



HAL
open science

Ambiguous incentives and the persistence of effort: Experimental evidence

Robin M. Hogarth, Marie Claire Villeval

► **To cite this version:**

Robin M. Hogarth, Marie Claire Villeval. Ambiguous incentives and the persistence of effort: Experimental evidence. *Journal of Economic Behavior and Organization*, 2014, 100, pp. 1-19. 10.1016/j.jebo.2014.01.006 . halshs-01098750

HAL Id: halshs-01098750

<https://shs.hal.science/halshs-01098750>

Submitted on 7 Jan 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

WP 1432

**Ambiguous Incentives and the Persistence of Effort:
Experimental Evidence**

Robin M. Hogarth, Marie Claire Villeval

November 2014

GATE Groupe d'Analyse et de Théorie Économique Lyon-St Étienne

93, chemin des Mouilles 69130 Ecully – France

Tel. +33 (0)4 72 86 60 60

Fax +33 (0)4 72 86 60 90

6, rue Basse des Rives 42023 Saint-Etienne cedex 02 – France

Tel. +33 (0)4 77 42 19 60

Fax. +33 (0)4 77 42 19 50

Messagerie électronique / Email : gate@gate.cnrs.fr

Téléchargement / Download : <http://www.gate.cnrs.fr> – Publications / Working Papers

Ambiguous Incentives and the Persistence of Effort:

Experimental Evidence

Robin M. Hogarth and Marie Claire Villeval

Abstract: When the assignment of incentives is uncertain, we study how the regularity and frequency of rewards and risk attitudes influence participation and effort. We contrast three incentive schemes in a real-effort experiment in which individuals decide when to quit: a continuous incentive scheme and two intermittent ones, fixed and random. In all treatments, we introduce a regime shift by withdrawing monetary rewards after the same unknown number of periods. In such an ambiguous environment, we show that less able and more risk averse players are less persistent in effort. Intermittent incentives lead to a greater persistence of effort, while continuous incentives entail exit as soon as payment stops. Randomness increases both earlier and later exiting. This selection effect in terms of ability and risk attitudes combined with the impact of intermittent rewards on persistence lead to an increase in mean performance after the regime shift when incentives are intermittent.

Keywords: Incentives, intermittent reinforcement, randomness, effort, quitting, learning, experiment

JEL Classifications: C92, D84, M54, J31

Contact: Robin M. Hogarth, Department of Economics and Business, Universitat Pompeu Fabra and Barcelona Graduate School of Economics, Ramon Trias Fargas, 25–27, 08005 Barcelona, Spain. E-mail: robin.hogarth@upf.edu
Marie Claire Villeval, University of Lyon, F-69007; CNRS-GATE, 93 Chemin des Mouilles, F-69130, Ecully, France. E-mail: villeval@gate.cnrs.fr

Acknowledgments: The authors are grateful for comments from the editor and two anonymous reviewers, Gary Charness, David L. Dickinson, Dirk Engelmann, Lorenz Goette, Donald A. Hantula, Natalia Karelaia, Peter Kuhn, Alvin Roth, Aldo Rustichini, Karl Schlag and participants at the World conference of the Economic Science Association in Copenhagen, the IMEBE conference in Barcelona, at the EMIR workshop on behavioral economics of incentives in Lyon, and at seminars at Universitat Pompeu Fabra, the Universities of Guelph, Heidelberg, Nijmegen, Rennes, Sydney, Tilburg, Monash University in Melbourne, Erasmus University in Rotterdam, and Paris School of Economics. The authors are particularly grateful to Lingyi Tan. They also thank Romain Zeiliger for programming the experiment. Financial support from the Agence Nationale de la Recherche (ANR BLAN07-3_185547 “EMIR” project) is gratefully acknowledged. Hogarth gratefully recognizes the financial support of the Spanish Ministerio de Economía y Competitividad (Grant ECO2012-35545). The revision of this paper was performed within the framework of the LABEX CORTEX (ANR-11-LABX-0042) of Université de Lyon, within the program “Investissements d’Avenir” (ANR-11-IDEX-007) operated by the French National Research Agency (ANR).

1. INTRODUCTION

Ambiguity is pervasive. In many economic activities agents learn the states of underlying processes by observing outcomes across time. For example, an investor in the stock market may observe fluctuations in the value of his shares. What is signaled by a drop in the value of the shares: a change in the fundamental value of the assets or just an insignificant random event? Should the investor sell his shares or keep them? Similarly, if an employee no longer receives a bonus in addition to her salary, should she consider that this is due to temporary circumstances or to a permanent change? Should she change her reference point?

In this paper we explore ambiguous situations where economic agents reap the benefits of engaging in an activity across time until – unknown to them – there is a shift in the underlying process and pursuing the activity is no longer profitable. Knowing that agents dislike income variability, we seek to illuminate when and why they persist or desist from participating in the activity when they bear all the cost of uncertainty. How long, for example, will the investor continue to place new orders and does this depend on the regularity of his previous outcomes? How long will an employee keep working in the same firm if she no longer receives a bonus? How is the decision affected by preferences regarding risk and ambiguity and/or the regularity with which bonuses have been paid in the past? Knowing when, if, or how the parameters governing an underlying process have changed is, however, not a trivial intellectual task. When agents engage in repeated activities across time, one typically assumes that learning results from feedback and beliefs are updated in a Bayesian manner. However, the structure of incentives might influence this process when the environment is ambiguous. For example, Bereby-Meyer and Roth (2006) have shown that the speed of learning in repeated games is affected by whether incentives involve probabilistic or continuous reinforcement.

Our goal is to draw attention to the likelihood that organisms will persist or desist in effort before and after regime shifts when the payment of rewards is ambiguous, conditional on individuals' ability and risk attitudes, and how the structure of incentives affects persistence and effort over time. Our hypothesis is that such an ambiguous environment will impact whether individuals are willing to persist or not in the activity and how long, according to their individual cost of effort, *i.e.*, ability level, risk/ambiguity aversion, and expected income. Another hypothesis is that the frequency and the regularity of incentives, or equivalently their concentration or spreading over time, influence persistence in the task because they condition the speed at which individuals update their beliefs. They therefore affect the identification of regime shifts in ambiguous environments.

For example, contrast two ambiguous situations where the same total amounts of non-contractual bonuses have been distributed to workers in addition to their salary prior to a regime shift but where afterwards no more bonuses are paid. In one group, the firm has distributed bonuses both regularly and frequently (*e.g.*, on a monthly basis). This corresponds to a *continuous* incentive scheme. In the second group, workers have been used to receiving less frequent or more concentrated but bigger bonuses in the same interval of time (*e.g.*, every six months on average). This corresponds to an *intermittent* incentive scheme. We conjecture that regardless of the scheme, both more risk or ambiguity averse workers and less able workers will exit this situation earlier than others in order to avoid ambiguity. Indeed, the less able workers, who have a higher marginal cost of effort, have a lower expected payoff and therefore a higher risk aversion than abler workers with a higher expected payoff. Thus, it is likely that these workers require higher risk premia and are more likely to drop out as soon as possible, before the regime shift, because their expected payoffs do not allow them to bear the full cost of uncertainty. Moreover, we conjecture that workers exposed to regular and frequent incentives will perceive a

regime shift more rapidly than those exposed to irregular and less frequent incentives. They will exit more rapidly after the regime shift than those exposed to intermittent incentives, because even with the same degree of loss aversion and the same reference payoff, they are able to update their beliefs more rapidly due to a different structure of incentives.

In this paper we investigate experimentally the persistence in effort in an ambiguous environment according to the individuals' ability level and degree of risk and ambiguity aversion, and the differential effects of continuous and intermittent (random and fixed) incentive schemes before and after the introduction of a regime shift. Participants in our experiment were required to perform a tedious task repeatedly and to decide when to stop performing the task definitively. As such, goal setting was endogenous. The task itself involved counting specific letters in written paragraphs where words had been combined randomly to form meaningless sentences. We compare three treatments in a between-subjects design. In the Baseline treatment, participants were paid a piece-rate and each period gave rise to actual payment. This corresponds to a continuous incentive scheme. In the Random Intermittent Incentive treatment, for each participant there was a one out of three chance that any particular period would be paid and we determined individually randomized sequences of payment for all these participants. In the Fixed Intermittent Incentive treatment, one period out of three was paid according to a fixed schedule. In all three treatments, payment was stopped after the 20th period. To introduce ambiguity, the participants were informed of the amount of the piece-rate but they received no *ex ante* information about the sequences of payments. They were only informed at the beginning of the session of their show-up fee and that some periods would be paid whereas others would not, and at the end of each period, whether it was paid or not. We also elicited participants' attitudes towards risk and ambiguity.

Our findings show that under ambiguity, more risk averse and less able participants are more likely to quit the task as soon as possible, *i.e.* before the regime shift, regardless of the treatment probably because they dislike having to bear all the risk. As well-known in contract theory, maintaining them longer would require insuring them, at least partly, against income variability. Moreover, in the Baseline treatment where it takes less time to identify the regime shift individuals exit soon after payment stops. This may be due to the large immediate gap between their income expectations and their updated beliefs about receiving a payment, which is likely perceived as a loss (see Koszegi and Rabin, 2006, for a model of reference-dependent preferences based on expectations, and see Crawford and Meng, 2011, and Abeler *et al.*, 2011, for tests of this theory). On the other hand, participants in the Random Intermittent Incentive treatment exit either before or long after the regime shift, with the Fixed Intermittent Incentive treatment in an intermediate position. Therefore, under ambiguity, both the frequency and randomness of incentives matter in the decision to exit a situation, with randomness having more extreme effects in both directions. The gap between the individuals' income expectations and their updated beliefs about the likelihood of payment is smaller than in the Baseline. Moreover, intermittent incentives lead to more persistent and higher effort levels. Indeed, when incentives are continuous, once payment stops individuals who do not exit immediately exert less effort than those under the other schemes.

The remainder of this paper is organized as follows. In section 2, we present the related literature. In section 3 we introduce our experimental design and state our theoretical predictions. Section 4 develops our results and section 5 discusses these results and concludes.

2. RELATED LITERATURE

Given their interest in incentive schemes, it is surprising that economists have not been more concerned with the effects of ambiguity, frequency and variability of incentives.¹ An early exception is Tustin and Morgan (1985) who proposed an optimization theory of choice when workers have to assign effort in the presence of concurrent fixed or uncertain reinforcement schedules that can be delivered after a specified period of time. More recently, Lazear (2006) studied whether it is more efficient to spread or concentrate incentives and to provide or not provide information. He shows that this depends on the responsiveness to incentives, and therefore on the cost of effort. Informing individuals about incentives (like speed inspections or school testing) dominates stochastic incentives for individuals who have a high cost of effort. For example, predictable tests provide the best incentives for less talented children to learn; less predictable tests are better for high ability students but result in no learning for less talented ones because their cost is too high. Like in this study, we are interested in the impact of intermittent incentives and ability on effort. In contrast, in our study individuals are not informed about the distribution of incentives over time.

There is also a recent literature examining the impact of ambiguity on incentives. For example, Kellner and Reiner (2012) study how ambiguity aversion influences the choice of payment schemes by agents. Ederer *et al.* (2013) investigate theoretically how the design of ambiguous incentive schemes with randomization may help to solve the gaming problem in a

¹ In contrast, there is an old literature in psychology on the impact of variable reinforcement schedules on behavior (see Jenkins and Stanley, 1950; Ferster and Skinner, 1957; Robbins, 1971; Hilgard and Bower, 1975). For example, Deslauriers and Everett (1977) showed that, to increase bus ridership, it was more effective to give tokens to roughly one in three as opposed to every passenger. In Golz (1992) participants exposed to continuous reinforcement in the acquisition phase invested less in the extinction phase –when rewards are withdrawn– than those who had experienced variable reinforcement. Extensions of this paradigm by Hantula and colleagues (Hantula and Crowell, 1994; DeNicolis Bragger *et al.*, 2003; Brecher and Hantula, 2005) showed that variable reinforcement schedules induce responses that are more immune to extinction than their continuous counterparts. We differ from this literature in that we use real monetary incentives and we control for preferences toward risk and ambiguity. Moreover, we study how risk attitudes mediate the impact of intermittent incentives on participation and effort.

principal-agent framework when agents try to exploit incentives in their own interest. The agents ignore which one of two types of ambiguous incentive schemes is used when they decide on their effort level. This ambiguity leads them to balance effort across tasks instead of concentrating effort on their preferred task. It imposes of course more risk on the agent. We also examine the impact of ambiguity in rewards on effort but we focus on participation decisions, we manipulate the frequency and regularity of incentives and we introduce a regime shifts in their payment.

Also related to our analysis, studies have investigated how the allocation of effort reacts to unexpected changes in incentives over time. This change is usually a variation in the level of compensation, not in its structure. For example, Camerer *et al.* (1997) explore income targeting by cabdrivers to explain the negative response of hours worked to an increase in daily wages. Others have suggested that variations of motivation within a workday could explain this negative correlation (Goette *et al.*, 2004; Goette and Huffman, 2006). A static model of reference-dependent preferences (Koszegi and Rabin, 2006, 2007) in which the individual's reference point is his rational expectation generated by recent events can also explain this behavior. In particular, Koszegi and Rabin show that regarding labor supply, due to loss aversion, a worker will be willing to quit earlier if he has received an unexpected gain earlier, but he will be more likely to continue if he has received a high expected income. Evidence of the link between reference points in expectations and effort provision is provided by Abeler *et al.* (2011) and Crawford and Meng (2011). In particular, Abeler *et al.* (2011) show that effort reacts to differences in earnings expectations in a task in which people are paid randomly with a fixed wage or a piece-rate. Individuals quit later when the fixed payment is higher. We also predict that a worker who does not receive a reward in a period will be more likely to quit the task in the next period if he expects to be paid (based on the previous sequence of payments) than if he anticipates that all periods are not paid. But in contrast with these studies, we consider ambiguity and focus on the

consequences of intermittent *vs.* continuous incentives on the mode of belief updating before and after a regime shift occurs. We also propose a dynamic model of belief updating.

Studying learning and incentives,² Bereby-Meyer and Roth (2006) compare two treatments of a prisoner’s dilemma game, one in which payoffs are deterministic and another in which they are probabilistic, holding the expected payoffs constant. Randomness in the link between a decision and its consequences slows the speed of learning, as cooperation decays more slowly when payoffs are noisy. Others who have shown awareness of possible effects of the structure of incentives include Lazear (1990), O’Flaherty and Komaki (1992), Eriksson *et al.* (2009), and Erev *et al.* (2010). We differ from this literature because we do not compare stochastic and deterministic incentives, as individuals are not aware *ex ante* of the frequency of incentives. Instead, we let them form expectations about the distribution of incentives and adjust their participation decision and effort accordingly.

3. EXPERIMENTAL DESIGN AND THEORETICAL PREDICTIONS

3.1. Experimental treatments

Our experiment consists of three treatments and is based on a between-subjects design. In the *Baseline*, each participant has to perform a task that consists of counting the occurrence of four different letters in a paragraph. This task is repeated in different periods across time. In each period, a different paragraph is displayed on the participant’s screen for two minutes. The paragraph consists of words that are randomly combined to form four meaningless sentences. The

² Cognitive neurosciences have also studied how the brain deals with unexpected outcomes and learns from them. In the ‘reward prediction error’ models, a release of dopamine (a neurotransmitter involved in behavior and cognition) signals to the brain the need to adjust when there are discrepancies between predicted and actual rewards (Schultz *et al.*, 1997; Montague *et al.*, 1996; Caplin and Dean, 2008). After repeated trials with a reward, the omission of an expected reward induces a negative prediction error and a depression of the dopamine neurons which encode this prediction error (for recent surveys on the neural basis of reinforcement learning, see Niv and Montague, 2009, and Balleine *et al.*, 2009).

letters to be counted appear successively. All the participants receive the same paragraphs in the same order. The difficulty of the task is comparable across periods. An answer is considered correct if it corresponds to the true value plus or minus one. The participants are immediately informed if a submitted answer is correct or not. The current number of their correct answers is displayed continuously. Advantages of this task are that it is tedious, it does not require any prior expertise and it leaves little room for learning.

Participants are informed that they have to complete a minimum of 15 periods after which they have to decide when to quit the task. The mean performance during these 15 periods gives us a measure of the individual's ability. It is made clear in the instructions that the end of the experiment is decided by the subject. There is a maximum of 35 periods (*i.e.*, 70 minutes) but this information is not made common. In other words, the participants can exit voluntarily between the 16th and the 35th periods. They are informed that not all periods are paid but they do not know in advance whether a period will be paid or not. Nor are they informed about the total number of paying periods. At the end of each period, they learn *ex post* whether correct answers provided in that period are or are not paid. The instructions specify that their performance does not affect the selection of a period for payment.

In the baseline treatment, all correct answers given in each of the first 20 periods are rewarded monetarily at the end of the session. After the 20th period, we introduce a regime shift (but do not announce it): the correct answers no longer give rise to paying points. When periods are paid, each correct answer is rewarded with one paying point, such that between 0 and 4 points can be earned in a period. There is no penalty for providing wrong answers. As soon as the participant decides to start a new period, an entry fee of 1 point has to be paid, whether correct answers are or are not paid in the period. Therefore, whereas payment of the entry fee is certain,

the reward is uncertain. We adopted this procedure to encourage participants to think more deeply about their entry and exit decision.

To summarize, although players have to enter the first 15 periods, after period 15 each participant must decide whether to start the period or not before the beginning of every subsequent period. If he or she starts the period, the participant is only informed whether the period gives rise to payment or not after the two minutes allowed for the period have elapsed. If the player decides not to start a new period, he or she has to confirm this decision and exit the task. Exiting is irreversible.³

The *fixed intermittent incentive treatment* (Fixed II, hereafter) is identical to the Baseline except that, during the first 20 periods, one period out of three is paid (periods 2, 5, 8, 11, 14, 17 and 20). In addition, in the paid periods each correct answer is rewarded by three paying points instead of one in the baseline treatment to keep earnings comparable at the end of the compulsory periods.⁴ Here too, no payment is made after the 20th period and this is not common information.

The *random intermittent incentive treatment* (Random II, hereafter) is identical to the Fixed Intermittent Incentive treatment except that, instead of paying one out of three periods in a regular, deterministic fashion, we pay exactly seven periods in the first 20 periods but which periods are paid are randomly determined for each subject. To facilitate comparisons between treatments, we impose the constraint that correct answers are paid for the 20th period (not made

³ Our design dissociates the period in which we allow people to exit the game (beginning of period 16) and the period in which we introduce the regime shift (beginning of period 21). Introducing the regime shift and the exit option at the same time could have introduced confounding factors and a possible experimenter demand effect. This distinction allows us to distinguish the impact of ambiguity from the impact of changes in the fundamental conditions. Finally, it is possible to control econometrically for the possible selection effect of subjects quitting between periods 16 and 21, as will be shown below.

⁴ In the baseline treatment, participants can earn a maximum of 80 paying points in the session (4 correct answers*1 point*20 paying periods), while in the two other treatments, the maximum is slightly higher, i.e., 84 (4 correct answers*3 points *7 paying periods). Therefore, the incentive to exert effort may be slightly higher in the intermittent incentive schemes. To make treatments identical on this point, we could have used a piece-rate of 2.86 instead of 3 in the intermittent incentive treatments. We reasoned, however, that this piece-rate would not be easily interpreted by the participants. In addition, since participants did not know in advance the sequence of payments, we do not believe that the difference we indicate affected behavior.

common knowledge). As before, no payment occurs after period 20. As in the Fixed II treatment, correct answers are paid three times more than in the Baseline. In the three treatments, the same paragraphs appear in the same order to keep the difficulty of the task constant across treatments.

We elicited attitudes towards risk and ambiguity at the beginning of the session by asking participants to price both a clear and a vague bet following a procedure similar to that of Fox and Tversky (1995).⁵ The random draws are done only at the end of the session in order not to influence behavior in the game. In all treatments, when participants decide not to start a new period or when the final period has been reached, they have to answer questions as to why they did or did not decide to exit. A demographic questionnaire follows.

3.2. Predictions

Before stating predictions about effort and exiting decisions, we consider the process of belief formation. Indeed, a critical input to effort and persistence in the task is the subjective expected probability that a period will be remunerated. In contrast with static models of reference-dependence based on expectations, we consider a dynamic adjustment of subjective beliefs that a period will be paid. At the beginning of each period, we assume that players assess a probability that the period will be paid based on their experience up to that point. We further assume that individuals are boundedly rational in that their judgments do not comply fully with Bayesian reasoning but with a less cognitively demanding Bayesian-like process.⁶ Assessments result from

⁵ The participants make a first set of 20 decisions between accepting a certain payoff and drawing a ball in an urn that contains 5 blue balls and 5 yellow balls (the lottery is risky). The amount of the certain payoff increases from 0.25 Eurocents to €5. One yellow ball drawn from the urn pays €5, a blue ball pays nothing. Then, the participants make a second set of 20 similar decisions except that the proportions of yellow and blue balls are now unknown (the lottery is ambiguous). In both sets of decisions, a risk neutral participant should choose to draw a ball until the certain payoff is equal to at least €2.5. An ambiguity averse (seeking) participant should switch from the lottery to the certain payoff for lower (higher) certain amounts when the lottery is ambiguous than when it is risky. One of the two sets is randomly drawn for real payoffs.

⁶ We did not elicit participants' estimates of the probability that effort in the next period will be rewarded because the use of proper incentives for these judgments would have confounded our ability to identify the reasons for persistence in the task. Because of hedging, participants who understood that there was no task payoff after period 20

a sequential anchoring-and-adjustment process whereby a participant's judgment is anchored on her previous assessment and updated by the experience of the latest period. More formally, we propose the model:

$$S_k = S_{k-1} + w_k[s(x_k) - S_{k-1}] \quad (1)$$

where S_k is the belief that the next period will be paid after experiencing k periods ($0 \leq S_k \leq 1$); S_{k-1} is the belief that the previous period would be paid; $s(x_k)$ reflects the latest evidence, that is, whether the k^{th} period was ($s(x_k)=1$) or was not paid ($s(x_k)=0$); and w_k is an adjustment parameter that determines how much the previous assessment is modified to account for the latest evidence. Eq. 1 is a special case of the model developed by Hogarth and Einhorn (1992) in the context of predicting order effects in belief updating. Here, assessments are weighted averages of past beliefs and the latest evidence. As such the model is structurally similar to the EWA model of reinforcement learning investigated by March (1996) and Camerer and Ho (1999). It can be thought of as an approximation to Bayesian updating except that it implies inconsistencies (*e.g.*, order effects) that are incompatible with Bayes. The adjustment parameter, w_k , should be sensitive to both the sign of the impact of the evidence, $[s(x_k) - S_{k-1}]$, and the level of the preceding belief, S_{k-1} . For $[s(x_k) - S_{k-1}] \leq 0$, w_k is assumed to be proportional to S_{k-1} . Larger anchors (*i.e.*, values of S_{k-1}) should be more sensitive to the impact of negative evidence than smaller ones, consistent with neuroscientific evidence. Thus,

$$w_k = \alpha S_{k-1} \text{ when } [s(x_k) - S_{k-1}] \leq 0 \quad (2)$$

Similarly, when $[s(x_k) - S_{k-1}] > 0$, w_k is assumed to be inversely proportional to S_{k-1} :

$$w_k = \beta(1-S_{k-1}) \text{ when } [s(x_k) - S_{k-1}] > 0 \quad (3)$$

where $0 \leq \alpha, \beta \leq 1$.

might still decide to stay until the end to earn prediction payoffs. Using instead a non-incentivized belief elicitation procedure would have produced biased and non-reliable estimates. For a study comparing the accuracy of incentivized vs. non-incentivized beliefs, see notably Trautmann and Kuilen (2011).

The model can therefore be written as:

$$S_k = S_{k-1} + \alpha S_{k-1} [s(x_k) - S_{k-1}] \text{ when } [s(x_k) - S_{k-1}] \leq 0 \quad (4)$$

and

$$S_k = S_{k-1} + \beta(1-S_{k-1})[s(x_k) - S_{k-1}] \text{ when } [s(x_k) - S_{k-1}] > 0 \quad (5)$$

The magnitudes of α and β are assumed to reflect the variability of incentives in the evidence experienced. When experienced variability is low, individuals will put more weight on the most recent evidence compared to situations where this is greater. This implies larger values of α and β . The reason is that, with low variability, not much information is needed to make an assessment. As variability increases, more past observations are needed and more weight is placed on the anchor, S_{k-1} . Thus, we hypothesize that α and β are greatest for the Baseline and larger for the Fixed II than the Random II treatment. That is:

$$\alpha_{Baseline} > \alpha_{FixedII} > \alpha_{RandomII} \text{ and } \beta_{Baseline} > \beta_{FixedII} > \beta_{RandomII}.$$

Figure 1 simulates belief curves for $\alpha_{Baseline} = \beta_{Baseline} = 0.8$; $\alpha_{FixedII} = \beta_{FixedII} = 0.5$; and $\alpha_{RandomII} = \beta_{RandomII} = 0.3$ over the 35 periods.⁷ The data for the Baseline and Fixed II treatments are calculated exactly whereas those of the Random II treatment are the result of averaging the outcomes associated with simulating ten random sequences.

(Insert Figure 1 about here)

The model predicts different patterns of beliefs across periods for the three treatments.⁸ In the Baseline, belief increases in the early periods but drops rapidly after payment stops and stays below the levels of the other treatments through period 35. In the Fixed II treatment, beliefs

⁷ One could, of course, also consider situations where $\alpha > \beta$ or $\alpha < \beta$ and variations of differences between the values of the parameters by treatments. We have systematically calculated several such variations to test the sensitivity of the model's predictions to our choices for α and β . In short, the model's predictions concerning the relative patterns of beliefs in the different treatments both prior to and after the regime shift are quite robust.

⁸ We acknowledge that Figure 1 is an inaccurate representation of the differences in the beliefs for the very early periods in that the differences between treatments for α and β would not be apparent until after experiencing several periods.

oscillate with the pattern of payment until after period 22 when there is a decline until the end of the game. As to the Random II treatment, averaging random draws, there is a constant pattern through period 20 after which there is a decreasing trend toward the end of the game. However, beliefs are higher than those in the other treatments after the regime shift. Here we assume that players update their beliefs sequentially but do not recognize temporal patterns. We present in Appendix A an alternative model that also implies sequential updating but is sensitive to the recognition of temporal patterns. It predicts the same patterns as those observed in Figure 1 and delivers the same predictions. It is possibly more realistic but less parsimonious.

Let us now consider predictions on effort and exiting decisions for before and after the regime shift. These predictions reflect not only participants' expectations that they will be paid but also their expected earnings.

Before the regime shift, exiting decisions are expected to depend mainly on individuals' risk attitudes and ability levels. Indeed, in our design the agents bear all the uncertainty associated with the payment of the output, which risk averse individuals will dislike. To some extent, it is as if they had to bear all the uncertainty of the luck component of output. Risk or ambiguity averse individuals are, therefore, more likely than risk-neutral subjects to exit the game earlier (*i.e.* before the regime shift), regardless of the treatment, because they need a higher risk premium to satisfy their participation constraint (on the trade-off between risk and incentives, see notably Prendergast, 1999). In contrast, provided the participation constraint is satisfied, risk should not affect effort in our design since the expectation of a piece-rate has a motivating effect. Moreover, participants with a higher marginal cost of effort should provide a lower effort for the given piece-rate compared to high ability participants, and therefore make lower earnings (when paid) than them. Given their lower expected payoff per period, less able participants should exhibit higher risk aversion. This should also motivate them to quit earlier than more able

participants. We also expect that before the shift, the exit decisions should be similar in the Baseline and the Fixed II treatment since expected payoffs based on beliefs should be similar. In the Random II treatment, we may expect some early exit decisions if the random draws concentrate all the bonuses in the early periods of the game. Finally, effort may differ across treatments if individuals are able to anticipate in the Fixed II treatment that after a period is paid no payment will be received in the next two periods.

After the regime shift, in contrast, exiting decisions are expected to depend mainly on the treatment. Indeed, the sorting by risk and by ability level should have already occurred as soon as exit was allowed. Individuals should exit just after the regime shift when they were previously continuously reinforced. Indeed, at the moment of the regime switch, participants in the Baseline may feel a strong disutility, as the actual payment is null although they still have high expectations, which may create a deep feeling of loss. Moreover, they can update their expectations more quickly than in the other treatments. Exit should occur later in the Intermittent Incentives treatment and especially so in the Random II treatment as shown by the simulation in Figure 1 where beliefs on payment remain always higher in this treatment. Due to the previous sorting by ability, mean performance should be higher after the regime shift than before.

We summarize our main hypotheses as follows:

H1. Before the regime shift, low skill participants and more risk or ambiguity averse individuals exit at a greater rate than high skill and less risk or ambiguity averse ones.

H2. Before the regime shift, exiting behavior is similar in the Baseline and the Fixed II treatments. In the Random II treatment, early exits may occur if the random payments are concentrated at the beginning of the game.

H3. Before the regime shift, in the Fixed II treatment individuals adjust their effort across periods, depending on whether the previous period or the two previous periods were paid or not.

This is also expected after the regime shift, with some noise.

H4. After the regime shift, individuals quit earlier in the Baseline than in the Fixed and Random II treatments, *i.e.* when incentives are continuous instead of intermittent.

H5. After the regime shift, individuals quit later in the Random than in the Fixed II treatment.

H6. Total earnings are higher in the Baseline than in the two intermittent treatments, due to the different exit patterns across treatments.

3.3. Experimental procedures

The experiment consists of 12 sessions conducted at the laboratory of the GATE (*Groupe d'Analyse et de Théorie Economique*) institute in Lyon, France. Between 12 and 19 individuals took part in each session, for a total of 210 participants invited via the ORSEE software (Greiner, 2004). The participants were undergraduate students from the local engineering and business schools. In each treatment, half of the sessions were run at 12 noon and half were run at 2pm, to control for the possibility that exiting decisions could be influenced by differing opportunity costs of time during the day. Table 1 displays summary information about the sessions.⁹

(Insert Table 1 about here)

Five participants were invited in a session every quarter of an hour. Upon arrival, a participant was randomly assigned to a computer terminal by drawing a tag from a bag. Therefore, when entering the laboratory most participants found others already working at the task, so that it was impossible to know how much time others had been in the laboratory. The

⁹ We ran five sessions of the baseline, five sessions of the Random II treatment but only two sessions of the Fixed II treatment (36 independent observations), because we considered the latter to be more of an intermediate treatment. It turned out that the treatment effects were so strong that we did not deem it necessary to collect more observations.

moment when a participant left the laboratory was not informative as to the period when he or she exited the task. This limits the possibility of peer effects on quitting.

Participants found the instructions for both the risky and ambiguous lotteries in their cubicles (Appendix B). Questions were answered in private. Then, after completing their decisions, participants received the instructions for the main task and could start by playing two practice periods to gain familiarity with the task. The experiment was computerized, using the REGATE platform. After deciding to quit and after completing the final demographic questionnaire, participants were informed on their screens to proceed to the payment room. There, random draws were made to determine the earnings in the risk and ambiguity elicitation tasks. Subjects were instructed to avoid disseminating information about the experiment in order to minimize information spillovers across sessions.

The experiment lasted on average 70 minutes, including the payment of participants, and participants earned an average of €14.29 (standard deviation €2.81), including a show-up fee of €4 and the payoff from either the risky or the ambiguous lottery.

4. RESULTS

In this section, we consider successively the exiting decisions, the performance levels and the earnings achieved.

4.1. Exit decisions

Figures 2a, 2b, and 2c show when participants in each of the three treatment groups exit the task. Figure 3 complements this by displaying the percentage of participants remaining after each period following the first 15 compulsory periods. Recall that payment is stopped after period 20.

(Insert Figures 2a, 2b, 2c, and 3 about here)

The figures reveal that the distributions of exiting behavior differ across treatments. In the Baseline, the distribution is symmetrical with a median (and mode) of 17% at the 22nd period (see Fig.2a). Compared to the intermittent incentive treatments, few participants exit the task when it continues to be paid but is no longer compulsory (*i.e.*, between periods 15 and 20). However, the effect of stopping payment is dramatic (see Fig.3). By the end of period 25, only 29% of Baseline participants remain in the task. At the end of the task (end of period 35, when required to stop), only 2% remain.

The distribution of exiting periods is strikingly different in the Random II treatment. The mode of the distribution is at period 35 when 20% of the players have to be forced to stop (Fig.2c). In addition, 36% of the participants exit between periods 15 and 20, while the corresponding percentages are 20% in the Baseline and 25% in the Fixed II treatment. Compared with the other treatments, randomness of incentives is associated with a less steep decline in the number of participants remaining over time, as predicted (see Fig.3).

The Fixed II treatment has a tri-modal distribution: 11% of the participants exit at period 15, 17% at period 26, and 8% at each of the 33rd, 34th, and 35th periods (see Fig.2b). Participants are somewhat slower to exit the task than those in the Random II treatment before period 25, but after this the rate of exiting increases and only 8% remain after period 35 (see Fig.3). The relatively high frequency of exits in periods 26 and 29 in this treatment suggests that, as in the Baseline in which many participants exit in periods 22 and 23, it takes two to three unpaid sequences of periods before many participants decide to quit.

Kolmogorov-Smirnov exact tests (K-S hereafter) show that the distribution of exiting periods is significantly different between the Baseline and both the Random II ($p<0.001$) and the Fixed II treatments ($p<0.001$). However, the distributions of the two Intermittent Incentive treatments are not significantly different ($p=0.222$). Mann-Whitney tests (at the mean values; M-W hereafter)

show qualitatively similar results. The average exiting period is lower in the Baseline than in both the Random ($p=0.052$) and the Fixed II treatments ($p=0.002$) with no difference between the two II treatments ($p=0.693$).¹⁰

To identify the determinants of the decision to exit the task, our econometric analysis considers separately the exiting decisions before the regime shift and after the regime shift. We first estimated a Probit model in which the dependent variable is the probability to be still at work at the end of period 20, *i.e.* before the regime shift (model (1)). Then, we estimated two Tobit models in which the dependent variable is the exiting period. Tobit models are required because the data are censored both on the left and on the right. Model (2) considers only the sub-sample of participants who were still at work at the end of period 20 and who exited after the regime shift, and model (3) includes the full sample.¹¹ The independent variables include dummy variables for each treatment, with the Baseline as the reference category. They include the number of periods paid before period 15 in the Random II treatment that exceeds a predicted number of 5 rewards (since 7 periods are paid before period 21), because players who experienced more paying periods at the beginning of the game might have perceived an earlier regime shift and exited earlier than the other players. The skill variable is the number of correct

¹⁰ Due to the specific distributions of exiting periods across treatments, a simple comparison of mean exiting periods would be quite misleading. The means (standard deviations) for the three treatments (Baseline, Random II, Fixed II, respectively) are: 23.18 (4.00), 25.48 (7.17), and 26.44 (6.90).

¹¹ We do not use duration models to analyze the exiting behavior because there is no variation per period in independent variables after period 20. Moreover, to control for a possible selection bias in the decision to exit before the regime shift, we also estimated a Heckman two-step model. In a first step we estimated with a Probit model the probability to be still working at the end of period 20 and in the second step, we estimated the determinants of the exiting period for those who stay at least 20 periods by means of an OLS model in which we included the Inverse of the Mills' Ratio obtained from the first equation. To identify the model, we dropped from the second equation the variable indicating the number of periods paid before period 15 minus 5 (regressions available upon request). Since the Inverse of the Mill's ratio was not significant, we do not need to use a two-step model. This is why we report in Table 3 separate Probit and Tobit regressions.

answers provided in the 15 compulsory periods.^{12, 13} We also include a dummy variable for the afternoon sessions since participants may have a higher opportunity cost for staying longer in the afternoon than during the lunch break. The rank of arrival variable captures the impact on the participant's exiting decision of observing other participants leaving.¹⁴ The risk index indicates the switch point in the risky lottery; the higher this index, the less risk averse the individual. The ambiguity aversion variable expresses the difference in switch points in the risky and the ambiguous lotteries. A positive value means that the player has switched from the random drawing to the certain equivalent earlier for the ambiguous than for the risky lottery, thereby indicating some ambiguity aversion. A dummy variable captures multiple switching in at least one lottery (22 out of 210 participants).¹⁵ We control for the players' belief about their family wealth relative to their schoolmates on a scale from 1 (among the 10% least wealthy families) to 10 (among the 10% wealthiest families). We include controls for age, gender, number of years of post-secondary education, and level of excellence on high school certificates (from 1 for "no distinction" to 4 for "very good") as a measure of cognitive abilities. The results of these regressions are displayed in Table 2.

(Insert Table 2 about here)

In Model (1), the probability of staying at least 20 periods is increased significantly by both skill and risk seeking preferences. Indeed, less able participants earn less in the paid periods and

¹² We consider the first 15 periods since all the participants play the same number of periods. Lacking systematic records of all performance for the first six sessions of the Baseline and Random II treatments, we calculated the average performance per paid period in the first 15 periods and multiplied this number by 15.

¹³ Note that the skill levels were almost similar across treatments. In the following, the three p -values correspond always to the comparison between the Baseline and the Random II treatment, then between the Baseline and the Fixed II treatment, and last between the Random and the Fixed II treatments. Regarding the score in practice periods, K-S tests indicate $p=0.270$, 0.269 , and 0.800 , while M-W tests yield $p=0.079$, 0.213 , and 0.784 . Regarding the cumulative score at the end of period 15, K-S tests indicate $p=1.000$, 0.975 , and 0.818 , while M-W tests yield $p=0.934$, 0.904 , and 0.977 . Last, considering the cumulative number of mistakes during the 15 periods of the compulsory task, K-S tests indicate $p=0.999$, 0.785 , and 0.594 , while M-W tests yield $p=0.724$, 0.890 , and 0.738 .

¹⁴ The rank of arrival is a better instrument to capture peer effects in exiting than the rank of departure because the former is exogenous while the latter is endogenous to the exiting decision.

¹⁵ In case of multiple switches, we determine risk attitude with the mean value of the switch decisions.

this makes them less willing to take the risk of waiting for another paying period because the likely amount of the reward based on their performance is not high enough to satisfy their participation constraint.¹⁶ This supports hypothesis H1. However, ambiguity aversion has a separate but positive effect on the probability to be working at least until period 20. One interpretation of this surprising finding is that ambiguity averse players stay longer, *ceteris paribus*, if they expect that this will help them become better informed about the returns on their efforts. In line with hypothesis H2, we find no significant difference between the Fixed II and the Baseline treatments. Supporting also H2, more periods paid before period 15 in the Random II treatment increase the likelihood of exiting the game before we stop rewarding effort. This is as if the subjects anticipated a regime shift earlier than the actual one. However, contrary to what was expected, there still exists a significant and negative effect of the Random II treatment on the probability to be still working in period 20. This suggests that some subjects exhibit a higher disutility associated with the randomness of payments that is distinct from risk or ambiguity aversion. Participating in an afternoon session decreases the probability of staying at least 20 periods by 14.4%, which can be explained by a higher opportunity cost of participating in the afternoon than during the lunch break. Last, participants who report being wealthier are more prone to early exiting (-2.7% per decile) as they possibly value less the payoff they can obtain from continuing in the experiment.

Model (2) indicates that, conditional on being still working at the time of the regime shift, the participants work significantly longer when rewards in the first 20 periods were intermittent as opposed to continuous (supporting H4). However, there is no significant difference between the coefficients of the Fixed and the Random treatments (in contrast to H5) ($p=0.361$). Less risk

¹⁶ Additional regressions with interaction terms between skill and each treatment (available upon request) show that the effect of skill on the exiting decision is not affected by the treatment.

averse participants exit later than the other participants, but, as predicted, skill has no more influence on the decision to exit the game after the regime shift. The estimation of model (3) on the full sample of participants leads to the same conclusions.

To complete this analysis, Table 3 reports the estimates of the determinants of the probability of working until the limit of the 35th period. Model (1) reports the estimates of an unconditional Probit model together with marginal effects, while model (2) corresponds to the second equation of a two-step Probit model with sample selection, in which the selection equation is the probability of working at least 20 periods. The independent variables are the same as in Table 2 and to identify the model with selection, we dropped from the second equation the variable indicating the number of periods paid before period 15 minus 5. We find that the probability of working until the end increases significantly in the Random II treatment while the impact of the Fixed II treatment is only marginally significant in the unconditional Probit model. In model (1), the marginal effects of these treatments are 19.26% and 12.66%, respectively. Both models indicate that the main force driving the latest possible exiting decision is the randomness of rewards rather than their frequency.¹⁷

(Insert Table 3 about here)

We summarize our main findings on exiting decisions as follows.

Result 1. Under ambiguity and regardless of the incentive structure, the less able and the more risk averse participants tend to quit as soon as possible while the more able and less risk averse stay longer. After the regime shift, skill levels do not make any difference but risk attitudes still play a role.

¹⁷ Note that the estimates of model (1) on the sub-sample of participants still at work at the end of period 20 give qualitatively the same findings (available upon request).

Result 2. Under ambiguity, a continuous incentive scheme allows individuals to adjust their behavior quickly to the new situation; they quit very early after the regime shift.

Result 3. Under ambiguity, intermittent incentives lead individuals to quit on average later after a regime shift than under a continuous incentive scheme.

Result 4: Compared to fixed intermittent incentives, random incentives lead to more extreme behavior in both directions, that is, a higher probability to exit the situation either before the regime shift or only at the very end of the game.

To investigate the reasons to continue or exit, Table A in Appendix C classifies the most important reasons cited by players for exiting the task before the 35th period broken down by treatment.¹⁸ It shows that three-fourths of all answers cite financial reasons (monetary tradeoffs and income targeting). Aspects related to the task (“don’t like task”, “boredom”, “too hard”) are only cited first some 20% of the time across all treatments. Table A also reports for each treatment the estimates of an OLS model in which the dependent variable is the number of the exiting period (excluding forced exits after period 35). The independent variables include the reason cited by the participants when exiting voluntarily, either in first, second or third position. The regressions show that task-related concerns are associated with earlier exit while monetary reasons (“no more payment” and opportunity costs) are linked to later exit. In the Random II treatment, the insufficient level of incentives and the difficulty to reach the income target are also linked with earlier exit, but not significantly so ($p=0.11$).

¹⁸ Table A only includes the players who exit voluntarily; those who worked until period 35 received a different questionnaire. The 17 participants in Random II who stayed until the end cited 26 “first” reasons for continuing. Ten (59%) reported that they hoped “to earn more” and another three (18%) said that they had not reached their “target income for the day.” “Like task” was cited first by three players. The remaining responses included the willingness not to stop before others and maximizing the numbers of correct answers.

4.2. Performance profiles

Figure 4 displays the evolution of the number of correct answers per period over time in each treatment. Table 4 summarizes the means and standard deviations of three sets of performance measures and the points participants earned for their performance. The first set of measures indicates the number of correct answers by blocks of periods. The second set shows the number of correct answers for periods earlier than period 21 conditional on whether the previous period (or the previous two periods) gave rise to payment or not; in periods 22 to 35, the previous period is always unpaid. This set of measures captures whether participants adjust their effort to any regular pattern in payments. For example, in the Fixed II treatment do they increase their effort in the period that follows two unpaid periods?¹⁹ The third measures indicate the mean total number of correct answers provided in each treatment.

(Insert Figure 4 and Table 4 about here)

Both Figure 4 and measures 1 in Table 4 indicate hardly any difference in performance per period between treatments in the first 20 periods. This is confirmed by K-S tests (Baseline vs. Random II treatment, $p=0.492$, Baseline vs. Fixed II treatment, $p=0.532$, Random vs. Fixed II treatments, $p=0.611$). The standard deviations are also very close across treatments. In contrast, once payment stops definitively the mean performance in the Baseline decreases rapidly whereas effort remains high in both intermittent incentives treatments. It should be acknowledged, however, that after period 23 only very few participants are still working. K-S tests show that the distribution of the number of correct answers after period 20 is significantly different between the Baseline and both the Random II ($p<0.001$) and the Fixed II treatment ($p=0.001$), but not between

¹⁹ Fig.4 and measures of scores have been computed with the data from the last six sessions (two per treatment), with $N=35$ in the Baseline, $N=35$ in Random II, and $N=36$ in Fixed II. The reason is that, by mistake, the number of correct answers was recorded only in the paid periods in the first six sessions (three with the Baseline and three with the Random II treatments). Therefore, we do not use these sessions for the calculation of performance. We note, however, that the distribution of exiting periods by treatment is the same in the first six and the last six sessions (K-S test, $p=0.398$ for the Baseline and $p=0.937$ for Random II).

the two II treatments ($p=0.995$).²⁰ This suggests that the few participants who do not exit early in the Baseline have little motivation but stay in case payment might restart. In the II treatments, performance increases slightly throughout the game, an important result from an organizational design perspective. Overall, the sum of total correct answers when exiting the game is 66.91 in the Baseline, 80.83 in the Random II treatment and 84.28 in the Fixed II treatment. If performance in the two intermittent treatments is not significantly different (M-W test, two-sided, $p=0.709$), both the Random and the Fixed II treatments differ from the Baseline ($p=0.069$ and $p=0.002$, respectively).

Figure 4 also indicates that performance fluctuates across periods in each treatment but no regular pattern emerges. Measures 2 in Table 4 suggest that in the Fixed II treatment participants do not adjust their effort when in the first 20 periods a period is always paid after two unpaid periods: performance is not higher after two periods unpaid than after one unpaid period (Wilcoxon test, $p=0.642$; for the Random II treatment, $p=0.646$). Similarly, performance does not decline after one period has been paid (Wilcoxon test, $p=0.712$). This does not support hypothesis H3. These results do not necessarily imply, however, that people did not recognize the sequence of payment. Several aspects may have played a role: demand effects cannot be excluded although we were cautious not to generate them (Levitt and List, 2007; Zizzo, 2010); students do not come to the laboratory to be inactive; the task was not that long and painful and there were no on-the-job leisure activities.²¹

²⁰ M-W tests on the number of correct answers by blocks of periods reach the same conclusions. For periods 1 to 20, the p -value is 0.819 for the comparison between the Baseline and Random II, 0.617 for the comparison between the Baseline and Fixed II, and 0.774 for the comparison between the two II treatments. For periods 21 to 25, the p -values are 0.084, 0.290 and 0.418, respectively. For periods 26 to 30, they are 0.016, 0.009, and 0.966. Last, for periods 31 to 35, they are 0.031, 0.097, and 0.550.

²¹ A visual inspection of the data shows that only two players reduce their effort after a paid period; however, they leave only in periods 29 and 30. Disentangling the various explanations would have required eliciting the beliefs that the next period would be paid. But incentivizing this elicitation could have biased the game.

To complement this analysis, we have examined the determinants of the total number of correct answers by estimating an OLS model (see Table 5).²² The independent variables are the same as in the regressions reported in Tables 2 and 3, except that we do not need to control for the number of paid periods in the first 15 periods in the Random treatment.²³ The results show that in comparison with the performance in the Baseline a player produces 18.74 units more in the Fixed II treatment (significant at 1%) and 12.23 units more in the Random II treatment (significant at 5%). Naturally they also show that skill influences the total performance (the coefficient of this variable is 21.72, significant at 1%).

(Insert Table 5 about here)

Combined with our descriptive analysis of performance, this leads to the next result:

Result 5. Before the regime shift, the combination of higher but less frequent incentives produces similar effort to the continuous payment of lower incentives; after the shift, people adjust their effort downward only when they were previously compensated continuously.

Result 6: In the Fixed II treatment, individuals do not adjust effort to the previous sequence of payments.

4.3. Earnings

Consider now the net amount of paying points earned. Table 4 indicates the total number of paying points (*i.e.*, net of entry fees) actually earned by the end of the session. Table 4 shows that the participants who are still involved in period 20 accumulated more points in the first 20 periods in the Random II treatment (mean=47.37) than in the Baseline (mean=42.01; M-W test, $p=0.001$), and more than in the Fixed II treatment (mean=41.89; $p=0.028$). There is no difference

²² We also estimated a two-step Heckman model with sample selection in which the selection equation explains the determinants of the probability to work until the 20th period (available upon request). The Inverse of the Mill's Ratio that was included in the second OLS regression was not significant, indicating that the data do not suffer from a selection bias. This is why we only report the simple OLS model in Table 5. Note also that we can use only the data from the last six sessions, as explained in footnote 19.

²³ We also checked that this variable has no significant influence on performance.

between the Baseline and the Fixed II treatments ($p=0.955$). This higher number of points earned in the Random II treatment, could be due to more early exits than in other treatments and to the fact that less able participants are more likely to exit first. On the other hand, the final net number of paying points earned by all the participants is slightly higher in the Baseline (mean=37.31) than in both the Random II (mean=35.65; M-W test, $p=0.113$) and the Fixed II treatments (mean=33.39; $p=0.051$). There is no difference between the two II treatments ($p=0.319$). This pattern of differences is clearly due to the fact that participants in the II treatments entered (and paid the fee) in more unpaid periods than those in the Baseline.

This supports hypothesis H6 and leads to our final result.

Result 7. Mean total payoffs are marginally higher in the Baseline than in the Intermittent Incentive treatments. In the Random II treatment, mean payoffs were the highest before the regime shift and become lower than in the Baseline at the end of the game.

5. CONCLUSION

The reaction of individuals to ambiguity and their perception of regime shifts are still not well understood. In our experiment, we show that less able and more risk averse individuals are less likely to accept working under high ambiguity. Participants with a higher marginal cost of effort have lower payoffs from undertaking the task and may be locally more risk averse than more able participants with a higher income. Higher risk aversion requires a higher risk premium to participate longer. Moreover, we show that both frequency and regularity in incentive schemes influence the speed at which individuals update their expectations and take the appropriate quitting decisions. Our experiment shows that individuals who are used to receiving rewards in every period react quickly to payment stopping by quitting the task. In contrast, individuals who are rewarded less frequently react differently if their payments were or were not regular. Irregular

payment induces both earlier exits and a higher likelihood of persisting in effort until the last period of the game, probably because it takes more time for individuals to interpret a long interval without rewards as a regime shift. The conjunction of the treatment and the skill effects explains that under intermittent incentives, mean performance is higher than with continuous incentives for the same overall wage cost.

Data concerning economic phenomena are often experienced across time under conditions of ambiguity. Consider stock market prices, interest rates, bonuses. Often these are interpreted as signals about the underlying state of an economic system (*e.g.*, the stock market or the economic situation of a firm) and are used to update expectations and reference points. To study this issue, dynamic models are required to explain how people update their reference points and expectations over time under ambiguity. In this paper, we modeled our participants' beliefs about periods being paid by using a simple, weighted-average form of Hogarth and Einhorn's (1992) belief updating model. Whereas more complex models are possible (see, *e.g.*, Appendix A), the model we use provides a simple account for our data. With this model and the experiment, we have shown that after experiencing different incentive schemes, human learning processes lead to different reactions (*e.g.*, quit or persist). Our findings are also important from an organizational perspective as they show that, after a regime shift, effort still increases over time when rewards were previously intermittent, although not when they were continuous. Whereas we need to remain cautious about the external validity of laboratory experiments (Levitt and List, 2007), notably due to the artificiality of our design, these findings illustrate how people deal with ambiguity in rewards and they could also matter for dynamic contracting problems in the presence of high risks (see notably Biais *et al.*, 2010).

A possible extension of this research would be to consider the case in which no regime shift ever occurs, which would allow us to observe during a longer period of time the impact of

intermittent vs. continuous incentives on exiting decisions and performance. In addition, although we were interested here in the exiting decision, it would be interesting to study how the degree of ambiguity affects the manner in which individuals self-select between different types of incentive schemes at the entry of a job. Finally, our analysis has been carried out using positive incentives. A natural extension would be to investigate how the frequency and regularity of negative incentives, such as audits and sanctions, might impact the evolution of behavior over time, and particularly the respect for rules and social norms.

REFERENCES

- Abeler, J., Falk, A., Goette, L., and Huffman, D. (2011). Reference Points and Effort Provision. *American Economic Review*, 101(2), 470-92.
- Balleine, B.W., Daw, N.D., and O'Doherty, J.P. (2009). In Glimcher, P.W., Camerer, C.F., Fehr, E., and Poldrack, R.A. (Eds.), *Neuroeconomics, Decision-Making and the Brain*. Amsterdam: Elsevier, 367-387.
- Bereby-Meyer, Y., and Roth, A. E. (2006). The Speed of Learning in Noisy Games: Partial Reinforcement and the Sustainability of Cooperation. *American Economic Review*, 96(4), 1029-1042.
- Biais, B., Mariotti, T., Rochet, J.C., and Villeneuve, S. (2010). Large Risks, Limited Liability, and Dynamic Moral Hazard, *Econometrica*, 78(1), 73-118.
- Bretcher, E. G., and Hantula, D. A. (2005). Equivocality and escalation: A replication and preliminary examination of frustration. *Journal of Applied Social Psychology*, 35(12), 2606-2619.
- Camerer, C.F., Babcock, L., Loewenstein, G., and Thaler, R. (1997). Labor Supply of New-York City Cabdrivers: One Day at a Time. *The Quarterly Journal of Economics*, 112(2), 407-441.
- Camerer, C.F., and Ho, T. H. (1999). Experience-weighted Attraction Learning in Normal Form Games. *Econometrica*, 67(4), 827-874.
- Caplin, A., and Dean, M. (2008). Dopamine, Reward Prediction Error, and Economics. *The Quarterly Journal of Economics*, 123(2), 663-701.
- Crawford, V.P., and Meng, J. (2011). New York City Cab Drivers' Labor Supply Revisited: Reference-Dependent Preferences with Rational-Expectations Targets for Hours and Income. *American Economic Review*, 101, 1912-1932.
- DeNicolis Bragger, J., Bragger, D., Hantula, D. A., and Kirnan, J. (1998). Hysteresis and uncertainty: The effect of uncertainty on delays to exit decisions. *Organizational Behavioral and Human Decision Processes*, 74 (3), 229-253.
- DeNicolis Bragger, J., Hantula, D. A., Bragger, D., Kirnan, J., and Kutcher, E. (2003). When success breeds failure: History, hysteresis, and delayed exit decisions. *Journal of Applied Psychology*, 88 (1), 6-14.
- Deslauriers, B. C., and Everett, P. B. (1977). Effects of intermittent and continuous token reinforcement on bus ridership. *Journal of Applied Psychology*, 62 (4), 369-375.
- Ederer, F., Holden R., and Meyer, M. (2013). Gaming and Strategic Ambiguity in Incentive Provision. Mimeo. Available at: <http://ssrn.com/abstract=2197597>.
- Erev, I., Ingram, P., Raz, O., and Shany, D. (2010). Continuous Punishment and the Potential of Gentle Rule Enforcement. *Behavioural Processes*, 84, 366-371.
- Eriksson, T., Poulsen, A., and Villeval, M.C. (2009). Feedback and Incentives: Experimental Evidence. *Labour Economics*, 16 (6), 679-688.
- Ferster, C. S., and Skinner, B. F. (1957). *Schedules of Reinforcement*, New York, NY: Appleton-Century-Crofts.

- Fox, C.R., and Tversky, A. (1995). Ambiguity Aversion and Comparative Ignorance. *The Quarterly Journal of Economics*, 110(3), 585-603.
- Goette, L., Huffman, D., and Fehr, E. (2004). Loss Aversion and Labor Supply. *Journal of the European Association*, 2(2-3), 216-228.
- Goette, L., and Huffman, D. (2006). Incentives and the Allocation of Effort over Time: The Joint Role of Affective and Cognitive Decision Making. *IZA Discussion Paper 2400*, Bonn.
- Golz, S. M. (1992). A sequential learning analysis of decisions in organization to escalate investments despite continuing costs or losses. *Journal of Applied Behavior Analysis*, 25, 561-574.
- Greiner, B. (2004). An online recruitment system for economic experiments. In *Forschung und wissenschaftliches Rechnen GWDG Bericht 63*, Ed. K. Kremer, and V. Macho. Göttingen: Gesellschaft für Wissenschaftliche Datenverarbeitung.
- Hantula, D. A., and Crowell, C. R. (1994). Intermittent reinforcement and escalation processes in sequential decision making: A replication and theoretical analysis. *Journal of Organizational Behavior Management*, 14 (2), 7-36.
- Hilgard, E. R., and Bower, G. H. (1975). *Theories of Learning*, 4th ed, Englewood Cliffs, NJ: Prentice-Hall.
- Hogarth, R. M., and Einhorn, H. J. (1992). Order Effects in Belief Updating: The Belief-Adjustment Model. *Cognitive Psychology*, 24, 1-55.
- Jenkins, W.O., and Stanley, J.C. (1950). Partial Reinforcement. A Review and a Critique. *Psychological Bulletin*, 47(3), 197-234.
- Kellner, C., Riener, G. (2012). The effect of ambiguity aversion on reward scheme choice. *DICE Discussion Paper 55*, University of Düsseldorf.
- Koszegi, B., and Rabin, M. (2006). A Model of Reference-Dependent Preferences. *Quarterly Journal of Economics*, 121 (4): 1133-1165.
- Koszegi, B., and Rabin, M. (2007). Reference-Dependent Risk Attitudes. *American Economic Review*, 97(4), 1047-1073.
- Lazear, E. P. (1990). The timing of raises and other payments. *Carnegie-Rochester Conference Series on Public Policy*, 33, 13-48.
- Lazear, E. P. (2006). Speeding, Terrorism, and Teaching to the Test. *The Quarterly Journal of Economics*, 121 (3), 1029-1061.
- Levitt, S.D., and List, J. A. (2007). What Do Laboratory Experiments Measuring Social Preferences Reveal About the Real World? *Journal of Economic Perspectives*, 21(2), 153-174.
- March, J. G. (1996). Learning to be risk averse. *Psychological Review*, 103(2), 309-319
- Montague, P.R., Dayan, P., and Sejnowski, T.J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *Journal of Neurosciences*, 16, 1936-1947.

- Niv, Y., and Montague, P.R. (2009). Theoretical and Empirical Studies of Learning. In Glimcher, P.W., Camerer, C.F., Fehr, E., & Poldrack, R.A. (Eds.), *Neuroeconomics, Decision-Making and the Brain*. Amsterdam: Elsevier, 367-387.
- O’Flaherty, B., and Komaki, J. L. (1992). Going beyond with Bayesian updating. *Journal of Applied Behavior Analysis*, 25, 585-597.
- Prendergast, C. (1999). The Provision of Incentives in Firms. *Journal of Economic Literature*, 37(1), 7-63.
- Robbins, D. (1971). Partial Reinforcement: A Selective Review of the Alleyway Literature since 1960. *Psychological Bulletin*, 76(6), 415-431.
- Schultz, W., Dayan, P., and P.R. Montague (1997). A neural substrate of prediction and reward. *Science*, 275, 1593-1599.
- Trautmann S, and Kuilen, G. van de (2011). Belief Elicitation: A Horse Race among Truth Serums, Tilburg University, Center for Economic Research Discussion Paper 2011-117.
- Tustin, R.D., Morgan, P. (1985). Choice of reinforcement rates and work rates with concurrent schedules. *Journal of Economic Psychology*, 6, 109-141.
- Zizzo D.J. (2010). Experimenter demand effects in economic experiments. *Experimental Economics*, 13, 75–98.

TABLES AND FIGURES

Table 1. Characteristics of the experimental sessions

Session number	Number of participants	Percentage of males	Treatment
1	12	42	Random Intermittent Incentives
2	19	47	Random Intermittent Incentives
3	17	29	Baseline
4	19	42	Random Intermittent Incentives
5	18	78	Baseline
6	19	32	Baseline
7	17	59	Random Intermittent Incentives
8	18	39	Random Intermittent Incentives
9	17	47	Baseline
10	18	72	Baseline
11	18	50	Fixed Intermittent Incentives
12	18	22	Fixed Intermittent Incentives

Table 2. Determinants of the exiting period

	Probability of still working in period 20 Probit (1)	Exiting period conditional on working in period 20 Tobit (2)	Exiting period Tobit (3)
<i>Baseline treatment</i>	Ref.	Ref.	Ref.
Fixed II treatment	.03 (.32) [.01]	6.16*** (1.06)	4.21*** (1.35)
Random II Treatment	-0.78*** (.24) [-.25**]	7.26*** (.94)	2.31** (1.07)
Random II treatment * {(Nb. of paid periods until period 15) - 5}	-.46** (.18) [-.14**]	.38 (.73)	-1.75** (0.84)
Afternoon sessions	-.46** (.22) [-.14**]	-.21 (.80)	-1.55# (.95)
Rank of arrival	-.01 (.02) [-.00]	-.11# (.07)	-.13 (.09)
Skill	.54*** (.20) [.16***]	.52 (.79)	2.41*** (.87)
Risk index	.09** (.04) [.03**]	.30* (.17)	.42** (.19)
Ambiguity aversion	.09** (.04) [.03**]	.06 (.17)	.26 (.18)
Multiple switches	-.36 (.32) [-.12]	-.86 (1.39)	-1.66 (1.60)
Relative wealth	-.09* (.05) [-.03*]	.13 (.20)	-.25 (.22)
Male	.31 (.22) [.09]	-1.17 (.79)	.39 (.95)
Age	-.02 (.05) [-.01]	.61 (.22)	.33 (.25)
Level at High School final exam	-.03 (.13) [-.01]	.04 (.45)	-.05 (.55)
Number of years of study	.12 (.10) [.04]	-.56 (.41)	.31 (.49)
Constant	-1.05 (1.44)	9.87 (6.34)	6.14 (6.69)
N	210	152	210
Left/right censored observ.	-	11/22	12/22
LR χ^2	43.34	72.74	39.01
$p > \chi^2$	<0.001	<0.001	<0.001
Log-likelihood	-102.08	-380.22	-616.66
Pseudo R ²	-	.09	.03

Note: Standard errors are in parentheses and marginal effects are in square brackets. *** means significant at the 0.01 level, ** at the 0.05 level, * at the 0.10 level, and # at the .11 level. In model (2), data are censored at period 20.

In model (3) data are left censored at period 16 (first period in which individuals are allowed to exit). In both Tobit models are right censored at period 35.

Table 3. Determinants of the probability of working until the 35th period (Probit model with sample selection)

Dependent variable: probability to reach the 35 th period	Unconditional Probit model (1)	Probit model with sample selection (2)
<i>Baseline treatment</i>	<i>Ref.</i>	<i>Ref.</i>
Fixed II treatment	.85* (.48) [.13]	.74 (.55)
Random II Treatment	1.29*** (.42) [.19]	1.66*** (.54)
Random II treatment *{(Nb. of paid periods until period 15) -5}	-.27 (.20) [-.04]	-
Afternoon sessions	-.32 (.28) [-.05]	-.06 (.41)
Rank of arrival	.00 (.03) [.00]	-.00 (.03)
Skill	.16 (.26) [.02]	-.11 (.31)
Risk index	.04 (.06) [.01]	-.01 (.08)
Ambiguity aversion	.05 (.05) [.01]	.01 (.08)
Multiple switches	-.36 (.48) [-.05]	-.20 (.48)
Relative wealth	-.00 (.07) [-.00]	.10 (.08)
Male	-.01 (.28) [-.00]	-.18 (.27)
Age	.13** (.06) [.02]	.14* (.07)
Level at High School final exam	.16 (.17) [.02]	.15 (.17)
Number of years of study	-.12 (.13) [-.02]	-.22* (.13)
Constant	-5.87*** (2.13)	-4.00 (3.21)
N	210	210
Left-censored observations	-	58
Wald χ^2	-	25.85
LR χ^2	26.41	-
$p > \chi^2$.02	0.02
Log-likelihood	-57.23	-149.61
Pseudo R ²	.19	-

Note: Standard errors are in parentheses and marginal effects are in square brackets. *** means significant at the .01 level, ** at the .05 level, and * at the .10 level.

Table 4. Summary statistics of task performance (in points)

Treatments	Baseline	Random II	Fixed II
<i>1) Mean performance per period by blocks of periods</i>			
Periods 1-15 (compulsory)	3.00 (1.00)	3.13 (0.88)	3.06 (0.94)
Periods 16-20	3.16 (1.05)	3.33 (0.82)	3.26 (0.83)
Periods 21-25	3.01 (1.03)	3.51 (0.61)	3.38 (0.85)
Periods 26-30	2.55 (1.00)	3.43 (0.72)	3.43 (0.76)
Periods 31-35	2.17 (0.98)	3.42 (0.69)	3.38 (0.87)
<i>2) Mean performance per period</i>			
<i>Periods <21</i>			
Previous period paid	3.04 (1.02)	3.27 (0.79)	3.10 (0.92)
Previous period unpaid	-	3.14 (0.90)	3.14 (0.91)
Two previous periods unpaid	-	3.17 (0.91)	3.12 (0.91)
<i>Periods >21 (previous period unpaid)</i>	2.78 (1.06)	3.45 (0.67)	3.42 (0.79)
<i>3) Total performance across session</i>	66.91 (18.25)	80.83 (29.56)	84.28 (28.88)
<i>4) Mean total paying points</i>			
End of period 20 if working in period 20	42.01 (9.81)	47.37 (8.33)	41.89 (11.49)
End of sessions	37.31 (11.47)	35.65 (10.78)	33.39 (12.57)

Note: Standard deviations are in parentheses. For measuring performance when the previous period was paid or unpaid before period 21, we consider periods 2 to 20, included. For measuring performance when the two previous periods were unpaid before period 21, we consider periods 3 to 20, included.

Table 5. Determinants of the total number of correct answers

Dependent variable: Total number of correct answers	OLS model (1)
<i>Baseline treatment</i>	<i>Ref.</i>
Fixed II treatment	18.74*** (5.66)
Random II Treatment	12.23** (5.19)
Afternoon sessions	-11.06** (4.40)
Rank of arrival	-.66 (.43)
Skill	21.72*** (3.45)
Risk index	.97 (.85)
Ambiguity aversion	1.04 (.77)
Multiple switches	-16.31* (.48)
Relative wealth	-.03 (1.02)
Male	1.85 (4.50)
Age	.07 (.89)
Level at High School final exam	-1.65 (2.60)
Number of years of study	2.28 (1.94)
Constant	-1.94 (27.24)
N	106
F	6.15
Prob>F	.00
Adjusted R ²	.39

Note: Standard errors are in parentheses. *** means significant at the .01 level, ** at the .05 level, and * at the .10 level.

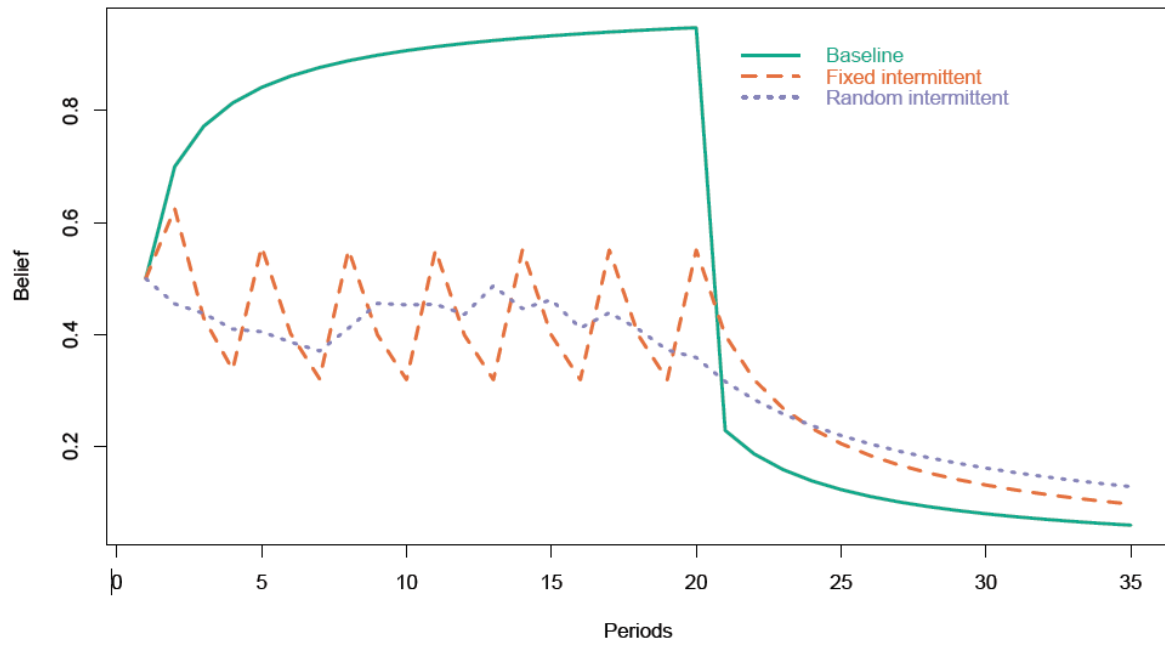


Fig. 1. Simulation of beliefs (with $\alpha_{Baseline} = \beta_{Baseline} = 0.8$; $\alpha_{FixedII} = \beta_{FixedII} = 0.5$; and $\alpha_{RandomII} = \beta_{RandomII} = 0.3$)

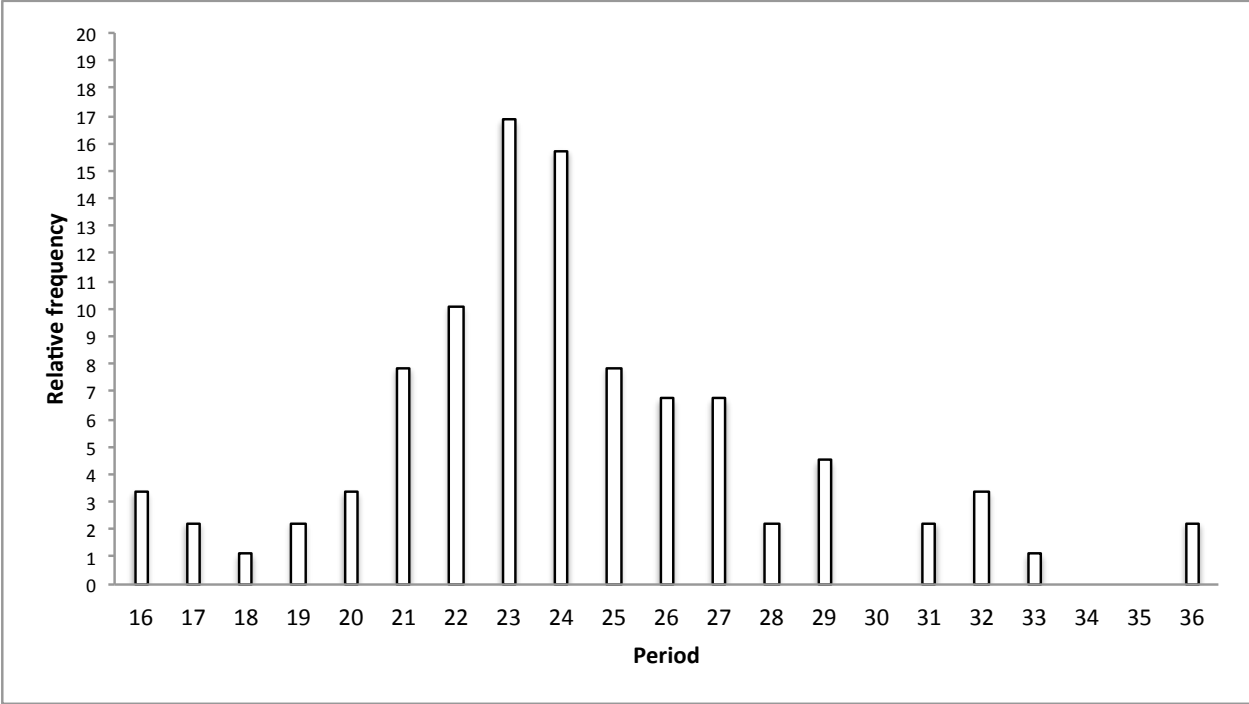


Fig.2a. Relative frequency of exits, by period in the baseline treatment

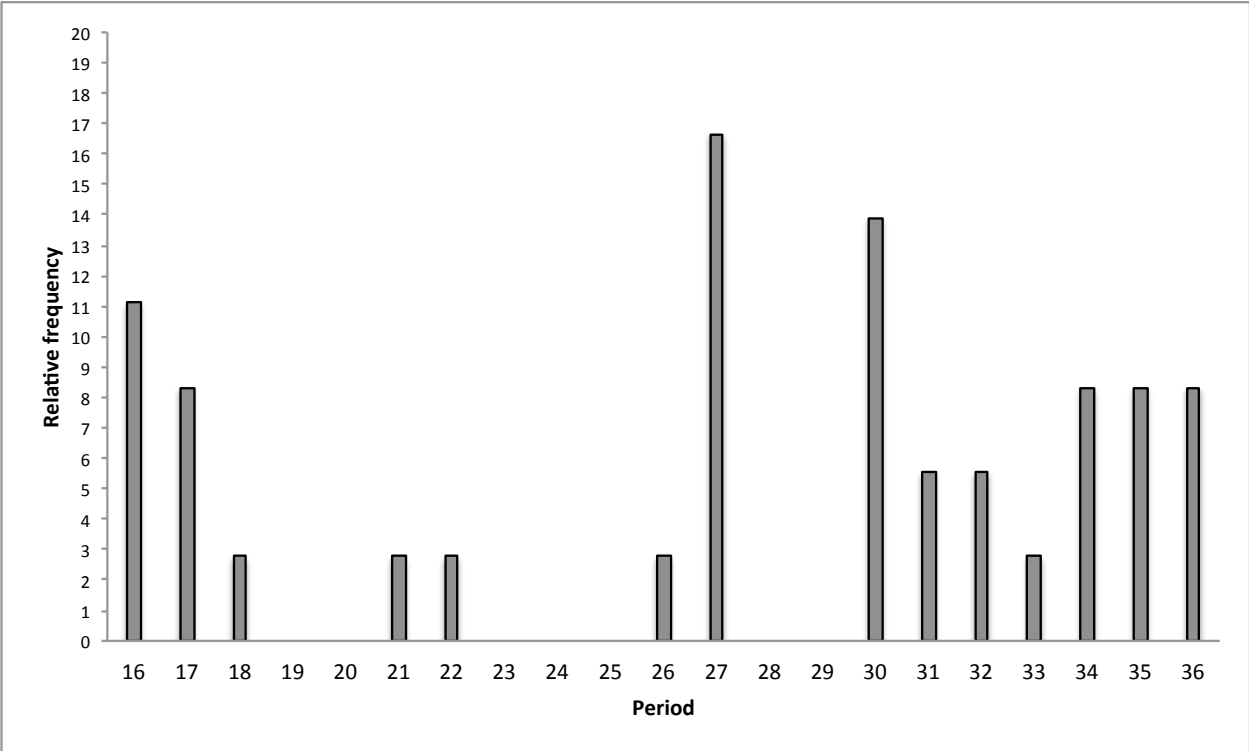


Fig.2b. Relative frequency of exits, by period in the Fixed II treatment

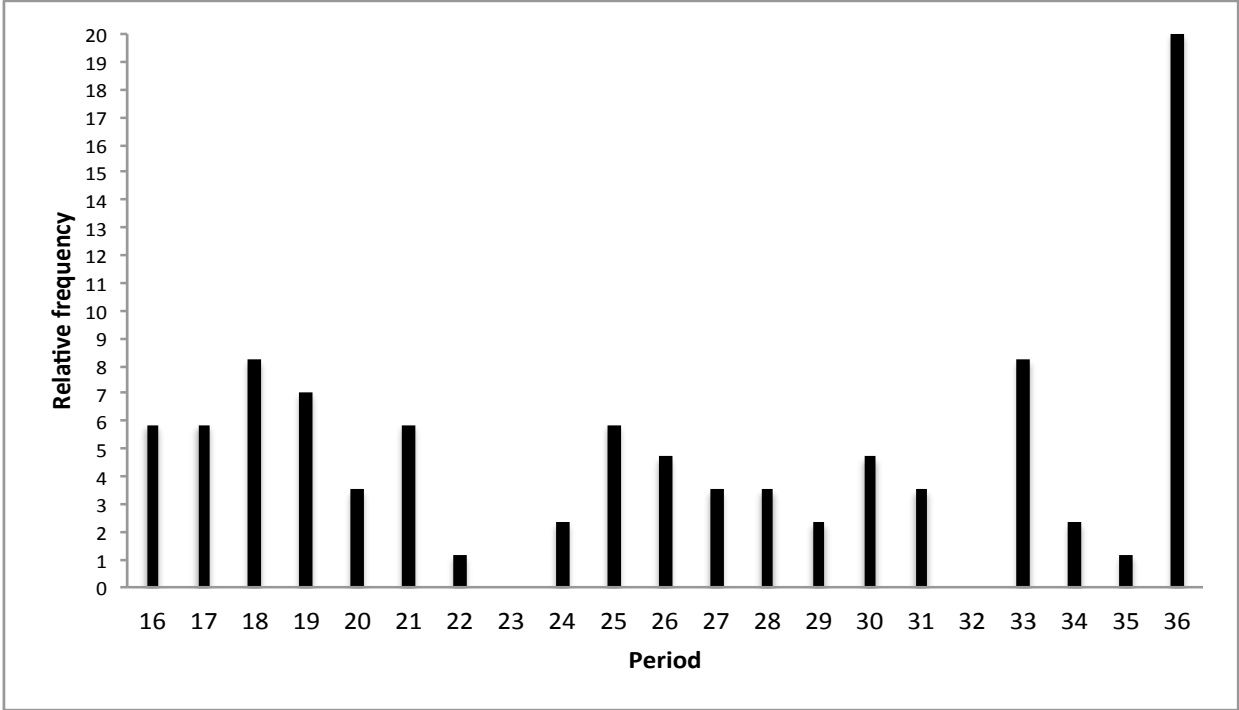


Fig.2c. Relative frequency of exits, by period in the Random II treatment

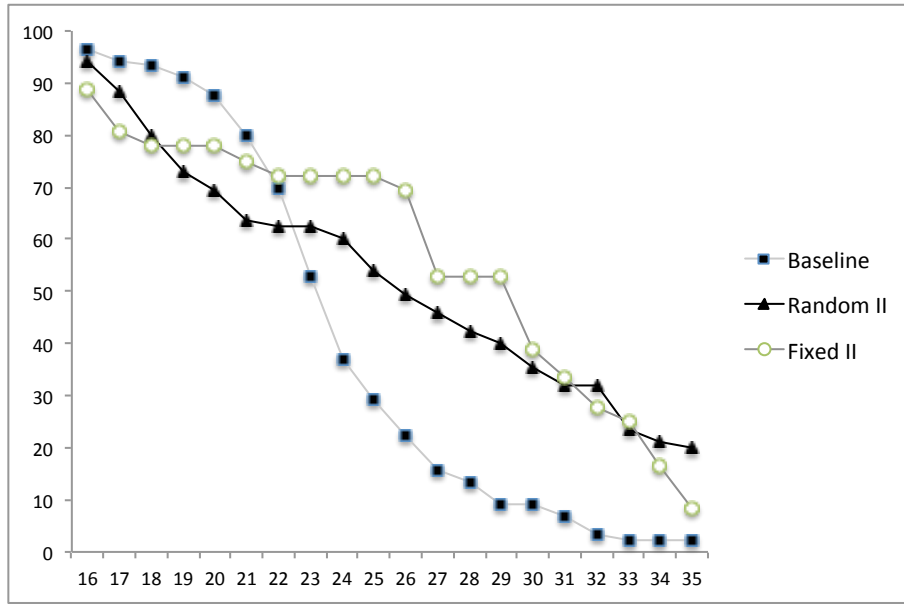


Fig.3. Percentage of participants remaining at the beginning of each period following the compulsory working time

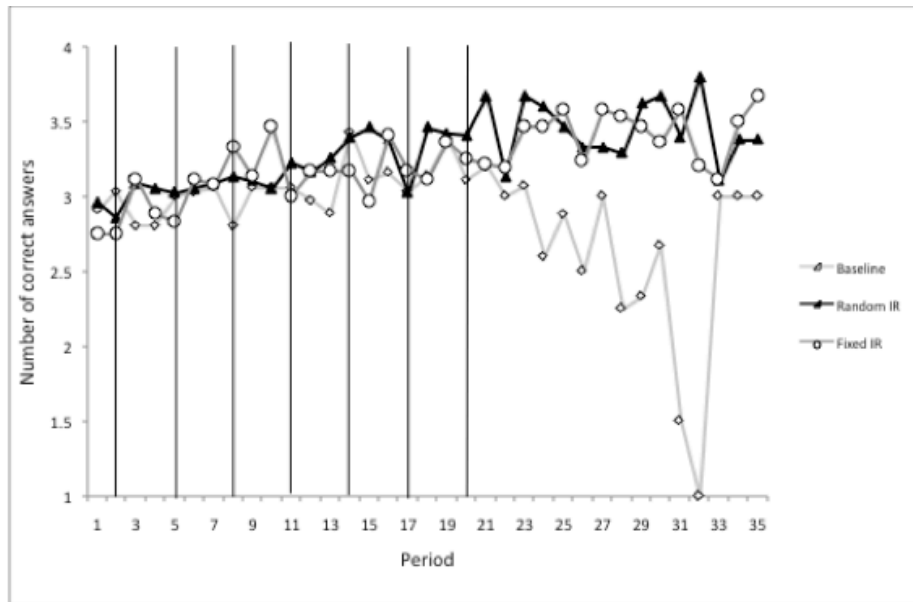


Fig.4. Evolution of the number of correct answers over time by treatment

Note: The vertical lines indicate the periods that were paid in the Fixed Intermittent Incentive treatment.