Working Paper Series

#2016-028

Mimetic behaviour and institutional persistence: A two-armed bandit experiment

Stefania Innocenti and Robin Cowan

Maastricht Economic and social Research institute on Innovation and Technology (UNU-MERIT)
email: info@merit.unu.edu | website: http://www.merit.unu.edu

Maastricht Graduate School of Governance (MGSoG)
email: info-governance@maastrichtuniversity.nl | website: http://www.maastrichtuniversity.nl/governance

Boschstraat 24, 6211 AX Maastricht, The Netherlands
Tel: (31) (43) 388 44 00

UNU-MERIT Working Papers
ISSN 1871-9872

Maastricht Economic and social Research Institute on Innovation and Technology
UNU-MERIT

Maastricht Graduate School of Governance
MGSoG

UNU-MERIT Working Papers intend to disseminate preliminary results of research carried out at UNU-MERIT and MGSoG to stimulate discussion on the issues raised.

Mimetic behaviour and institutional persistence: a two-armed bandit experiment*

Stefania Innocenti†

UNU-MERIT, Maastricht University

Robin Cowan‡

UNU-MERIT, Maastricht University
BETA, Strasbourg University
IUF, Institut Universitaire de France

Abstract

Institutions are the result of many individual decisions. Understanding how agents filter available information concerning the behaviour of others is therefore crucial. In this paper we investigate whether and how agents' self-efficacy beliefs affect mimetic behaviour and thus, implicitly the evolution of institutions. We propose an experimental task, which is a modified version of the two-armed bandit with finite time horizon. In the first treatment, we study in detail individual learning. In the second treatment, we measure how individuals use the information they gather while observing a randomly selected group leader. We find a negative relation between self-efficacy beliefs and the propensity to emulate a peer. This might ultimately affect the likelihood of institutional change.

Keywords: Emulation; mimicry; laboratory experiment; self-efficacy; institutional change
JEL: B52; C13; C91; D02; D03; D83

*We would like to thank Eleonora Nillesen, Martin Strobel, Arno Riedl, Matthew Embrey and Nico Rasters for their precious advice. We also acknowledge the helpful comments of all the participants to the BEElab proposal meeting in Maastricht in February 2015, the ASFEE 2015 conference in Paris, the MLSE seminar in Maastricht in December 2015, the seminar at the University of Florence and IMEBESS 2016 conference in April 2016. This research was supported by the Institut Universitaire de France.

†E-mail address: `innocenti@merit.unu.edu`; Corresponding author.

‡E-mail address: `cowan@merit.unu.edu`

1 Introduction

In social sciences, it is by now commonly accepted that "institutions" refer to "the rules of the game in society or, more formally, [to] the humanly devised constraints that shape human interactions" (North, 1990, p.3). An institution is a rule or a set of rules which frames how people, a group or a population interact with one another. Technologies, language codes, artistic spirits, fashions are institutions, and their adoption is the result of multiple individual decision making processes (McElreath et al., 2005). Nevertheless, one question remains highly debated across literatures: Why do people agree on some institutional structures, abide by them and perpetuate the rules of behaviour they prescribe?

One might argue that humans calculate the costs and benefits of certain behaviours and, when the advantages are evidently higher than the costs, they choose a specific institutional structure and follow the rules that it imposes. Social psychologists do not see rule compliance as a mere result of a cost-benefit analysis. They believe that humans also follow rules because of some basic cognitive needs such as accuracy, affiliation and social identity. The human mind is a sense-making machine (Kahneman, 2011), its goal is to gain cognitive clarity. Humans strive for accuracy, explanation and meaning. Contrarily, inconsistencies generate discomfort. For these reasons, humans tend to adapt to the patterns of behaviour they observe. Moreover, humans follow rules because of their deep desire for affiliation and social identity. Being rejected or disapproved induces fear; humans need to belong (Baumeister and Leary, 1995). According to the social identity perspective, humans behave so as to maintain their identities as self-consistent and socially integrated (Cialdini and Goldstein, 2004).

Many institutions can be seen as emergent, in that they arise endogenously as agents modify their behaviour, adapting to the behaviour of others. Thus, a key building block of institutional change resides precisely in these behavioural changes, attempting to satisfy these different, and sometimes conflicting cognitive needs and desires. Responding to, and perhaps mimicking others can be seen as a form of social learning. But at the same time, learning processes and behavioural adjustments, involving non-mimetic behaviour, ignite gradual, incremental institutional change (Mahoney and Thelen, 2009). Because institutional change at its base involves some "coordinated" change in the behaviour of many individuals, observing others' behaviour and possibly learning from it, are central in the micro-processes that lie behind any institutional change. In fact, institutional transformation depends on people's interpretation of existing rules and their effort to alter them (Mahoney and Thelen, 2009).

Thus, in order to understand how rules of behaviour might change, studying how agents filter available information concerning the behaviour of others becomes of crucial importance. This paper, building on contributions from psychology and economics, sheds some light on the reasons that individuals follow existing institutional structures and perpetuate them over time, instead of deviating from them. We investigate how cognition shapes the likelihood of mimetic behaviour, and thus, implicitly, whether individuals might or might not bring about new behaviours, contributing to institutional evolution or, on the contrary, to its persistence.

Imagine a group of small farmers has two distinct types of seed. If individuals act alone and have no reference point, they discover the best seed solely on the basis of their individual successes and mistakes. However, if they are able to observe others, they could use these observations in some way to add to their own information regarding which seed performs better. But to what extent will they mimic the behaviour of a neighbour, following by the behavioural path he/she set, or deviate from it? And what role do individual introspective beliefs, particularly about their own self-efficacy, play in this?

These are the questions this paper wishes to answer. More specifically, we are interested in understanding whether self-efficacy beliefs affect individual propensity to emulate a single target subject. Our innovative experimental design seems able to uncover another possible micro-mechanism to identify the reasons for which individuals abide by a pre-established institutional structure and have no intention to change it. We use a specific experimental task which is a modified version of the common two-armed bandit game with finite time horizon. Agents, in a laboratory, make a series of consecutive choices from which they derive real payoffs for two consecutive treatments. We consider individual learning as instrumental to observational learning. We assume that the two phenomena are complementary. Therefore, we study individual learning patterns in detail using data from treatment one. In treatment two we observe how individuals learn when observing the actions and rewards obtained by a randomly selected group "leader" who plays before everybody else. The leader can be considered a "focal point" given his pioneer nature. Our goal is to understand how people process the information retrieved by observing others' experience. Our hypothesis is that weakly self-efficacious people use the leader as a role model and do not deviate from his action, implicitly abiding by the behavioural path he sets. Conversely, more self-efficacious people might be less apt to follow the leader's action and thus more inclined to experiment new patterns of behaviour thus sparking institutional change. We find significant evidence to support this hypothesis: in easy environmental conditions, higher self-efficacy beliefs reduce the propensity to emulate and induce higher reliance on individual learning processes. Results are inconclusive for more volatile environments.

The remainder of the paper is organised as follows. Section 2 briefly reviews the institutional economics literature, connects it to the literature on learning and to the social psychology perspective. Section 3 presents the general design of our experiment. Section 4 elaborates the behavioural hypotheses. Section 5 presents our estimation strategy. Firstly we treat individual learning in great detail and selects, within a set of possible candidate models, the one which seems to best fit the individual learning process used by our subjects. Taking stock of these estimates, we then estimate the emulation parameters. We test our main hypothesis in section 6. Lastly, section 7 concludes and stresses the limitations of this study.

2 Related literature

Research on institutions in different disciplines has channeled renewed attention towards their functioning and evolution.

Within the rational-choice framework, two main approaches exist (Kingston and Caballero, 2009): institutions-as-rules; and institutions-as-equilibria. The two approaches mainly differ regarding the mechanisms behind the institutional genesis and evolution. The institutions-as-rules approach takes a functionalist perspective. Institutions are considered the rules of the game in a society (North, 1990) which are generally enforced by members of relevant groups. Their emergence and change can be the fruit of a centralised process as well as of an evolutionary one. This approach thus studies the formation of institutions but abstains from providing an explanation for the reasons for which people follow a certain pattern of behaviour.

Conversely, the importance of motivations stands at the core of the institution-as-equilibria approach. This perspective focuses on how interacting agents create a structure that gives them the motivation to act in a manner which conforms to the structure itself, and thus perpetuates it. An institution which does not provide incentives to follow it carries the seeds of its own destruction. A rule in fact serves as a coordination device which shapes the behaviour we can expect others to have. The beliefs people hold, which embed the institutions they are part of, ultimately motivate others' behaviours. In this perspective, rules represent a social construct which induces behavioural coordination and creates social order. Thus, institutions get formed and changed endogenously (Greif and Kingston, 2011; Greif and Laitin, 2004; Kosfeld et al., 2009).

These two approaches treat change differently but a subtle commonality exists.

According to the institutions-as-rules perspective, institutions are responsive to the interests and needs of their creators (Greif and Kingston, 2011). The process of change depends on the intentions of the players to enact institutional change and on their understanding of it (North, 2005; Mantzavinos et al., 2004). Agents have mental models which reflect their understanding of reality. These models are used to evaluate the possibility and desirability of rule changes. Overtime, mental models are revised. In this way agents might alter their perceptions about possible new rules, leading them to try to change the old ones. The revision of mental models occurs as a consequence of a learning process (Denzau and North, 1994; Mantzavinos et al., 2004).

The institutions-as-equilibria perspective underlines another causal link between beliefs and change. As said beliefs motivate behaviour and shape expectations about the behaviour of others. Beliefs are relevant because individuals might have limited information concerning the surrounding environment or the strategies of the other players. Rules coordinate people's beliefs. Thus, rules aggregate knowledge and information. Consequently, rule change is affected by the way in which agents process available pieces of information.

The two approaches thus affirm that the key to institutional change is the evolution of mental models or the evolution of motivations, that is to say, learning.

Learning implies acquiring new information until the modification or reinforcement of existing knowledge, or beliefs occurs. When an individual learns something, he acquires some new signals concerning the state of nature. This new information might either corroborate previous understanding, strengthening the conviction of a behaviour, or it might contradict, it inducing the agent to discard old knowledge, and possibly re-consider some of his behaviour. As most other animals, humans change their behavioural repertoire through personal experience or observing others.

At the beginning of the 1990s, learning became a subject of study in economics and political science. Banerjee (1992) and Bikhchandani et al. (1992) — who proposed models which have since become known as herding and information cascades models respectively — first pointed out the importance of mimetic behaviours. In both models subjects take sequential actions and have to discover, given a set of possible actions, which is the best strategy. The main difference between the two models resides in the amount of information the subjects have. In the Banerjee (1992) model some subjects hold more information than others, while in Bikhchandani et al. (1992) all subjects are equally informed. These models made clear that subjects might ignore their private information and rather mimic the behaviour of their predecessor(s).

These models are closely related to observational learning models elaborated in anthropology (see Henrich and McElreath (2003) for a review). Social learning abilities are adaptive and, in some circumstances, learning from others can reduce individual understanding of the environment and thus lead to the diffusion of maladaptive behaviours. These models nevertheless stress that observational learning in the form of mimicry, and individual learning coexist.

This same idea has been studied in social psychology by Bandura. In his view the missing piece to decode human behaviour resides into people's introspective beliefs on their capabilities. The acquisition and retention of new patterns of behaviour, i.e. learning, is the result of both experience and observation of others.

New behaviours are not learned automatically, but they are rather cognitively mediated by self-corrective adjustments which depend on the feedback individuals receive from their own performance or that of their peers. Human behaviour is thus the result of the interplay between external sources of influence and (crucially) the beliefs people hold on their own capabilities to produce desired effects by their own actions. The latter is referred to as an agent's belief in his or her own self-efficacy.[1] Behavioural changes thus depend on the self-efficacy beliefs that individuals develop continually integrating information from 4 different sources. Firstly, self-efficacy beliefs change with the result of personal attempts to control contingencies. Secondly, self-efficacy beliefs are influenced also by the observations of the behaviour of others and the consequences of those behaviours. Vicarious experiences generally have weaker effects on self-efficacy beliefs than do personal experiences (Bandura, 1997). Thirdly, efficacy beliefs are influenced by verbal persuasion, though this tends to be less potent and enduring than personal and vicarious experiences. Lastly, individuals might associate their performance with their emotional and physiological states (Bandura, 1989, 1997)

This paper, drawing on these literatures, investigates a possible micro foundation of mimetic behaviours in an experimental fashion. It assesses a link between self-efficacy beliefs and mimetic behaviours and thus implicitly institutional change. More precisely, following Bandura's theory, we assume self-efficacy mediates individual behaviour. Behavioural outcomes affect learning processes and feed back into future actions. We first study individual learning strategies in detail. We select the best possible model able to predict how our experimental subjects learned over time from the outcomes of their own experimentation, and how they used their information to make decisions. Second, we focus on emulation. We investigate the conditions under which self-efficacy beliefs condition the propensity to emulate a pioneer leader who represents a focal point. We believe that this link

[1]Self-efficacy is concerned with people's perceived capabilities to achieve some goals. Self-efficacy differs from other concepts such as those of self-esteem, locus of control or outcome expectancies. In his 1997 book Bandura claims that whilst efficacy is a judgment of one's capability, self-esteem is a judgment of self-worth. According to Bandura (1997, p. 11) "there is no fixed relationship between beliefs about one's capabilities and whether one likes or dislikes oneself". Even locus of control and self-efficacy are, according to Bandura (1997, p. 20), "entirely different phenomena". Locus of control is concerned with the beliefs that behavioural outcomes depend on one's own actions or on forces beyond personal control and "cannot by any stretch of imagination be considered the same as beliefs about whether one can produce certain actions (self-efficacy)" (Bandura, 1997, p.20). Perceived self-efficacy is also different from outcome expectancies. Self-efficacy has to do with people's confidence that they can perform a certain action if they wish to. Outcome expectations are judgments about the outcomes that are likely to flow from the performances one puts in place. They depict the expectations one has on outcomes given the behavioural choices he has decided to make (See Bandura (1997) on this specific point).

implicitly affects the likelihood of institutional persistence.

3 General design of the experiment

Our experiment consists of three parts.

First, our subjects were presented with a questionnaire to assess their self-efficacy level. This is standard questionnaire used to rate participants' confidence to perform certain behaviours in a variety of circumstances (Schwarzer and Jerusalem, 1995; Schwarzer et al., 1997; Judge and Bono, 2001). It is not task specific, it is rather meant to evaluate how competently one can perform across a variety of hypothetical situations. It deals with people's perceived potential. We treat this information as static and exogenous. We do not consider any feedback-loop between learning and self-efficacy beliefs, this seems reasonable given the short duration of the experiment. Second, subjects answered 10 questions aiming at measuring their problem-solving abilities and math skills. We selected some questions from the standard Graduate Record Examination (GRE) test. Subjects were allowed neither calculators nor pens and paper. Finally, subjects participated in the experimental game which is a modified version of the common two-armed bandit problem with finite time horizon.

The bandit problem takes its name from the common slot machines which can be found in casinos. In order to play, the gambler inserts a coin and pulls one of the machine's available handles (or arms) initiating the spinning of some flywheels. When the flywheels stop, a combination is displayed and the player receives a payoff. The subsequent gamble starts from the last combination obtained, nevertheless the player can choose to pull another of the available arms. At any trial, the gambler compares the scores obtained through the chosen handles. His objective is to maximise, over a series of pulls, the expected payoff.[2]

The reasons for choosing such experimental task are several. First, this setting allows us to assess how people make reiterative choices and to check for their consistency over time. Moreover, small modifications allow us to control how the possibility to mimic a peer and other contextual mutations affect individual behaviour. Lastly, binary choices lie at the core of the two-armed bandit. Although Simon underlined the importance of binary choice experiments to test, for example, utility-maximising principles as early as 1959, to the best of our knowledge only very few experiments in economics have followed this prescription (Banks et al., 1997; McElreath et al., 2005; Gans et al., 2007).

[2] An optimal strategy for infinite time horizon bandit problems and exponential discount rates can be found calculating the Gittins index (Gittins, 1979). However, this strategy is not optimal in our case because our problem has a finite time horizon and the type of discounting is for us unknown.

The game is divided into 2 treatments whose order was not randomised and which differ from one another mainly in terms of the information participants received. In treatment one, subjects sequentially chose between two alternative colours (A and B) for three sets of twenty rounds, thus making 60 binary choices. We will call a set of 20 rounds sub-setting. To avoid people carrying their priors over sub-settings, we changed every 20 rounds the two colours between which subjects had to choose. Each sub-setting had a "preferred" colour, in the sense that it yielded a higher expected payoff. The preferred colour was set randomly at the start of the sub-setting, was the same for all subjects, and unchanged during the sub-setting. Participants were informed that there was a difference in expected payoffs to the two colours, but not what the expected payoffs were, nor which was higher. They were informed that payoffs ranged between 1 and 18 units. Subjects played by selecting one colour in each round. After the colour was selected agents were informed about the score received, and reminded which colour they had chosen. Only the most recent choice and payoff were displayed. Before the first round, no information was displayed. Payoffs were drawn from normal truncated distributions, with fixed mean and variance, bounded between 1 and 18.[3] Thus a sub-setting can be characterised by a quadruple: $(\mu^A, \mu^B, \sigma^{A^2}, \sigma^{B^2})$.

The mean of the more rewarding colour, for instance A, was fixed and equal to $\mu^A = 13$ units, while that of the less rewarding one was $\mu^B = 10$ units. The variance was the same for both distributions and determined the difficulty of the decision-making environment. The three sub-settings were distinguished by their variances in the payoff distribution: $\sigma^{A^2} = \sigma^{B^2} \in \{0.25, 4, 16\}$. When the variance is low it is easy to learn which is the best colour. Conversely, when the variance is very high, detecting the more rewarding colour is very difficult. Each group of subjects played in all three sub-settings, but the sequence of variance values came in random order and changed across groups. The reason for opting for a dynamic environmental task is the fact that in prolonged static situations little is being learned as nothing disrupts the execution of the activity.

Treatment two of our experiment mainly differs from the previous one in the information given to participants. As in treatment one, the experimental task consisted of repeated a binary choices. Participants had to select one of two colours for three consecutive sub-settings of twenty rounds each. Subjects were randomly assigned to a group of 4 or 5 people whose identity remained unknown. Each group had a leader who was randomly selected and faced the same environmental conditions as everybody else. Leaders played exactly as in treatment one: they did not observe others' behaviour, only their own choices and payoffs were shown to them, exactly as in treatment one. Non-leaders had different information. For them

[3]We checked that the mean and variance in the sample of payoffs were close to those of the underlying normal distribution.

this treatment resembles what Bikhchandani et al. (1992) define as an "observable signals" scenario. Leaders played first. Non-leaders were immediately informed of their leader's choice and payoff. Starting from round two onwards, non-leaders also observed the outcome of their previous personal choice as in treatment one.

As in treatment one, the payoffs were drawn from normal truncated distributions. The mean of the distribution from the most rewarding colour was set equal to 13, whilst the one for the worse option to 10. The pairs of colours between which the individuals had to make a choice changed every twenty rounds, at the end of each sub-setting. One of the two colours was always on average better than the other. The best option was randomly decided and changed every sub-setting. The variance of the payoff distributions, as in the previous treatment, it was set to be either low, medium or high. The order was randomised across groups.

The experiment has been programmed in PHP and administered via computer. All instructions were displayed on the screen. One hundred and seventy five undergraduate students (74 females and 101 males) participated in all parts of this experiment. The experiment was run at Maastricht University during the first week of May 2015. It lasted about one hour and the average payment was of 18.2 euros.

4 Behavioural hypothesis development

The goal of this experiment is to disentangle the relation between self-efficacy and a specific form of observational learning, that is mimicry. Specifically we aim at understanding the extent to which highly self-efficacious people mimic a peer and whether this attitude changes depending on environmental conditions.

One way in which people learn, within a social context, is observing a peer's action. The simplest form of observational learning is mimicry (Byrne, 2002)[4]. It consists in observing and copying the results obtained by a single target. From the followers' perspective, the target becomes a point of reference or focus. The

[4]Concepts like imitation, emulation and mimicry have been highly debated in several disciplines that include neuroscience, psychology, anthropology, sociology and animal behaviour. These concepts mostly have to do with the act of copying somebody's action or performance. Nevertheless, subtle differences amongst the 3 concepts exists. Mimicry is defined as imitation for its own sake. It is considered the less cognitively demanding form of imitation. Imitation implies that a new action is being learned by observing another subject performing it. It also requires a purpose and a means/ends structure. Emulation instead implies that the subject observes somebody acting so as to achieve a goal and tries to achieve that same goal by whatever means. Refer to Hurley and Chater (2005, pp. 1-52) for a detailed review of the literature on the topic. In this paper we will focus on mimicry and emulation because the action taken by the leader is not new in the eyes of the observer. Thus the substantive novelty element at the basis of imitation is absent in our case. We will also, with some approximation, use the words mimicry and emulation interchangeably.

emulator does not necessarily interact with the target, but he can instrumentally use the target's responses to gain extra information about parts of the surrounding environment (Tomasello et al., 1987), or he can unquestionably mimic the target's action. Emulation does not imply any change in the behavioural method: through emulation, subjects do not enlarge their behavioural repertoire.[5]

As mentioned in Section 2, a two-way relation between self-efficacy and social learning exists. On the one hand, self-efficacy evolves vicariously simply observing strategies being modelled by peers. Observers rely on what they see; they elaborate their strategy and form generalized perceptions of their coping abilities. On the other hand, reliance of observational learning is also affected by self-efficacy. High self-efficacy beliefs help people evaluate the effectiveness of the techniques peers use to handle certain events. This consequently helps the observer to assess the steadiness and predictability of the environment (Bandura, 1982). Whilst non-self-efficacious individuals use the observed subject as a role model, self-efficacious people seldom substitute individual learning with emulation. The two phenomena simply complement each other. Observational learning, and emulation more precisely, helps confirm or disprove individual experiences (Bandura, 1989). This behaviour allows individuals to reduce the costs, in terms of effort and risk, of individual trial and error learning (Rendell et al., 2010). It reduces the noise in the individually obtained estimates concerning the state of the world, and avoids incurring in costly mistakes others have already made.

Reliance on emulation also varies depending on environmental conditions or task difficulty. People adaptively fine tune their attention to the behaviour of peers. As the non-stationarity of the environment makes individual learning more difficult, the propensity to rely on observational learning may increase (McElreath et al., 2005). Although it is hard to say *a priori* whether this applies more strongly for highly or weakly self-efficacious individuals, Bandura and Locke (2003, p. 96) claimed that "self-doubt about one's performance efficacy provides incentives to acquire the knowledge and skills needed to master the challenges". Thus, when facing difficult tasks, individuals with high self-efficacy beliefs (and thus lacking self-doubt) presumably feel sufficiently prepared for a challenge and thus less motivated to prepare. Vancouver et al. (2001), using control theory, tested whether individuals with high self-efficacy beliefs devote fewer resources to solve a difficult maze because the discrepancy between the desired level of preparedness and the perception of preparedness is smaller compared to individuals with low self-

[5] Our understanding of emulation is different from the imprinting theory proved by the famous Lorenz experiment in 1935. In our case, following behaviours might emerge because the target subject is a pioneer and explores the environment before everybody else. The target has a temporal advantage in understanding the best choice to be made. Thus, the emulator might be prone to follow. Nonetheless, this is different from the natural instinct that induces goslings to imprint on the first large moving object they see.

efficacy beliefs. A negative relation between between self-efficacy and performance was found by Vancouver et al. (2001, 2002) However, Bandura and Locke (2003) stated that this negative relation found by Vancouver et al. (2001, 2002) was likely to be a misleading consequence of the task used (i.e., decoding game called Mastermind). A learning task rather than a performance task was in Bandura and Locke (2003)'s opinion more appropriate to test this relation.

On the basis of the literature mentioned above, we expect self-efficacious people to rely on emulation to a lesser extent in cases of low environmental volatility. Thus, in easy conditions, we expect to find a negative relation between self-efficacy and the propensity to emulate a target subject. However, in cases of higher environmental variability, it is not clear what effect self-efficacy will have on the propensity to emulate a peer.

5 Theoretical models and estimation strategy

In order to test the above mentioned predictions and study observational learning processes and mimetic behaviours more specifically, we have first to study individual learning patterns. We observed the choices made and payoffs obtained in each round by each agent in treatment one. We select three plausible individual learning models. For each of the three prior-posterior updating rules, we estimate the individual learning parameter (β) and select the best fitting model. Subsequently, we move to the analysis of treatment two whose goal is to understand how people filter available information concerning the behaviour of others. As we consider observational and individual learning as complements, we select two emulation models which respect this assumption. We will use the individual learning parameters that we estimated in treatment one. Using the vector of individual choices taken and reward obtained by agents in treatment two, we estimate the emulation parameters (α or θ).

5.1 Treatment one: Individual learning

We use data from treatment one to detect patterns of individual learning. More specifically, we are interested in uncovering the strength of individual belief of being correct when making binary choices, and in understanding how this relation changes depending on the environmental volatility.

This treatment can be considered "a black box" (Nax et al., 2016). Players take actions and receive payoffs.[6] No information apart from the result of the individual performance is provided to the players.

[6] Differently from Nax et al.'s baseline case, in our case the payoff structure does not depend on others' choices.

Consequently, learning is the result of an asocial process. We consider this treatment as the best possible background to study individual learning. To achieve this goal we assume that, when choosing one of the two colours, individuals behave according to standard logit model.

This means that the probability a participant chooses colour A (Pr_A) in round t is equal to :

$$Pr_t^A = \frac{e^{\beta \hat{\mu}_t^A}}{e^{\beta \hat{\mu}_t^A} + e^{\beta \hat{\mu}_t^B}} \qquad (1)$$

where $\hat{\mu}_t^A$ and $\hat{\mu}_t^B$ represent the subjects current estimates of the mean payoff of colour A and B respectively.

The logit model initially proposed by Luce (1959), is widely used in economics. Generally, it is meant to explain how best responding individuals maximize their expected payoffs based on the distribution of scores they obtained in previous periods. The parameter β is usually interpreted as a measure of rationality (Belloc and Bowles, 2013). The larger is β, the smaller the probability that the individual will deviate from the best response. When $\beta = 0$ the agent chooses randomly between the two alternatives with probability 0.5. As β goes to infinity, the individual never deviates from the best response and the choice is, in that sense, optimal.

In our case, β measures the strength of the belief concerning the estimated average reward for each colour ($\hat{\mu}^A$ and $\hat{\mu}^B$), or differently the strength of the belief of being correct. If $\beta = 0$, the agent has little faith in his own estimation, the difference in the mean of the payoff distributions is neglected and the choice is made randomly. If β goes to ∞, the agent firmly holds onto his estimates and the colour which is thought to be, on average, the most rewarding is always chosen. Intermediate β values could testify the presence of experimentation which might lead to learning. The estimated payoff means for the two colours ($\hat{\mu}^A$ and $\hat{\mu}^B$) are calculated basis of the set of payoffs received over the previous rounds.

There exist several methods that individuals could use to update estimated payoff means. In order to understand how participants made their choices, we fit three different and minimally parametrised learning models which are presented in Table 1. This selection follows the previous literature and more precisely McElreath et al. (2005). These are widely used models in the literature. There is no *a priori* reason to believe one or the other is more likely to apply given our experimental task. Thus, we examine the three alternatives. All three models use a different updating rule to revise the estimated average reward obtained by colour A at time t, $\hat{\mu}_t^A$, considering the choices made $x_1^A ... x_{20}^A$ and payoffs obtained in each round, $y_1^A ... y_{20}^A$.

Table 1: Theoretical models

Model	Updating rule	Free parameters
Running Average	$\hat{\mu}_t^A = \frac{N_{t-1}^A \hat{\mu}_{t-1}^A + y_{t-1}^A}{N_t^A}$	β
Memory Decay	$\hat{\mu}_t^A = r\hat{\mu}_{t-1}^A + (1-r)y_{t-1}^A$	β, r
Bayesian Updating	$\hat{\mu}_t^A = \frac{\frac{\hat{\mu}_{t-1}^A}{\hat{\sigma}_{t-1}^{2^A}} + \frac{y_{t-1}^A}{\sigma^2}}{\frac{1}{\hat{\sigma}_{t-1}^{2^A}} + \frac{1}{\sigma^2}}$ $\hat{\sigma}_t^{2^A} = (\frac{1}{\hat{\sigma}_{t-1}^{2^A}} + \frac{1}{\sigma^2})^{-1}$	β

The first model reported in Table 1, assumes that the individual updates the expected value of the mean of the payoff distributions simply calculating an average of the observed payoffs. At every period the agent adds to his prior $\hat{\mu}_{t-1}^A$ the new payoff he obtained (y_{t-1}^A) and normalises for the number of choices of kind A he has made (N_t^A). The second model, as its name expresses, updates the estimated mean attributing heterogeneous weights (r) to the observed payoffs. It is better known as an adaptive expectations' model or as a weighted running average. Thus, if $r = 0$ information obtained from earlier rounds is completely ignored, and the estimate is equal to the most recent payoff observation. If $r = 1$, the contrary applies. Lastly, under the third model, individuals estimate the distribution mean in a Bayesian fashion assuming that the long-run variance of the payoff distribution (σ^2) is known. In this case, the importance of the most recent payoff obtained is a function of the estimated variance ($\hat{\sigma}_t^{2^A}$) and the real long-run variance (σ^2). For all models analysed, when colour A is not chosen at round t, its estimated payoff mean is assumed to stay equal to previously formulated posterior belief ($\hat{\mu}_t^A = \hat{\mu}_{t-1}^A$). All models share the unknown parameter β which can be estimated with maximum likelihood techniques. The Memory Decay model has an additional parameter, r, which needs to be estimated.

5.1.1 Estimation of the individual learning parameter β

The logit model described by Equation 1 allows us to specify the probability of observing the data conditional on previous choices made and corresponding payoffs obtained. We can fit these models either on an individual basis - obtaining for each individual estimates for the parameters that maximise the likelihood of observing the vector of choices made by each player - or across individuals, pooling the data together, and obtaining one value of the parameter estimate for the entire population of subjects. We first fit the models on pooled data, assuming all participants use the same updating rule. Considering the matrices ($i \times t$ with $t = 1...20$ and i individuals $i = 1...175$) of the observed set of choices made over time by the entire subject pool ($\mathbf{X}$) and corresponding payoffs values obtained ($\mathbf{Y}$), and the scalar of learning parameter β, and discounting factor (r) in the case of the second model, [7] the likelihood function can be written as follows:

$$Pr(\mathbf{X}, \mathbf{Y}|\beta, r) = \prod_{t=1}^{20} Pr(\vec{x}_t, \vec{y}_t|\beta, r, \mathbf{X}_{t-1}, \mathbf{Y}_{t-1}) \tag{2}$$

where $\vec{x}_t$ and $\vec{y}_t$ are the vectors containing all choices made at time t by all individuals, and $\mathbf{X}_{t-1}$ and $\mathbf{Y}_{t-1}$ are the matrices containing past choices and payoffs obtained by all individuals. The logarithm of the likelihood function takes the following form

$$\hat{l} = \ln \mathcal{L}(\beta, r; \mathbf{X}, \mathbf{Y}) = \sum_{t=1}^{20} \ln Pr(\vec{x}_t, \vec{y}_t|\beta, r, \mathbf{X}_{t-1}, \mathbf{Y}_{t-1}) \tag{3}$$

We fit each model reported in Table 1 to the data to retrieve the values of the parameter that maximizes the joint likelihood of observing the pooled data $(\hat{\beta}, \hat{r})$. In order to estimate the parameters we carry a numerical grid search. We used flat homogeneous priors. We set the initial values of $\hat{\mu}$ equal to 9.5 for both colours. The justification for this choice resides in the fact that students were communicated, during the introductory instructions, that the possible payoffs ranged between 1 and 18 and that one of the two options was on average always more rewarding than the other. For the case of the Bayesian model, whose assumption is that agents also know the long-run variance of the payoff distribution, we set σ^2 equal to three values of the real variance of the payoff distributions.

[7] For model 1 and 3 the maximization of the joint likelihood of observing the pooled data is made considering uniquely the β parameter.

Table 2: Goodness of fit measures: pooled data

Variance level	Low	Medium	High
Run. Average			
Log-Lik	1516.89	1995.33	2140.91
$\hat{\beta}$	0.63	0.50	0.49
	(0.018)	(0.018)	(0.022)
AIC	3035.77	3992.67	4283.82
Δ	0.37	0.18	0.12
w	0.03	0.00	0.00
Mem. Decay			
Log-Lik	1512.28	1852.81	1994.63
$\hat{\beta}$	0.60	0.42	0.28
	(0.026)	(0.018)	(0.020)
AIC	3028.57	3709.61	3993.27
Δ	0.38	0.24	0.18
$\hat{r}$	0.61	0.26	0.40
	(0.07)	(0.03)	(0.04)
w	0.97	1.00	1.00
Bayes			
Log-Lik	1520.69	2027.56	2237.38
$\hat{\beta}$	0.73	0.54	0.51
	(0.021)	(0.020)	(0.029)
AIC	3043.39	4057.11	4476.76
Δ	0.37	0.16	0.08
w	0.00	0.00	0.00

Standard errors in parenthesis.

We report in Table 2 the fits of each model on the pooled data for the three variance values. The parameter estimates are shown together with the estimators of the standard error of our parameters obtained taking the square root of the diagonal elements of the inverse of the Hessian matrix [8]. Table 2 also displays some goodness of fit measurements which allow us to compare the models scrutinised (see Burnham and Anderson (1998)). The Akaike Information Criterion (AIC) is the natural logarithm of the likelihood of observing the data plus twice the number of free parameters in the model. There is not a threshold value for the AIC, smaller values indicate a generally better fit. Δ represents the predictive power of each model when compared to a random one. It is calculated as the ratio between the

[8]These values are correct if the observations are independent, their number is large and the models are correct.

negative value of the log-likelihood of the model analysed, m, and the log-likelihood of a model wherein individuals choose randomly ($\Delta_m = 1 - LL_m/LL_{random}$). The value of Δ varies between 1, when the fit of model m is perfect, and 0 when the fit is the same as the random model. The Akaike weights, w, represent instead a way to make comparisons within a specific the set of models. The weight (w_m) for model m within a set of n models (in our case 3) is calculated as follows:

$$w_m = \frac{\exp(-0.5(AIC_m - AIC_{min}))}{\sum_{m=1}^{n} \exp(-0.5(AIC_m - AIC_{min}))} \tag{4}$$

AIC_{min} represents the smallest AIC in the set of model considered. A possible way to interpret the Akaike weights is to think about them as the probability that a given model is correct, thus the highest value of w corresponds to the best fitting model.

According to all goodness of fit measures, Memory Decay seems to be the best model. The AIC values for this model are the lowest regardless the higher number of free parameters. Δ values are the highest for all variance values. Moreover, the Akaike weights (w) show that the probability that people use the Memory Decay rule of updating, compared to the other available models, is approaching or equal to 1 for all variance values.

The estimates of β show that, for all models considered, choices become more random with the increase of the variance in the payoff distributions. This implies that the extent to which individuals believe in their estimates of the payoff means, given the scores obtained from the two colours, decreases - as demonstrated by the declining β – when the variance of the payoff distributions are high or the environment is highly volatile. This first result is in line with what was found by McElreath et al. (2005). This pattern also emerges when looking at the dynamics of the frequency of correct answers per each variance value.

Figure 1: Share of correct answers per round by variance

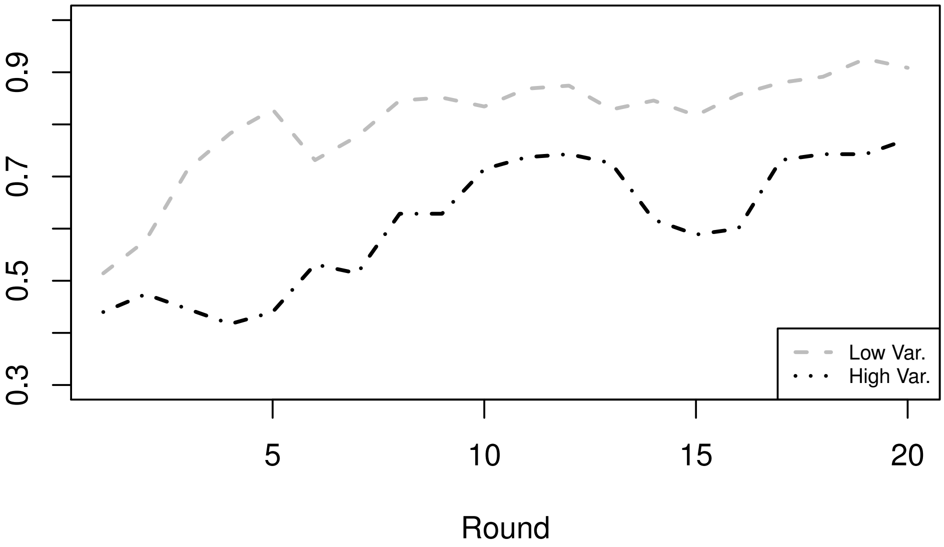

Figure 1 shows that the proportion of correct decisions at each round declines with increasing task difficulty, i.e. variance value. The medium variance case is not shown in this figure. It lies in between these two curves. The decline of β values per sub-setting is also found when analysing the ex-post elicitation of beliefs concerning the average reward of each colour. At the end of every sub-setting, participants were asked to give their best guess concerning the average reward obtained by choosing each of the colours. As can be seen from Figure 2, in case of low variance, the distribution of the answers is nicely peaked around the real means, indicated in blue. Conversely, when the environmental volatility increases, some subjects have more difficulty in understanding which are the real, correct means of payoffs' distributions.

Figure 2: Distribution of posterior beliefs on the means of payoff distributions by variance value: Treatment 1

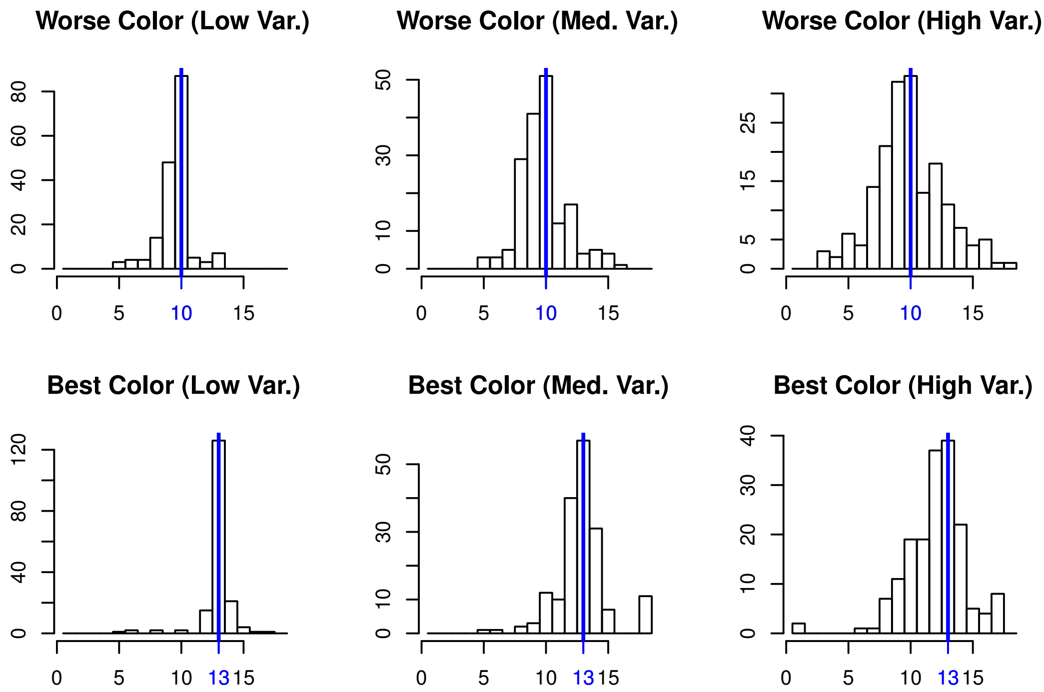

Table 2 also shows that the estimate of the r parameter – the additional unknown parameter for the Memory Decay model – declines, although not monotonically, with increasing variance. Agents pay higher attention to older scores in the low variance case. High variance induces agents to consider lastly obtained payoffs as more informative than older ones.

We also fit the models on an individual basis, allowing the three models presented in table 1 to be possible alternatives to explain the choices made by each participant. In this case, given the vector of choices made by each individual i over 20 rounds ($\vec{x}^i$) and corresponding payoffs received ($\vec{y}^i$) as well as β^i and r^i, for the Memory Decay model, the log-likelihood function can be written as follows:

$$\hat{l} = \ln \mathcal{L}(\beta^i, r^i; \vec{x}^i, \vec{y}^i) = \sum_{t=1}^{20} \ln Pr(x_t^i, y_t^i | \beta^i, r^i, \vec{x}_{t-1}^i, \vec{y}_{t-1}^i) \tag{5}$$

where x_t^i and y_t^i are the choice made and payoff received by individual i at round t and $\vec{x}_{t-1}^i$ and $\vec{y}_{t-1}^i$ represent the vector of length $t-1$ of past choices made and corresponding payoffs.

We estimated the individual learning parameter (β) for each model and each variance value thus obtaining nine $\hat{\beta}$ per agent. In order to retrieve the values of the parameter that maximises the joint likelihood of observing each vector of individual data, we proceeded numerically as we did in the pooled estimation case. We arbitrarily specified the upper bound of the grid within which the algorithm should have carried out the search. It turned out that, in some cases, $\hat{\beta}$ was taking values equal to the upper bound of this grid search, regardless of the value of the upper bound. These values might be divergent. This problem, which mostly emerged for the Memory Decay case, could be due to the fact that $\hat{\beta}$ is computed over only 20 rounds and thus unreliable estimates are produced. This implies that even if a true value of β existed, simply because of statistical variation, our estimation strategy would be unable to estimate it correctly. Thus, in order to test our estimation method, we ran 20000 Montecarlo simulations. We set the true β and r equal to the estimates obtained fitting the Memory Decay model to the pooled data as reported in table 2. We used the payoff distributions used in our experiment, and create fictitious data. We then fit the same Memory Decay model to this data. We checked the distribution of the estimator ($\hat{\beta}$) and it emerged that in case of low variance in the payoff distribution, in 26 % of the cases $\hat{\beta}$ takes extreme values, diverging from the true β. In case of medium variance, the proportion of extreme values declines to 10% reaching 11% when the payoffs' volatility is high (see Appendix 1). These proportions are in line with what we observe in our experimental data. Our estimates are divergent in 37 cases (21%) in the low variance sub-setting, in 31 cases (18%) for medium variance and in 21 (12%) cases when the payoff variance is high. Given these results, we decided to eliminate subjects whose estimated β took on extreme values, and perform our individual learning analysis with a reduced sample size.

A simple count shows that the Memory Decay represents the most frequently used updating rule. As shown in Table 3 this applies to over 68% of individuals in case of low variance, 83% in medium variance and 66% in the high variance case. Bayesian updaters represent 24% of the sample population in low variance, 3% in medium variance and 20% in volatile environments.

Table 3: Number of individuals using the different models

	Low Var.	Med. Var.	High Var.
Running avg.	9	20	20
Mem. Decay	95	120	102
Bayesian Updating	34	4	32
Total	138	144	154

This result points to a wide usage of the Memory Decay updating rule among our agents. As a robustness check, in Table 4 we show the correlations among all individual estimates of β obtained. The correlations across models are high for all variance values. More precisely, the individual estimates of β obtained fitting the Memory Decay model to the data are highly correlated with the estimates obtained fitting the other two candidate models. In both the low and medium variance cases the correlation coefficients are significantly around 70%. In the high variance case, the correlation coefficients decline but remain above 50%.

Table 4: Correlation Individual β across models by variance value

	Low Var.		Med. Var.		High Var.	
	Run.Avg.	Mem. Dec.	Run.Avg.	Mem. Dec.	Run.Avg.	Mem. Dec.
Run.Avg.						
Mem. Dec.	0.75***		0.70***		0.69***	
Bayes	0.93***	0.69***	0.90***	0.81***	0.53***	0.51***

We also analysed within models correlation coefficients (see Table 5). The table shows that the individual learning parameter β varies with the variance of the payoff distributions. Thus, we keep carrying out our analysis for the three levels of environmental volatility.

Table 5: Correlation Individual β within models by variance value

	Run. Avg.		Mem. Dec.		Bayes	
	Low Var.	Med. Var.	Low Var.	Med. Var.	Low Var.	Med. Var.
Low Var.						
Med. Var.	0.18**		0.26***		0.21***	
High Var.	0.19**	0.28***	0.26***	0.29***	0.10	0.18**

Given the simple count reported in Table 3 as well as the high correlation coefficients across estimated β, the Memory Decay model seems to have a clear advantage in predicting individual choices. We decided to consider the individual $\hat{\beta}$ obtained by fitting this model to the data of treatment one to carry the rest of our analysis. [9]

5.2 Treatment two: Emulation

The goal of this second treatment is to understand whether and how people use newly available information concerning the choices and results obtained by a random group leader. As in the previous treatment we will investigate which, among

[9] As a robustness check (not shown) we also carried our analysis with individual $\hat{\beta}$ obtained by fitting the running average, and noticed that the general results do not change.

the possible candidate models of emulation, best predicts the decisional behaviours of our subjects.

This treatment can be considered an "observable signals scenario" (Bikhchandani et al., 1992). A leader plays before everybody else, and his actions and payoffs are observed by his group members from round one onwards. Agents take action after the leader, and from round two onwards they are also presented with the outcome of their previous choice.

As displayed in Table 6, we restrict our analysis to two possible ways in which emulative behaviours can be modelled. Both models are very simple and share two characteristics. They consider sequential choices, each player acts in fact after the leader has made his/her choice. This means that each group member, after his/her choices, has two observations from possibly the same process from which he could estimate the means of the payoff distributions. Moreover both models assume that individuals weight the signals gained observing the leader and use them to validate or invalidate private information or simply mimic them. In both ways, individual trial and error is reduced.

Table 6: Theoretical models

Model	Updating rule	Parameters
Nested model	$Pr_t^A \propto (1-\alpha)L_t^A + \alpha X_t^A$	α
Additive model	$\hat{\mu}'^A = (1-\theta)\hat{\mu}_{i,t-1}^A + \theta y_{j,t}^A$ $Pr_{i,t}^A = \frac{e^{\beta\hat{\mu}'^A}}{Z}$ $\hat{\mu}_{i,t}^A = r\hat{\mu}'^A + (1-r)y_{i,t}^A$	θ

First, in line with McElreath et al. (2005) this situation can be modelled using a nested probability model. Specifically, the probability of choosing colour A at time t is given by

$$Pr_t^A \propto (1-\alpha)L_t^A + \alpha X_t^A \tag{6}$$

where X^A is an indicator variable taking the value 1 if the leader chose colour A at time t, and is 0 otherwise. L_t^A is the probability that subject i would choose colour A at t were he playing alone, as defined in Equation 1. This probability is calculated in a conservative manner. We relied on our previous estimates from treatment one, using the individual estimates of β and r obtained by fitting the Memory Decay model to the data from treatment one. The relevant unknown parameter to be estimated in equation 6 is α. This parameter measures the individual propensity to emulate. If α is zero, the model reduces to the simple individual learning process. The agent fully relies on individual learning and any information provided by the leader's action is dismissed. In case α is 1, it means that the individual perfectly mimics the behaviour of the target subject.

Second, emulation is modelled as an additive process according to which individuals add the information retrieved from the leader's action to their private information. The principle here is that the quality of the information gathered from the observation of the leader is exactly the same as that of any other agent. The leader's signal is not *a priori* any better than anyone else's. Thus, at the end of round one, group members hold two pieces of information from possibly the same payoff distribution. The relevant unknown parameter to be estimated is, in this case, θ which measures again the individual propensity to emulate. More specifically, θ captures how much individuals value the experience of the leader. In this case, for estimating the parameters, the sequence we assume is the following. At the start of each round the subject has, for each colour, A and B, a prior belief of the mean payoffs $(\hat{\mu}^A, \hat{\mu}^B)$. The leader plays (j) and the non-leader observes the leader's choice, for instance A, and the corresponding payoff $y^A_{j,t}$. According to the updating rule of Table 6, the subject updates his estimate of mean payoff for colour A to $\hat{\mu}'^A$. Based on the pair $(\hat{\mu}'^A, \hat{\mu}^B)$ the subject makes his choice following the standard logit model presented in Equation 1 and reported in Table 6. We use the previously estimated individual β to calculate this probability. Observing his own choice and the corresponding payoff obtained $(y^A_{i,t})$, the subject updates his estimate of the mean payoffs to a new couple $(\hat{\mu}^A, \hat{\mu}^B)$ following the Memory Decay rule (see Table 1). As discussed in the previous Section, statistically, the Memory Decay model works best for updating based on individual information. Thus in this step we use that rule, applying for each subject the value of Memory Decay, r, fitted from the first, individual play, treatment. We assume that in period $t = 0$ our agents start off with flat priors (uniform on $[1, 18]$), on the unknown mean of the payoff distributions.

We fit the two models to the individual data from treatment 2 and estimate the two emulation parameters for each variance value.

First we report simple information concerning the choices made by our subjects in treatment 2.

When comparing the vector of "correct" answers provided in treatment one, i.e. when playing individually, with those of treatment 2, i.e. when led by a group leader, it can be seen that on average performance increases (see Figure 3). The

Figure 3: Comparison correct answers per round by variance: Individual and Leader treatment

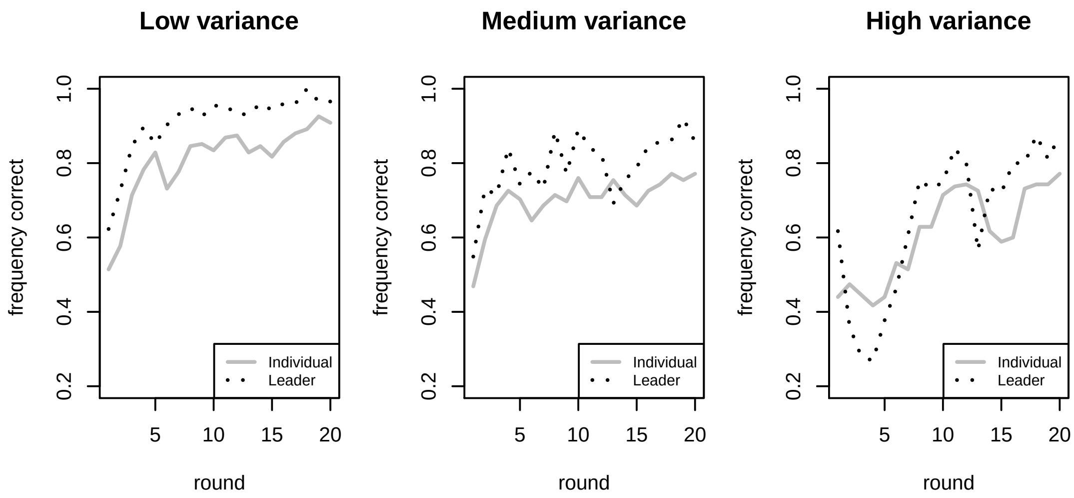

share of correct answers per round in treatment 2 is generally higher than that registered in the individual treatment both in case of low and medium variance. In the high variance case, the effect of playing with a leader is unclear. Signals are mixed, and while the payoffs distribution had the same moments both with and without the leader, there is no temporal pattern. This is almost certainly driven by the fact that in the high variance case, subjects were simply unable to detect which colour was superior, due to the large noise in the signals they receive. Information could not be extracted either from their own signals or from the signals given by the leader's play.

We also report, in Figure 4, the distribution of posterior beliefs on the means of the payoff distributions retrieved by end-of-sub-setting questionnaires. After 20 rounds, students were asked to give their best guess of the average rewards given by the 2 colours they had to choose from. The same pattern noticed in Figure 2 is present. In a more volatile environment, fewer individuals correctly guess the real means of the payoff distributions.

Figure 4: Distribution of posterior beliefs on the means of payoff distributions by variance value: Treatment 2

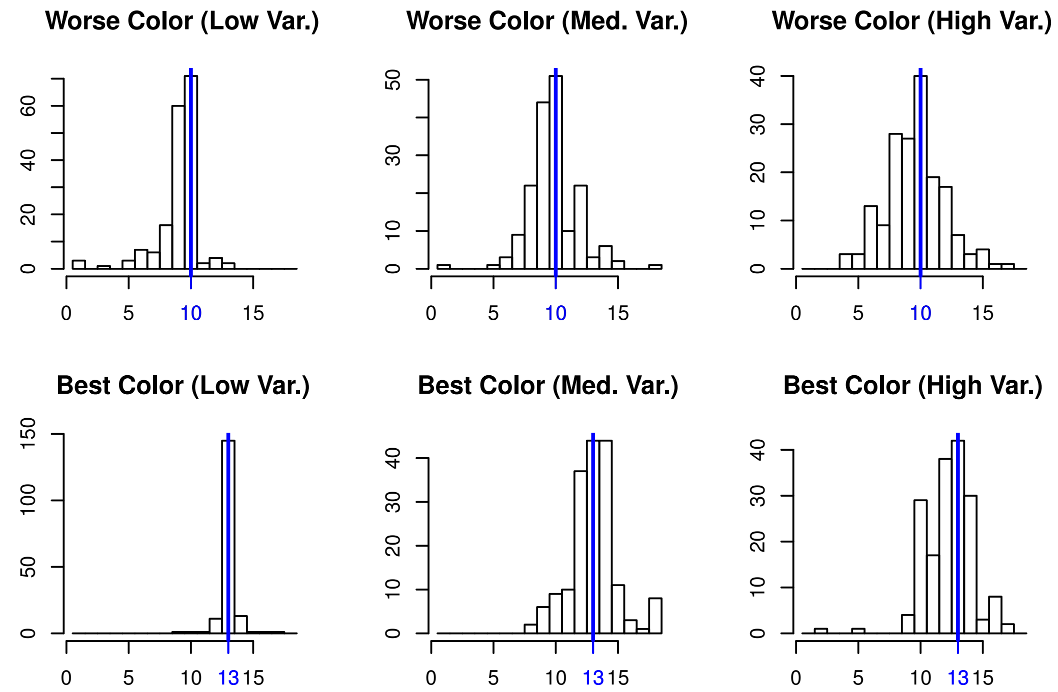

Table 7 reports the number of individual whose posterior belief on which one was the most rewarding colour was correct. These posterior beliefs were elicited by a between sub-setting questionnaire. It can be seen that in the low and medium variance a high number of individuals guessed correctly which one was the colour which yielded the highest scores independently from the fact of being led by the group leader. In case of high variance, the additional signal provided by the observation of the leader, instead, increases the number of correct beliefs.

Table 7: Number of individuals whose posterior on which was the best colour was correct: Individual and Leader treatment

	Individual Treatment	Leader Treatment
Low Var.	171	173
Med. Var.	163	163
High Var.	149	157

6 Results

We are now at the last step of our analysis. This concerns the relation between the propensity to emulate and self-efficacy perceptions. Specifically, we try to detect whether self-efficacious people emulate a peer and how this changes depending on environmental conditions. As mentioned in Section 4, drawing on psychology literature, one can expect to find a negative relation between propensity to emulate and self-efficacy beliefs only in case of little environmental volatility. In case of medium or high variance, we do not have a clear prediction regarding what the effect of self-efficacy will be on the tendency to emulate a peer. We do know that higher variance in the payoffs introduces noise in the signals received by the players, thus making the relationship between self-efficacy and propensity to emulate harder to detect.

In Table 8, we report the results of our preferred OLS regressions. Our dependent variable, i.e. the logarithm of θ or α, captures individuals' propensity to emulate. The logarithm of the results of the self-efficacy questionnaires represent our main regressor. We use math abilities as a control. In fact, it could be argued that higher or lower reliance on the leader's performance could, in this specific game setting, be related to individuals' logic and mathematical skills. This, to some extent, also controls for people's (in)abilities to understand the task at stake. Additionally, one might argue that emulating the leader's action could be plausible especially if the leader is obtaining good results from his/ her choices. For this reason, we control for the score obtained by the group leader.

In this regression analysis we excluded the 38 leaders (since they are playing individually and have no leader to follow or not) and those whose β estimates were at the extreme values, as discussed in Section 5.1.1.[10] Thus, the sample size reduces to 109, 113 or 122 people depending on the variance value.

[10] In some cases, the individuals whose β values took on extreme values are also leaders in treatment 2.

Table 8: Regression table

			Low Variance			
dep. var.	log(θ)	log(θ)	log(θ)	log(α)	log(α)	log(α)
Const	8.86 (4.39)*	9.22 (4.56)*	-233.84 (67.61)***	6.19 (4.33)	6.66 (4.39)	-208.91 (58.52)***
log(self-eff.)	-3.08 (1.29)**	-3.24 (1.38)**	-3.39 (1.53)**	-2.57 (1.26)**	-2.78 (1.31)**	-2.91 (1.29)**
log(math)		0.13 (0.39)	0.19 (0.34)		0.17 (0.35)	0.23 (0.35)
log(scoreleader)			43.98 (12.19)***			39.00 (10.63)***
R^2	0.04	0.04	0.14	0.03	0.03	0.13
obs.	109	109	109	109	109	109
			Medium Variance			
Const	2.91 (5.9)	4.01 (6.12)	-6.10 (33.78)	-1.89 (4.99)	-0.77 (5.07)	-99.64 (27.64)***
log(self-eff.)	-1.43 (1.71)	-1.96 (1.81)	-1.98 (1.81)	-0.27 (1.45)	-0.8 (1.51)	-1.07 (1.45)
log(math)		0.47 (0.39)	0.48 (0.39)		0.48 (0.33)	0.44 (0.33)
log(scoreleader)			1.86 (6.10)			18.22 (5.29)***
R^2	0.01	0.03	0.02	0.00	0.02	0.11
obs.	113	113	113	113	113	113
			High Variance			
Const	-3.65 (4.2)	-2.84 (4.23)	-58.21 (18.70)***	-4.67 (4.04)	-4.41 (4.01)	-29.94 (26.93)
log(self-eff.)	0.58 (1.22)	0.15 (1.26)	0.26 (1.23)	0.62 (1.17)	0.49 (1.18)	0.46 (1.18)
log(math)		0.45 (0.33)	0.30 (0.31)		0.14 (0.32)	0.13 (0.30)
log(scoreleader)			10.08 (3.33) ***			4.67 (4.88)
R^2	0	0.02	0.08	0	0	0.01
obs.	122	122	122	122	122	122

Robust standard errors in parenthesis $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

The results of the log-log regressions can be easily interpreted as elasticities.[11] As it can be seen in the upper part of Table 8, 1% increase in self-efficacy corresponds to about a 3% decrease in the propensity to emulate. This result is stable regardless of the social learning model used. Its significance is not affected when controlling for math abilities or the leader score. In easy environments, our hypothesis is thus supported: weakly self-efficacious individuals are more apt at following the leader's action. In the medium and high variance cases instead, we cannot reject the null hypothesis of no significant relation between self-efficacy beliefs and propensity to imitate the target.

7 Conclusions

Human beings strive for coherence but they are boundedly rational or 'cognitive misers'. They follow rules for purely instrumental reasons (i.e. cost-benefit analysis calculations) as well as for non-instrumental ones. In addition to the non-instrumental reasons behind rule following behaviours mentioned in the introduction, this paper contributes to uncover another possible one: low self-efficacy beliefs. Low self-efficacy beliefs affect people's propensity to emulate a pioneer peer, and this in turn might condition institutional evolution.

We designed an experimental laboratory game which is a modified version of the two-armed bandit problem with finite time horizon. We retrieved information on students' self-efficacy beliefs and considered this information as static and exogenous. The main goal of this paper was to evaluate to what extent introspective beliefs concerning one's abilities to understand an control external environments affect the propensity to mimic a peer's actions. To correctly assess mimetic behaviours, we first studied how individuals play when alone. We assumed people choose between two competing options following a standard logit model. Within a set of alternative updating rules, we selected the one that best fits the choices made by our subjects. Second, we developed some computational models to study emulation, i.e. the simplest form of observational learning. Our models do not consider mimicry as a substitute to individual learning. We consider emulative behaviours as the by-product of individual learning and the opportunity to observe a peer.

Our results point to a mildly significant negative relation between self-efficacy beliefs and emulative behaviours in scarcely volatile environments. Regardless of the model used, more self-efficacious agents seem follow their group leader to a lesser extent. This establishes a nexus between individuals' introspective beliefs and the way in which agents learn when given the possibility to observe a peer's

[11]The results of lin-lin, log-lin and lin-log specifications can be found in Appendix 2.

actions and outcomes. Our study, taking into account human cognitive capacities, seems to reveal another possible micro-mechanism able to explain why people abide by the rules.

We believe that our finding sheds some light on the reasons why institutional change can be difficult. Given the design of our experiment, the leader is providing a focal point for behaviour. His pioneer nature makes his behaviour salient, and displaying the leader's actions and outcomes is likely to prompt compliance. The mere fact that all group members see the choice made by the leader transforms this choice into a focal behaviour. This is a common feature of institutions. Sheingate (2009) claims that institutional change occurs when individuals, leveraging on ambiguities, use their agency and creativity to establish new precedents for action and transform the way institutions work and power is allocated (Sheingate, 2009). Our study complements this analysis. Self-efficacy, defined by Bandura (1989) as the most central mechanism affecting human agency, affects the extent to which people are prone to deviate from established behavioural paths which in our experiment are set by the leaders. Low self-efficacy implies considering the leader as a role model and thus prompts subjects to mimic perfectly his/her actions. In a society characterised by scarce volatility and where the general population has low self-efficacy beliefs, people are more inclined to follow, and to mimic the behaviour of prominent individuals regardless of the goodness of their actions. This suggests that in these situations, institutions will tend to be strong, and difficult to change from below. Following paths of behaviour established by leading individuals seems likely. And, unless an alternative visible leadership emerges and sets an example of behaviour that can be successfully imitated, bottom-up change is less likely to happen. This corroborates the empirical results obtained by Bernard et al. (2011) who found evidence that fatalistic beliefs and low self-efficacy in rural households in Ethiopia reduce the demand for long-term loans and for loans for productive purposes, thus possibly stemming long-term poverty. Aspiration improved and investment behaviour changed once individuals were invited to watch documentaries about people belonging to similar communities who had succeeded in agriculture or small business (Dercon et al., 2014). It also echoes a recent explanation provided by Acemoglu et al. (2014) of the fact that individuals, despite the low level of development experienced, remain loyal and unquestionably respect the village chief. Villagers have internalised a patron-client relationship. In our words villagers, given frequent exposure to coercion by village chiefs, might feel that behavioural outcomes depend on forces outside their control and that they do not have the capacity to reach higher well-being. This leads to aspiration failures: they accept the *status quo* together with the rules of behaviour this imposes. Our result is also in line with Gorodnichenko and Roland (2010, 2011) according to which societies wherein intellectual and affective autonomy is promoted are more innovative.

By contrast, when self-efficacy perceptions are high and environments relatively stable and certain, individuals are willing to try new behaviours and deviate from the prescriptions of focal behaviours. While this may pose challenges for the details of what a new institution would look like, it does make existing institutions more susceptible to change and more specifically to displacement (the dismissal and replacement of existing rules with new ones) and layering (the introduction of new rules on top of or aside to existing ones) (Streeck and Thelen, 2005). Results are instead inconclusive for medium and highly volatile environments.

We nevertheless want to stress some of the *caveats* of our analysis. First, we were able to explore only some of the possible models able to capture individual learning and mimicry. Second, as for any other laboratory experiment some external validity concerns apply. These results should be validated in a natural or quasi-natural environment. Third, we studied a very specific form of observational learning which is mimicry. We restricted our analysis to a case in which people observe a peer and possibly mimic his action. We did not consider the effect of social interactions in the form of communication for example. Taking this into account would surely represent a key element to formulate a realistic representation of institutional evolution. Lastly, one should carefully acknowledge our assumption of considering self-efficacy as an exogenous and static variable. Bandura himself claims that whilst self-efficacy affects behaviour and learning, learning processes feedback into self-efficacy perceptions. Considering this agency-structure feedback loop would be extremely important. Such attempt would be in line with the steps taken by many institutional economists and sociologists towards an integration of social, cognitive and institutional analysis (Steinmo, 2008; Mahoney and Thelen, 2009).

8 Appendix 1

Figure 5: Distribution $\hat{\beta}$: Montecarlo simulations (Low Variance)

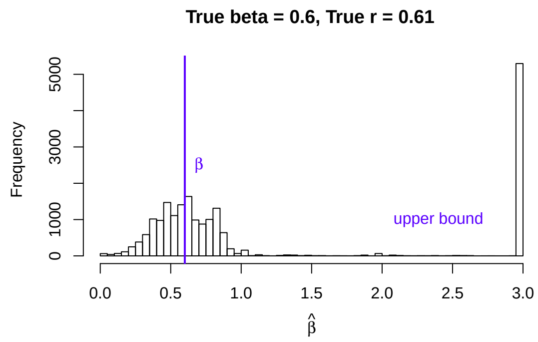

Figure 6: Distribution $\hat{\beta}$: Montecarlo simulations (Medium Variance)

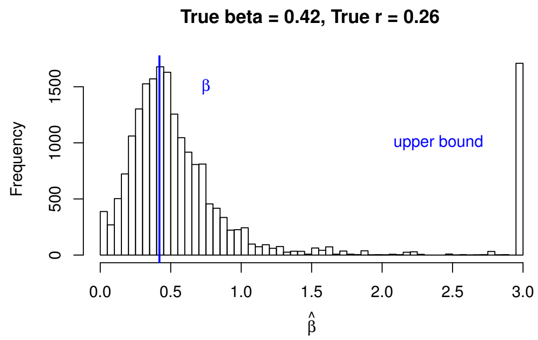

Figure 7: Distribution $\hat{\beta}$: Montecarlo simulations (High Variance)

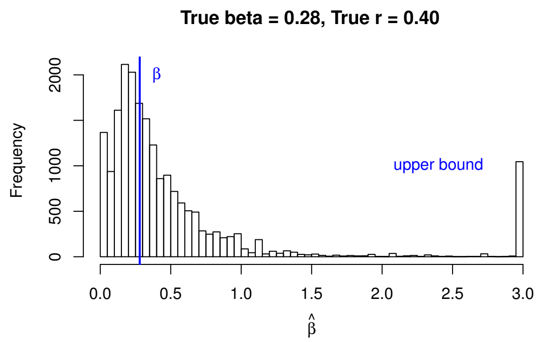

9 Appendix 2

Table 9: Regression table: Alpha Low and Medium Variance

			Low Variance			
dep. var.	α	α	α	$\log(\alpha)$	$\log(\alpha)$	$\log(\alpha)$
Const	0.65 (0.2)***	0.64 (0.21)***	-6.88 (1.89)***	-0.14 (1.39)	-0.15 (1.41)	-39.04 (10.71)
self-eff.	-0.01 (0.006)**	-0.01 (0.006)**	-0.01 (0.006)**	-0.08 (0.04)*	-0.08 (0.05)*	-0.09 (0.05)*
math		0.00 (0.01)	0.00 (0.01)		0.03 (0.09)	0.05 (0.08)
leaderscore			0.03 (0.007)***			0.15 (0.04)***
R^2	0.03	0.03	0.20	0.03	0.03	0.12
obs.	109	109	109	109	109	109
Const	1.67 (0.63)**	1.71 (0.62)**	-39.67 (10.48)***	6.19 (4.33)	6.66 (4.39)	-208.91 (58.52)***
log(self-eff.)	-0.42 (0.18)**	-0.44 (0.18)**	-0.46 (0.17)**	-2.57 (1.26)**	-2.78 (1.31)**	-2.91 (1.29)**
log(math)		0.01 (0.05)	0.03 (0.05)		0.17 (0.35)	0.23 (0.35)
log(leaderscore)			7.5 (1.90)***			39.00 (10.63)***
R^2	0.04	0.04	0.20	0.03	0.03	0.13
obs.	109	109	109	109	109	109
			Medium Variance			
dep. var.	α	α	α	$\log(\alpha)$	$\log(\alpha)$	$\log(\alpha)$
Const	0.25 (0.18)	0.23 (0.18)	-2.76 (0.84)***	-2.32 (1.51)	-2.45 (1.49)	-20.37 (5.08)***
self-eff.	0 (0.01)	0 (0.01)	0 (0.01)	-0.02 (0.05)	-0.03 (0.05)	-0.03 (0.05)
math		0.02 (0.01)	0.01 (0.01)		0.13 (0.08)	0.11 (0.09)
leaderscore			0.01 (0.003)***			0.07 (0.02)***
R^2	0	0.02	0.16	0	0.02	0.12
obs.	113	113	113	113	113	113
Const	0.31 (0.6)	0.46 (0.61)	-15.97 (4.52)***	-1.89 (4.99)	-0.77 (5.07)	-99.64(27.40)***
log(self-eff.)	-0.03 (0.18)	-0.11 (0.18)	-0.15 (0.177)	-0.27 (1.45)	-0.8 (1.51)	-1.07 (1.35)
log(math)		0.07 (0.04)	0.06 (0.04)		0.48 (0.33)	0.44 (0.32)
log(leaderscore)			3.03 (0.82)***			18.22 (5.29)***
R^2	0	0.02	0.16	0	0.02	0.11
obs.	113	113	113	113	113	113

Robust standard errors in parenthesis $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 10: Regression table: Alpha High Variance

	High Variance					
dep. var.	α	α	α	$\log(\alpha)$	$\log(\alpha)$	$\log(\alpha)$
Const	0.03 (0.16)	0.03 (0.16)	-0.09 (0.62)	-3.27 (1.25)**	-3.31 (1.26)**	-7.96 (5.02)
self-eff.	0.01 (0.01)	0.01 (0.01)	0.01 (0.01)	0.02 (0.04)	0.02 (0.04)	0.02 (0.04)
math		0 (0.01)	0 (0.01)		0.03 (0.07)	0.03 (0.07)
leaderscore			0 (0.002)			0.02 (0.02)
R^2	0.01	0.01	0.01	0	0	0.01
obs.	122	122	122	122	122	122
Const	-0.3 (0.5)	-0.3 (0.51)	-1.01 (3.49)	-4.67 (4.04)	-4.41 (4.01)	-29.94 (26.93)
log(self-eff.)	0.15 (0.15)	0.14 (0.15)	0.14 (0.15)	0.62 (1.17)	0.49 (1.18)	0.46 (1.18)
log(math)		0 (0.04)	0 (0.04)		0.14 (0.32)	0.13 (0.30)
log(leaderscore)			0.13 (0.64)			4.67 (4.88)
R^2	0.01	0.01	0.01	0	0	0.01
obs.	122	122	122	122	122	122

Robust standard errors in parenthesis $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 11: Regression table: Theta Low and Medium Variance

Low Variance

dep. var.	θ	θ	θ	$\log(\theta)$	$\log(\theta)$	$\log(\theta)$
Const	1.21 (0.34)***	1.21 (0.35)***	-5.23 (2.90)*	1.41 (1.42)	1.39 (1.44)	-42.21 (12.11)***
self-eff.	-0.02 (0.01)**	-0.02 (0.01)**	-0.02 (0.01)**	-0.1 (0.05)**	-0.1 (0.05)**	-0.1 (0.05)**
math		0.01 (0.02)	0.01 (0.02)		0.02 (0.1)	0.04 (0.1)
leaderscore			0.02 (0.011)**			0.17 (0.05)***
R^2	0.04	0.04	0.08	0.03	0.03	0.13
obs.	109	109	109	109	109	109
Const	2.91 (1.06)**	2.99 (1.1)**	-33.01 (15.73)**	8.86 (4.39)*	9.22 (4.56)*	-233.84 (67.61)***
log(self-eff.)	-0.7 (0.31)**	-0.74 (0.33)**	-0.75 (0.34)**	-3.08 (1.29)**	-3.24 (1.38)**	-3.39 (1.53)**
log(math)		0.03 (0.09)	0.04(0.09)		0.13 (0.39)	0.19 (0.34)
log(leaderscore)			6.51 (2.87)**			43.99(12.19)***
R^2	0.04	0.04	0.08	0.04	0.04	0.14
obs.	109	109	109	109	109	109

Medium Variance

dep. var.	θ	θ	θ	$\log(\theta)$	$\log(\theta)$	$\log(\theta)$
Const	0.81 (0.36)**	0.79 (0.36)**	-0.13 (1.20)	-0.44 (1.69)	-0.57 (1.7)	-2.04 (6.29)
self-eff.	-0.01 (0.01)	-0.02 (0.01)	-0.16 (0.01)	-0.05 (0.05)	-0.07 (0.05)	-0.06 (0.05)
math		0.02 (0.02)	0.003 (0.00)		0.12 (0.09)	0.12 (0.09)
leaderscore			0.003 (0.00)			0.00(0.02)
R^2	0.02	0.03	0.03	0.01	0.03	0.03
obs.	113	113	113	113	113	113
Const	1.83 (1.23)	1.97 (1.26)	-3.58 (6.33)	2.91 (5.9)	4.01 (6.12)	-6.01 (33.78)
log(self-eff.)	-0.42 (0.36)	-0.49 (0.37)	-0.50 (0.36)	-1.43 (1.71)	-1.96 (1.81)	-1.98 (1.81)
log(math)		0.06 (0.08)	0.05 (0.07)		0.48 (0.39)	0.47 (0.39)
log(leaderscore)			1.02 (1.14)			1.86 (6.10)
R^2	0.02	0.02	0.03	0.01	0.03	0.02
obs.	113	113	113	113	113	113

Robust standard errors in parenthesis $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 12: Regression table: Theta High Variance

	High Variance					
dep. var.	θ	θ	θ	$\log(\theta)$	$\log(\theta)$	$\log(\theta)$
Const	0.16 (0.3)	0.13 (0.3)	-1.23 (0.79)	-2.33 (1.32)*	-2.46 (1.33)*	-12.39 (3.52)***
self-eff.	0.01 (0.01)	0.01 (0.01)	0.01 (0.01)	0.02 (0.04)	0.01 (0.04)	0.01 (0.04)
math		0.02 (0.02)	0.02 (0.02)		0.1 (0.08)	0.05 (0.07)
Leader score			0.00 (0.00)			0.04 (0.01)***
R^2	0.01	0.02	0.04	0	0.02	0.08
obs.	122	122	122	122	122	122
Const	-0.38 (0.98)	-0.23 (0.99)	-7.91 (4.15)*	-3.65 (4.2)	-2.84 (4.23)	-58.21 (18.70)***
log(self-eff.)	0.24 (0.29)	0.16 (0.3)	0.17 (0.29)	0.58 (1.22)	0.15 (1.26)	0.26 (1.23)
log(math)		0.08 (0.07)	0.06 (0.06)		0.45 (0.33)	0.30 (0.31)
log(Leaderscore)			1.39 (0.73)*			10.08 (3.33)***
R^2	0.01	0.02	0.04	0	0.02	0.08
obs.	122	122	122	122	122	122

Robust standard errors in parenthesis $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

References

Acemoglu, D., T. Reed, and J. A. Robinson (2014). Chiefs: Economic development and elite control of civil society in sierra leone. *Journal of Political Economy 122*(2), 319–368.

Bandura, A. (1982). Self-efficacy mechanism in human agency. *American psychologist 37*(2), 122–147.

Bandura, A. (1989). Perceive self efficacy in the exercise agency: the psychologist. *Bulletin of the British Psychological Society 10*, 411–424.

Bandura, A. (1997). *Self-efficacy: The exercise of control.* New York : W.H. Freeman.

Bandura, A. and E. A. Locke (2003). Negative self-efficacy and goal effects revisited. *Journal of applied psychology 88*(1), 87.

Banerjee, A. V. (1992). A simple model of herd behavior. *The Quarterly Journal of Economics 107*(3), 797–817.

Banks, J., M. Olson, and D. Porter (1997). An experimental analysis of the bandit problem. *Economic Theory 10*(1), 55–77.

Baumeister, R. F. and M. R. Leary (1995). The need to belong: desire for interpersonal attachments as a fundamental human motivation. *Psychological bulletin 117*(3), 497–529.

Belloc, M. and S. Bowles (2013). The persistence of inferior cultural-institutional conventions. *The American Economic Review 103*(3), 93–98.

Bernard, T., S. Dercon, and A. S. Taffesse (2011). Beyond Fatalism - An empirical exploration of self-efficacy and aspirations failure in Ethiopia.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy 100*(5), 992–1026.

Burnham, K. P. and D. R. Anderson (1998). *Model selection and multimodel inference: a practical information-theoretic approach.* Springer Science & Business Media.

Byrne, R. W. (2002). Imitation of novel complex actions: what does the evidence from animals mean? *Advances in the Study of Behavior 31*, 77–105.

Cialdini, R. B. and N. J. Goldstein (2004). Social influence: Compliance and conformity. *Annual Review of Psychology 55*, 591–621.

Denzau, A. T. and D. C. North (1994). Shared mental models: ideologies and institutions. *Kyklos 47*(1), 3–31.

Dercon, S., T. Bernard, K. Orkin, and A. S. Taffesse (2014, April). The Future in Mind: Aspirations and Forward-Looking Behaviour in Rural Ethiopia. (WPS/2014-16).

Gans, N., G. Knox, and R. Croson (2007). Simple models of discrete choice and their performance in bandit experiments. *Manufacturing & Service Operations Management 9*(4), 383–408.

Gittins, J. C. (1979). Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society. Series B (Methodological) 41*(2), 148–177.

Gorodnichenko, Y. and G. Roland (2010). Culture, institutions and the wealth of nations. *National Bureau of Economic Research*.

Gorodnichenko, Y. and G. Roland (2011, May). Which dimensions of culture matter for long-run growth? *American Economic Review 101*(3), 492–98.

Greif, A. and C. Kingston (2011). Institutions: Rules or equilibria? In N. Schofield and G. Caballero (Eds.), *Political economy of institutions, democracy and voting*, pp. 13–43. Springer.

Greif, A. and D. D. Laitin (2004). A theory of endogenous institutional change. *The American Political Science Review 98*(4), 633–652.

Henrich, J. and R. McElreath (2003). The evolution of cultural evolution. *Evolutionary Anthropology: Issues, News, and Reviews 12*(3), 123–135.

Hurley, S. L. and N. Chater (2005). *Perspectives on Imitation: from neuroscience to social science*. Cambridge, Mass. ; London : MIT.

Judge, T. A. and J. E. Bono (2001). Relationship of core self-evaluations traits - self-esteem, generalized self-efficacy, locus of control, and emotional stability - with job satisfaction and job performance: A meta-analysis. *Journal of Applied Psychology 86*(1), 80–92.

Kahneman, D. (2011). *Thinking, fast and slow*. Macmillan.

Kingston, C. and G. Caballero (2009). Comparing theories of institutional change. *Journal of Institutional Economics 5*(02), 151–180.

Kosfeld, M., A. Okada, and A. Riedl (2009). Institution formation in public goods games. *The American Economic Review 99*(4), 1335–1355.

Luce, R. (1959). *Individual Choice Behavior a Theoretical Analysis*. New York : Wiley.

Mahoney, J. and K. Thelen (2009). *Explaining institutional change: ambiguity, agency, and power*. Cambridge University Press.

Mantzavinos, C., D. C. North, and S. Shariq (2004). Learning, institutions, and economic performance. *Perspectives on politics 2*(01), 75–84.

McElreath, R., M. Lubell, P. J. Richerson, T. M. Waring, W. Baum, E. Edsten, C. Efferson, and B. Paciotti (2005). Applying evolutionary models to the laboratory study of social learning. *Evolution and Human Behavior 26*(6), 483–508.

Nax, H. H., M. N. Burton-Chellew, S. A. West, and H. P. Young (2016). Learning in a black box. *Journal of Economic Behavior & Organization 127*, 1–15.

North, D. C. (1990). *Institutions, institutional change and economic performance*. Cambridge university press.

North, D. C. (2005). Institutions and the process of economic change. *Management International 9*(3), 1–7.

Rendell, L., R. Boyd, D. Cownden, M. Enquist, K. Eriksson, M. W. Feldman, L. Fogarty, S. Ghirlanda, T. Lillicrap, and K. N. Laland (2010). Why copy others? insights from the social learning strategies tournament. *Science 328*(5975), 208–213.

Schwarzer, R., J. Bäßler, P. Kwiatek, K. Schröder, and J. X. Zhang (1997). The assessment of optimistic self-beliefs: comparison of the german, spanish, and chinese versions of the general self-efficacy scale. *Applied Psychology 46*(1), 69–88.

Schwarzer, R. and M. Jerusalem (1995). Generalized self-efficacy scale. In J. Weinman, S. Wright, and M. Johnston (Eds.), *Measures in health psychology: A users portfolio. Causal and control beliefs*, pp. 35–37. Windsor, UK: NFER-NELSON.

Sheingate, A. (2009). Rethinking rules. In J. Mahoney and K. Thelen (Eds.), *Explaining Institutional Change: Ambiguity, Agency, and Power*, pp. 168–203. Cambridge University Press.

Simon, H. A. (1959). Theories of decision-making in economics and behavioral science. *The American Economic Review 49*(3), 253–283.

Steinmo, S. (2008). Historical institutionalism. In D. Della Porta and M. Keating (Eds.), *Approaches and Methodologies in the Social Sciences: A Plurali*, pp. 118–138. Cambridge University Press.

Streeck, W. and K. A. Thelen (2005). Institutional change in advanced political economies. In W. Streeck and K. A. Thelen (Eds.), *Beyond continuity*, pp. 1–39. Oxford University Press.

Tomasello, M., M. Davis-Dasilva, L. Camak, and K. Bard (1987). Observational learning of tool-use by young chimpanzees. *Human evolution 2*(2), 175–183.

Vancouver, J. B., C. M. Thompson, E. C. Tischner, and D. J. Putka (2002). Two studies examining the negative effect of self-efficacy on performance. *Journal of Applied Psychology 87*(3), 506.

Vancouver, J. B., C. M. Thompson, and A. A. Williams (2001). The changing signs in the relationships among self-efficacy, personal goals, and performance. *Journal of Applied Psychology 86*(4), 605–620.

The UNU-MERIT Working Paper Series

2016-01 *Mexican manufacturing and its integration into global value chains* by Juan Carlos Castillo and Adam Szirmai

2016-02 *New variables for vocational secondary schooling: Patterns around the world from 1950-2010* by Alison Cathles

2016-03 *Institutional factors and people's preferences in social protection* by Franziska Gassmann, Pierre Mohnen & Vincenzo Vinci

2016-04 *A semi-endogenous growth model for developing countries with public factors, imported capital goods, and limited export demand* by Jan Simon Hallonsten and Thomas Ziesemer

2016-05 *Critical raw material strategies in different world regions* by Eva Barteková and René Kemp

2016-06 *On the value of foreign PhDs in the developing world: Training versus selection effects* by Helena Barnard, Robin Cowan and Moritz Müller

2016-07 Rejected Afghan asylum seekers in the Netherlands: Migration experiences, current situations and future aspirations

2016-08 *Determinants of innovation in Croatian SMEs: Comparison of service and manufacturing firms* by Ljiljana Bozic and Pierre Mohnen

2016-09 *Aid, institutions and economic growth: Heterogeneous parameters and heterogeneous donors* by Hassen Abda Wakoy

2016-10 *On the optimum timing of the global carbon-transition under conditions of extreme weather-related damages: further green paradoxical results* by Adriaan van Zon

2016-11 *Inclusive labour market: A role for a job guarantee* scheme by Saskia Klosse and Joan Muysken

2016-12 *Management standard certification and firm productivity: micro-evidence from Africa* by Micheline Goedhuys and Pierre Mohnen

2016-13 *The role of technological trajectories in catching-up-based development: An application to energy efficiency technologies* by Sheng Zhong and Bart Verspagen

2016-14 *The dynamics of vehicle energy efficiency: Evidence from the Massachusetts Vehicle Census* by Sheng Zhong

2016-15 *Structural decompositions of energy consumption, energy intensity, emissions and emission intensity - A sectoral perspective: empirical evidence from WIOD over 1995 to 2009* by Sheng Zhong

2016-16 *Structural transformation in Brazil, Russia, India, China and South Africa (BRICS)* by Wim Naudé, Adam Szirmai and Nobuya Haraguchi

2016-17 *Technological Innovation Systems and the wider context: A framework for developing countries* by Hans-Erik Edsand

2016-18 *Migration, occupation and education: Evidence from Ghana* by Clotilde Mahé and Wim Naudé

2016-19 *The impact of ex-ante subsidies to researchers on researcher's productivity: Evidence from a developing country* by Diego Aboal and Ezequiel Tacsir

2016-20 *Multinational enterprises and economic development in host countries: What we know and what we don't know* by Rajneesh Narula and André Pineli

2016-21 *International standards certification, institutional voids and exports from developing country firms* by Micheline Goedhuys and Leo Sleuwaegen

2016-22 *Public policy and mental health: What we can learn from the HIV movement* by David Scheerer, Zina Nimeh and Stefan Weinmann

2016-23 *A new indicator for innovation clusters* by George Christopoulos and Rene Wintjes

2016-24 *Including excluded groups: The slow racial transformation of the South African university system* by Helena Barnard, Robin Cowan, Alan Kirman and Moritz Müller

2016-25 *Fading hope and the rise in inequality in the United States* by Jo Ritzen and Klaus F. Zimmermann

2016-26 *Globalisation, technology and the labour market: A microeconometric analysis for Turkey* by Elena Meschi, Erol Taymaz and Marco Vivarelli

2016-27 *The affordability of the Sustainable Development Goals: A myth or reality?* By Patima Chongcharoentanawat, Kaleab Kebede Haile, Bart Kleine Deters, Tamara Antoinette Kool and Victor Osei Kwadwo

2016-28 *Mimetic behaviour and institutional persistence: a two-armed bandit experiment* by Stefania Innocenti and Robin Cowan