
Wald χ^2	16.62**					

Regressions include session controls. s.e. in parentheses. *, **, *** signif. at 10%, 5% and 1\%.

Table 2: Predictors of Earnings in the Learning-by-Doing Experiment

	OLS		Zero Skewness Log Transform Marginal Effects	
# of safe choices	-1.124***	-1.706***	-0.929***	-1.411***
	(0.353)	(0.470)	(0.287)	(0.440)
# of times paid to avoid ambiguity	-0.2013	-0.394	-0.288	-0.301
	(0.291)	(0.346)	(0.242)	(0.266)
Practiced once	-2.194	-2.499	-0.769	-1.375
	(1.872)	(2.301)	(1.446)	(1.769)
Practiced twice	0.371	-0.398	0.172	-1.153
	(2.000)	(2.428)	(1.547)	(2.045)
Practiced three times	-0.180	-1.273	-0.567	-1.122
	(2.434)	(2.644)	(1.591)	(1.783)
Practiced four times	-7.411***	-8.990***	-4.716**	-6.009**
	(2.636)	(3.038)	(2.016)	(2.522)
Socio-economic controls	No	Yes	No	Yes
F-test for joint significance	4.25***	4.50***		
Shapiro-Wilks [p-value]	0.0002	0.0008	0.9662	0.9569
R^2	0.3201	0.4269	0.2777	0.4099
Observations	72	69	72	69

s.e. in parentheses. *, **, *** signif. at 10%, 5% and 1%. Marginal effects are evaluated at mean *X*s. Skewness parameter (k)=-18.5997, 95% conf. int. for k= [-18.960,-18.525]

Figure 1: Five Options Risk Preference Measurement Instrument

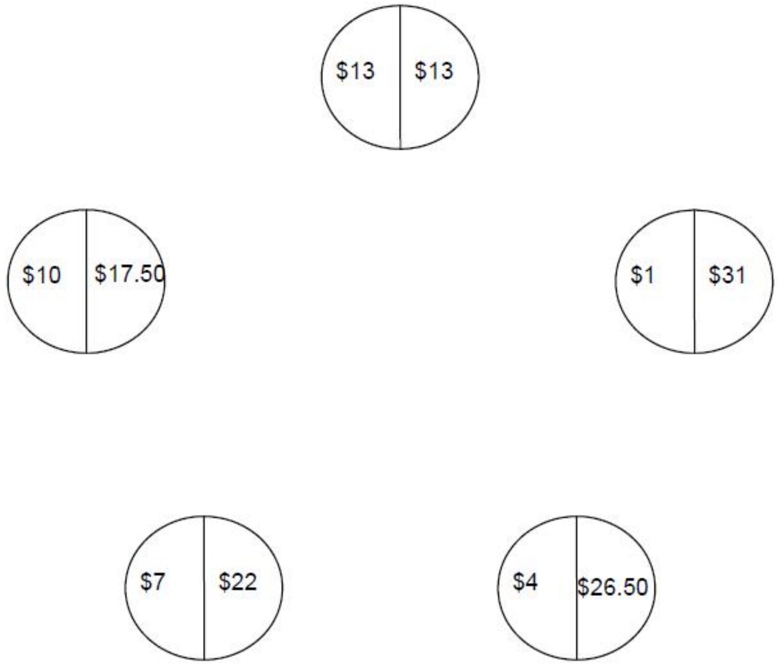

Figure 2: Decomposing Five Options Instrument into a Series of Binary Options Instrument

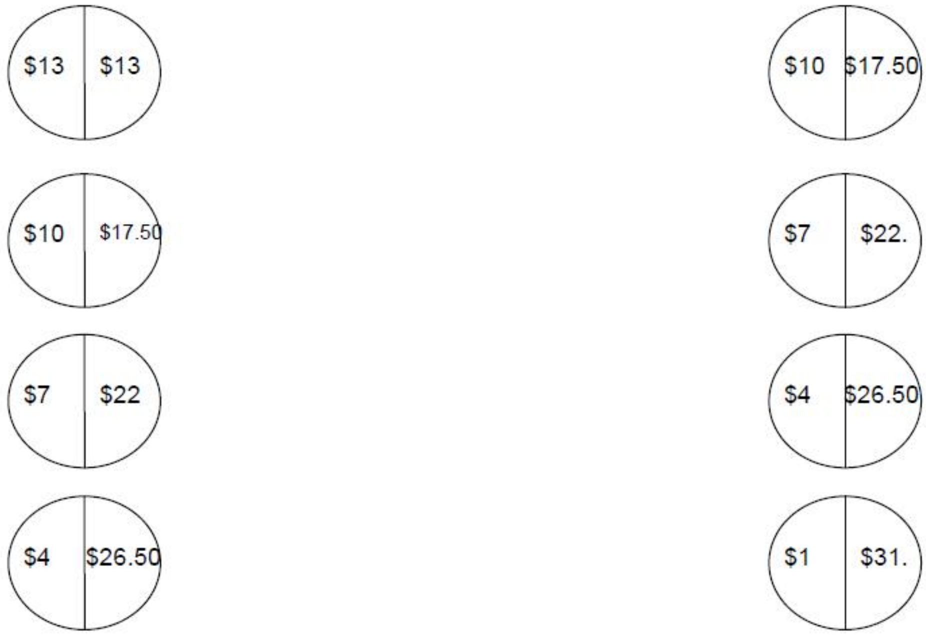

Figure 3: Binary Choices to Reveal Preferences for Ambiguity

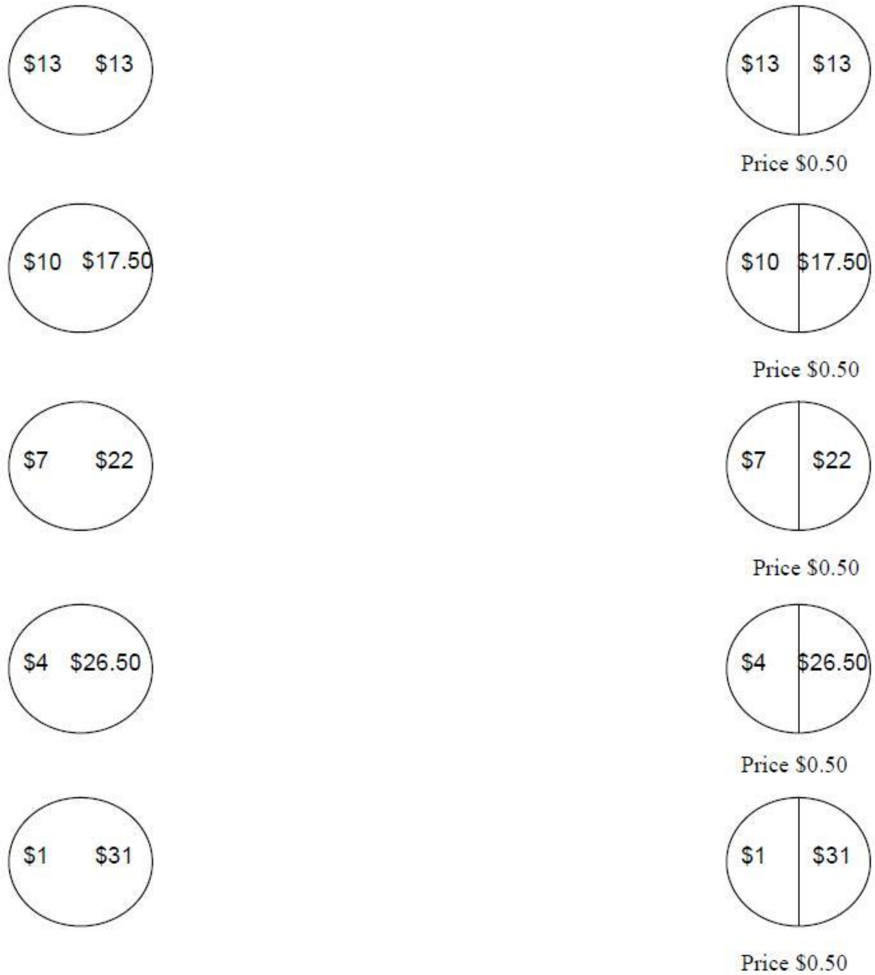

Figure 4: Average Payoff by Switchpoint in the Learning-by-Doing Game

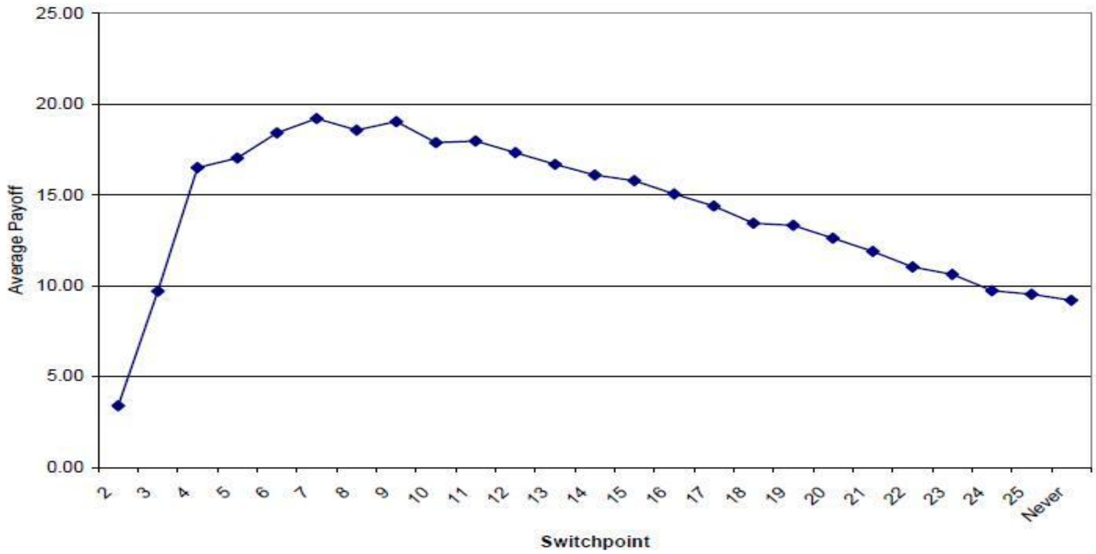

Figure 5: Distribution of Safe Choices in the Binary Game

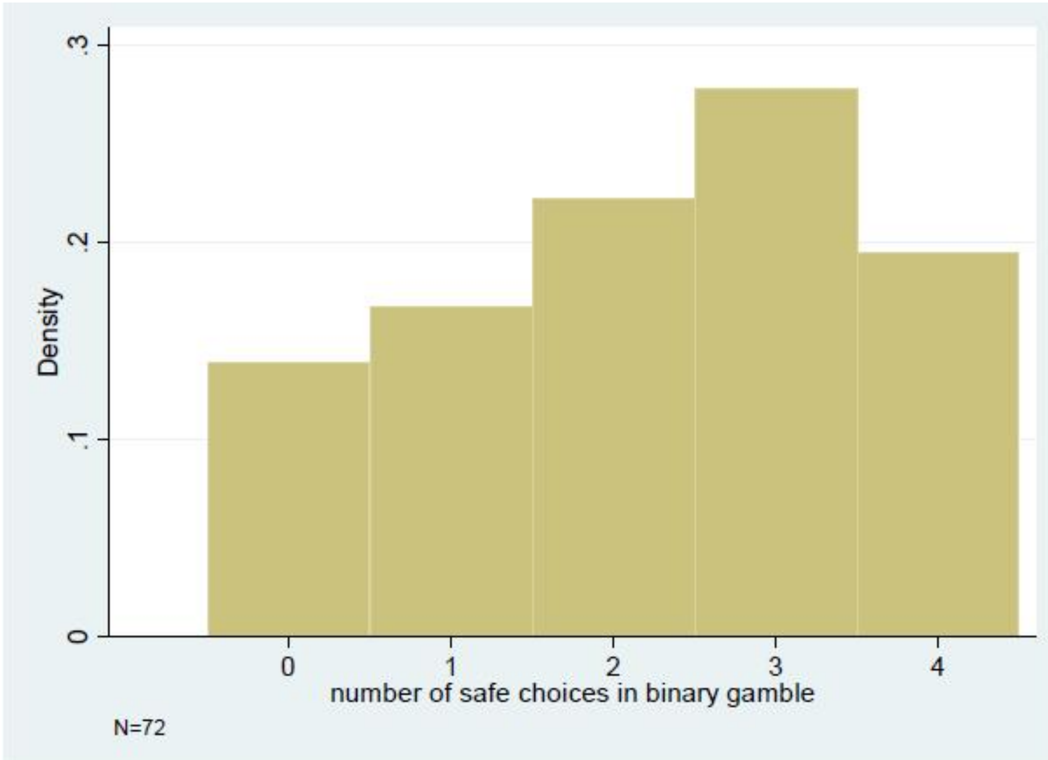

Figure 6: Distribution of the Number of Times Subjects Paid to Avoid Ambiguity}

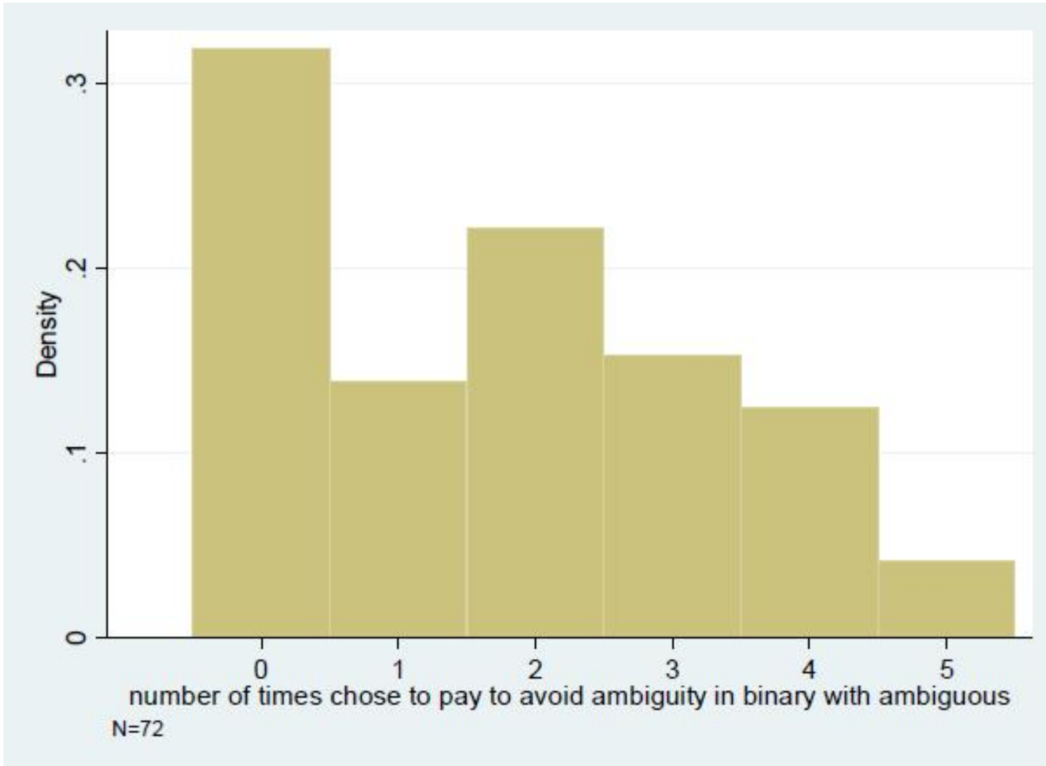

Figure 7: Distribution of the Number of Times Subjects Practiced the Learning-by-Doing Game

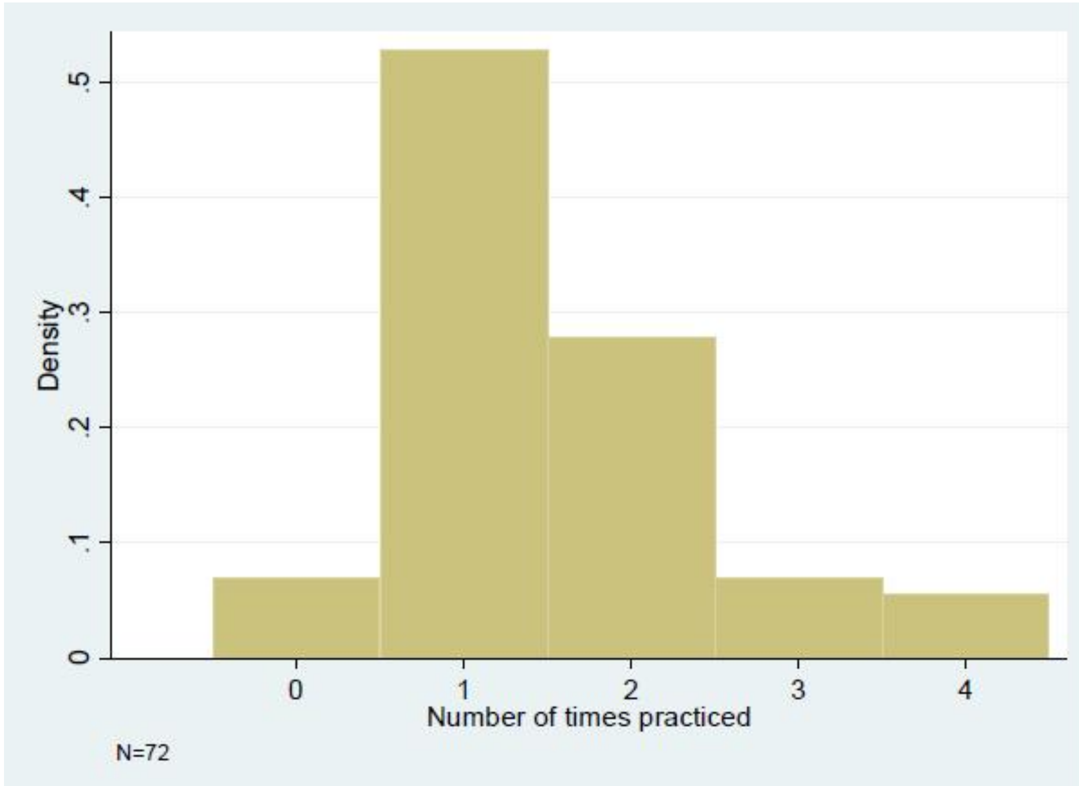

Figure 8: Distribution of the Payoffs in the Learning-by-Doing Experiment

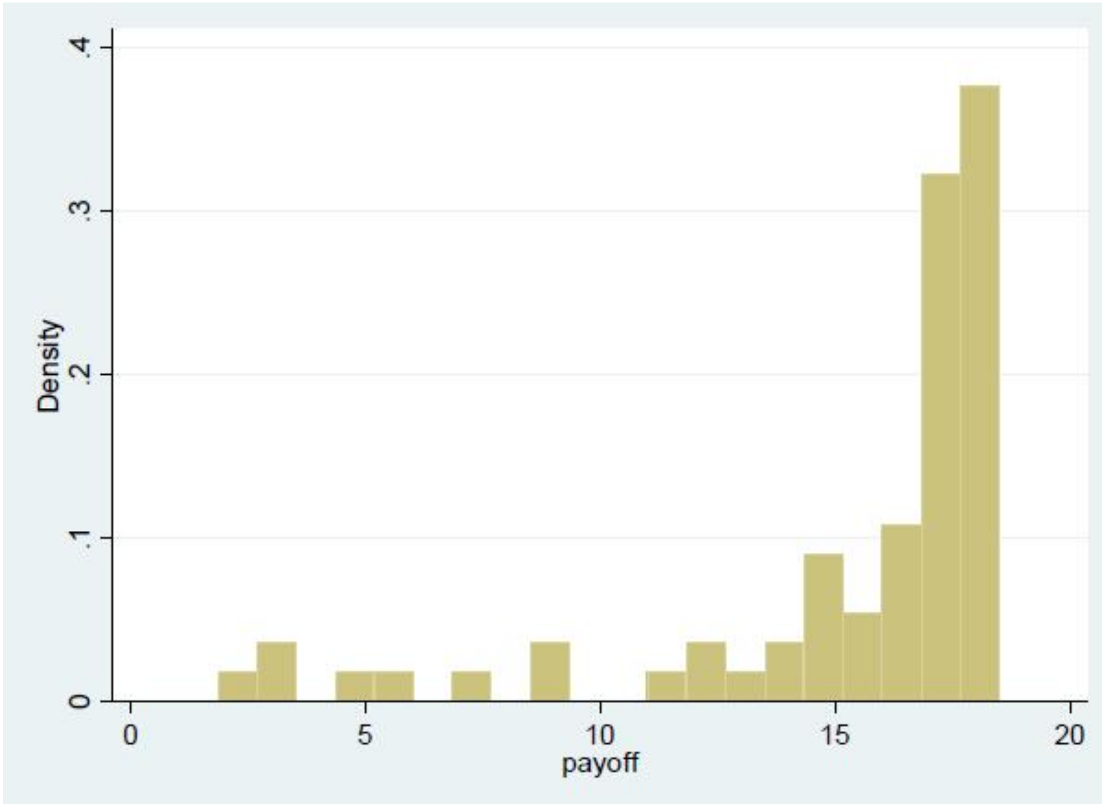

Figure 9: Zero-Skewness Logarithmic Transformation

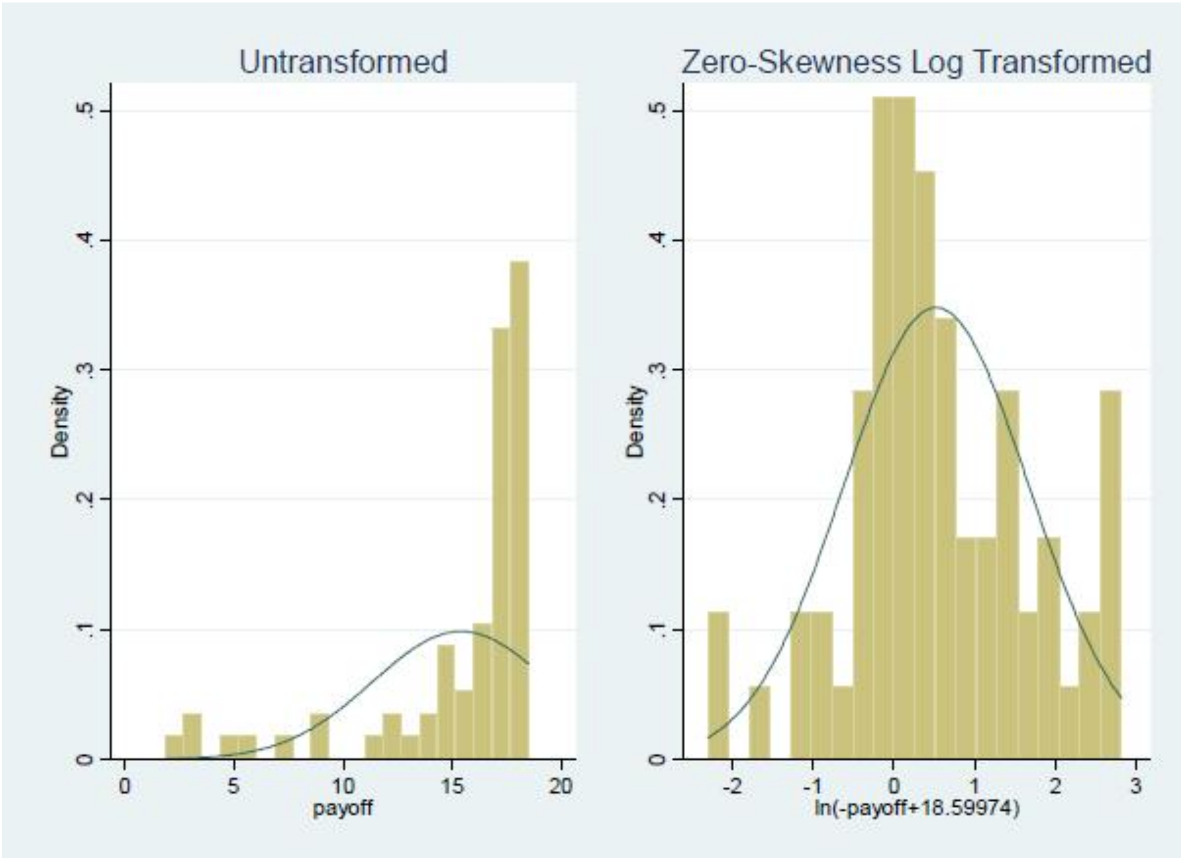

Figure 10: Example of Subjects' decisions (subject 115)

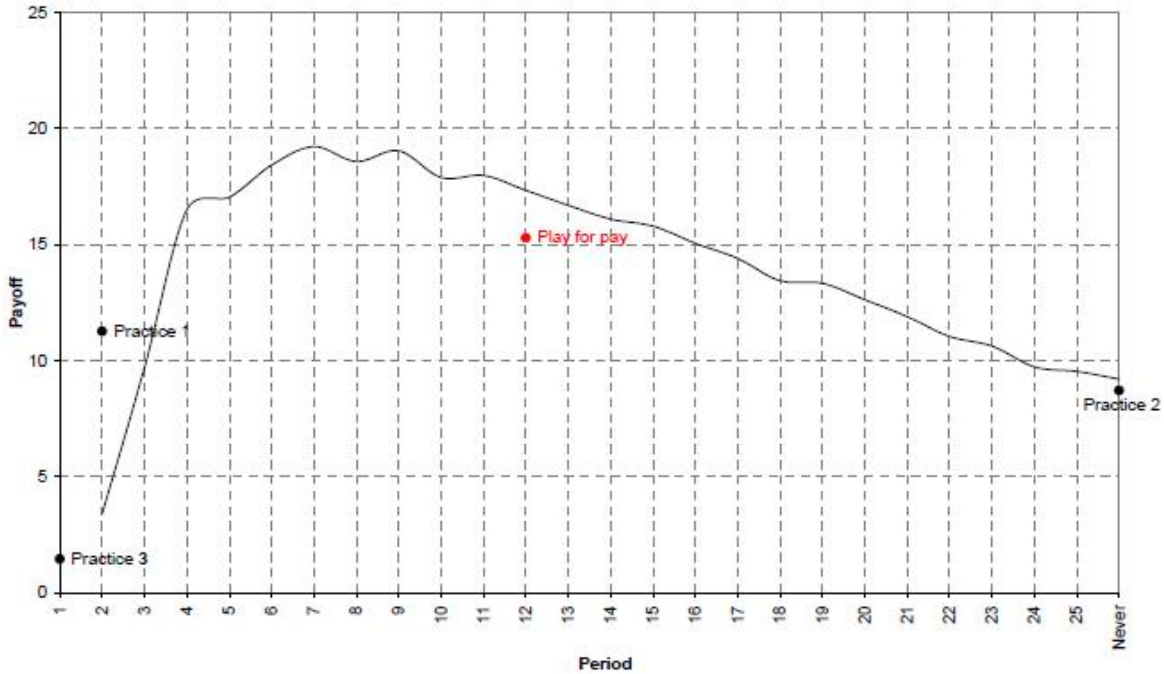

Figure 11: Example of Subjects' decisions (subject 111)

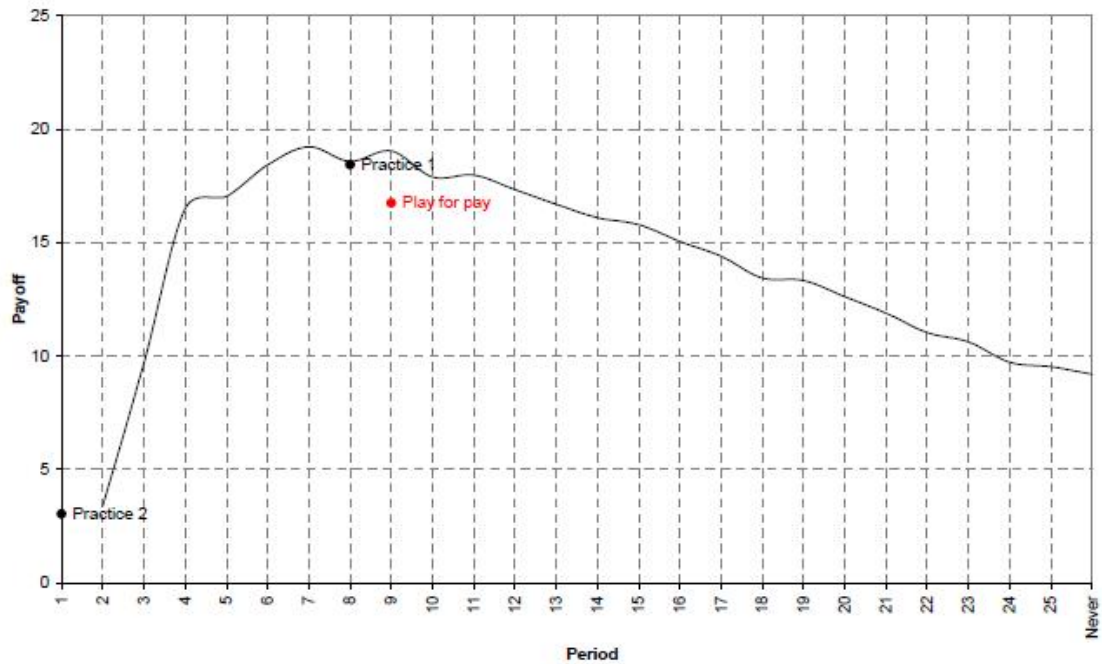

Figure 12: Difference in Best Practice Strategy and Strategy Paid for Pay

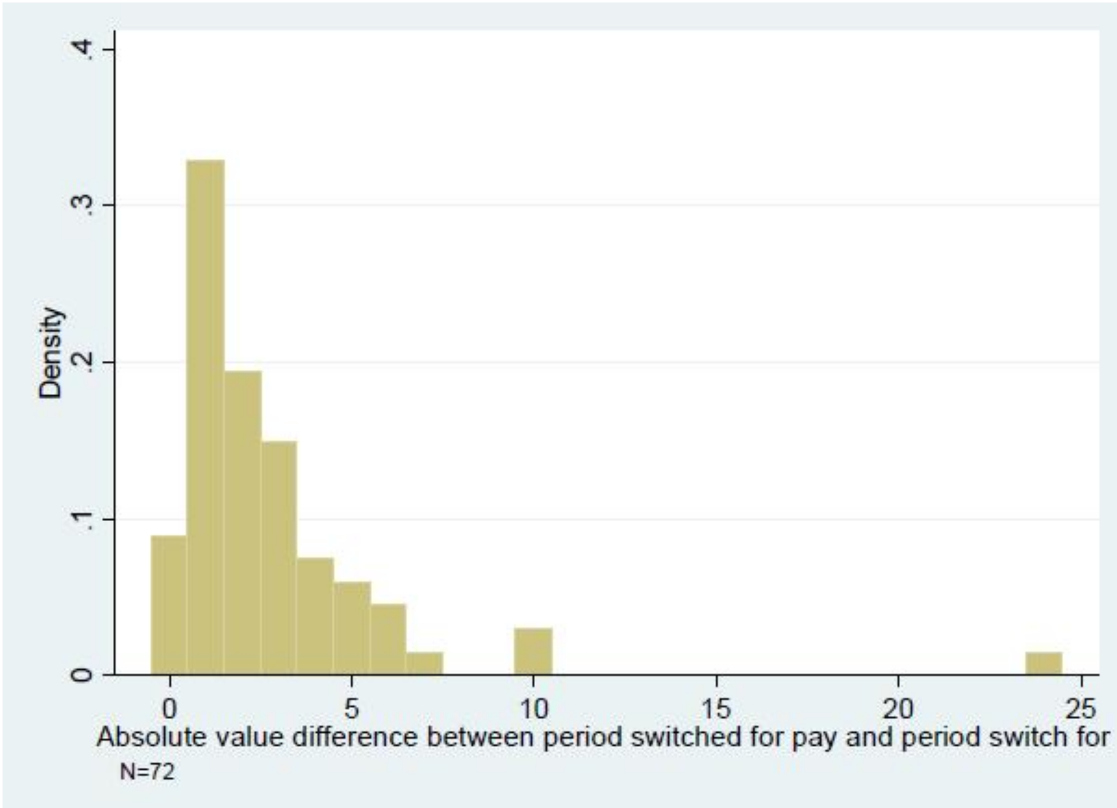

Appendix

Sample Descriptive Statistics (N=72)

	Mean	Standard deviation
Sex (Female =1)	0.347	0.479
Working	0.333	0.475
Highest degree is secondary	0.028	0.165
Highest degree is undergraduate	0.708	0.458
Highest degree is graduate	0.250	0.436
Mother tongue is English	0.361	0.484
Mother tongue is French	0.208	0.409

1130, rue Sherbrooke Ouest, bureau 1400, Montréal (Québec) H3A 2M8
Tél. : 514-985-4000 • Téléc. : 514-985-4039
www.cirano.qc.ca • info@cirano.qc.ca

Centre interuniversitaire de recherche en analyse des organisations
Center for Interuniversity Research and Analysis on Organizations