

Over-confidence and Low-cost Heuristics: An Experimental Investigation of Choice Behavior

Simona Cicognani and Luigi Mittone

Abstract

This paper extends choice theory by allowing for the interaction between cognitive costs and imitative dynamics. The authors experimentally investigate the role of imitation when participants face a task which is costly in cognitive terms. In order to disentangle different choice dynamics, they devise a laboratory experiment with a novel experimental task in which they model the choice of different alternatives through high or low cognitive costs and feedback information provided to subjects. Their results provide evidence for imitative behavior driven by under-confidence about own skills and by the beliefs regarding others' performance. They also find a temporal pattern in the distribution of choices, both in the high-cost and low-cost cognitive conditions, that may represent another cognitive shortcut.

(Published in Special Issue *Economics, Psychology and Choice Theory*)

JEL C91 D80 D81 D83

Keywords Belief; cognitive costs; decision making; laboratory experiments

Authors

Simona Cicognani, ✉ School of Social Sciences, University of Trento, Italy,
simona.cicognani@unitn.it

Luigi Mittone, Cognitive and Experimental Economics Laboratory, University of Trento, Italy

Citation Simona Cicognani and Luigi Mittone (2015). Over-confidence and Low-cost Heuristics: An Experimental Investigation of Choice Behavior. *Economics: The Open-Access, Open-Assessment E-Journal*, 8 (2014-36): 1—35. <http://dx.doi.org/10.5018/economics-ejournal.ja.2014-36>

1 Introduction

In every day life decision makers frequently face cognitive costs when choosing among different options. Traditional economic theories of rational behavior disregard cognitive costs and assume that economic agents process costly information fairly easily, since they are always able to select the utility-maximizing option among different ones. In contrast, evidence gathered by some psychologists and economists supports the idea that decision makers systematically violate the assumptions of rational choice theory (Tversky and Kahneman, 1974; Camerer, 2003; Bicchieri, 2006). In particular, the axiom of complete preferences, which entails that an individual is able to compare and express a preference relation between any two objects, requires conditions such as extraordinary computational and cognitive skills that are rarely met in practice. It follows that, in many circumstances, economic actors are boundedly rational and make use of simplifying heuristics, either conscious or unconscious, when they process information that carries cognitive costs (Gigerenzer and Gaissmaier, 2011).

In this paper we investigate a potential cognitive shortcut faced by decision makers: the influence of other agents in their reference community. A growing strand of literature argues that many decisions are affected by social interactions, also known as “peer effects” (Bernheim, 1994; Glaeser and Scheinkman, 2001). Peer effects represent the channel through which individuals belonging to a reference group influence the behavior of individuals in the same reference group, and they can be conceived as “an average intra-group externality that affects identically all the members of a given group” (Calvo-Armengol et al., 2009, p. 1239). The perspective from which we consider peer effects here is that according to which the behavior of other individuals in a reference community may represent a simplifying heuristic, that allows to economize on decision costs, especially when choices carry high cognitive costs. A potential channel through which this heuristic operates is the belief system about the quality of others’ choices.

Economists’ interest in the connection between social interactions and cognitive costs is very limited and relatively recent. An attempt in this direction is represented by the theoretical axiomatization of Hayakawa (2000), who supports the idea of social capital as a source of low-cost heuristics: choice decisions made by an individual depend crucially on the organization of her perceptions of available

choices. To the extent that such perceptions are affected by social elements, they cannot be considered independent of the particular social environment where decisions take place, and, in such a context, bounded rationality is a tool for the formation of endogenous preferences.

Along these lines, Earl and Potts (2004), in investigating the social side of choice in complicated situations make a sharp distinction between high-level preferences and low-level ones. While the former relate to innate behavior, the latter are the result of a learning and specialization process, and they can be acquired in the market for preferences, by simply using and buying other people's learning. Economic agents are indeed cognitively normal individuals who may access and acquire specialized consumer knowledge, regarded as recommendation of someone having relevant expertise on some specific problems. This is what happens in highly uncertain speculative markets, as highlighted by Earl et al. (2007), in which economic agents make use of decision rules picked up from others in an attempt of interpreting and gathering new relevant information. Heuristics are hence socially transmitted.

Conversely, very large and heterogeneous is the extent of contributions on peer effects *per se*. The literature in this field aims at measuring and disentangling endogenous effects and at solving the identification problem: individuals in a group influence each other and it is not easy to state who influences whom.¹ While this represents quite a complex problem to overcome empirically, through field data, the experimental research allows to avoid it, due to a complete control of the information provided to experimental subjects and to the absence of self-selection of subjects into groups. More specifically, in a laboratory setting the experimenter is able to control for different information scenarios, and hence to regulate the direction of influence of knowledge.

A significant amount of experimental work has modeled peer effects by conveying information to participants about the behavior of experimental subjects in previous sessions. Fortin et al. (2007) investigate a tax evasion scenario and divide participants into groups: at the beginning of each period subjects receive feedback about the amount evaded and the number of evaders among the members of their group in the previous period. Beugnot et al. (2013) devise both a

¹ This issue represents the "reflection problem" introduced by Manski (1993).

recursive and a simultaneous source of influence from others' behavior. In the former, each participant has one or two peers (belonging to the previous baseline treatment) and receives information at the beginning of the session about peers' personal information, average piece-rate and performance in the task;² in the latter, interactions occur in real time among players, in order to recreate real-like working situations. Overall, individuals are found to be significantly influenced by others' performance.

A more sophisticated network structure, which also entails network formation during the experiment, is proposed by Conte et al. (2009). The authors investigate which are the factors underlying network formation in an experiment in which, over a minimum of 15 rounds, connections between two players can be set up during the game if both players mutually agree. Network formation involves positive externalities to the extent that both direct and indirect links accrue benefits, but direct links involve costs. Given these payoff rules, players aim at maximizing their own earnings by forming indirect and direct links. They find empirical ground for best-response behavior, besides reciprocator and opportunistic patterns of playing the network formation game.

Despite these network studies, existing experimental literature does not seem to examine and disentangle the underlying factors that drive individuals to conform to others' choices, which is the focus of this work.

A potential motivation for being influenced by others' behavior may act through the beliefs regarding others' outcome: we may imitate those individuals that we consider as successful. This is what Apesteguia et al. (2007) and Offerman and Sonnemans (1998) indicate among their results and it is also investigated in our experiment through an incentivized elicitation of beliefs regarding others' performance. Apesteguia et al. (2007) in a symmetric Cournot game distinguish between participants and roles played by participants: participants can play in three roles and are not told with whom they are matched within groups of nine subjects each. The authors find evidence in favor of imitation of more successful individuals: the degree of imitation they find is increasing in the difference between own and others' payoff. Moreover, higher imitating patterns arise with respect to the individuals with whom one interacts, rather than with respect to individuals who

² Participants must perform a multiplication task without calculators and pens.

play in the same role but in different groups. A similar result is found by Offerman and Sonnemans (1998), with an experiment about investment decisions that depend on different states of the world. Rather than imitation of choices, they test for the presence of imitation of incentivized beliefs: least-performing individuals are found to imitate more the judgements and beliefs of more successful subjects.

Our study is also related to the theoretical and applied literature on crowd-following or herd behavior, which, starting from the seminal contribution of Banerjee (1992), considers economic agents who have incomplete information about the quality of a good or service and hence let the choices or evaluations made by others to guide their own actions.³

This paper builds on contributions such as Fortin et al. (2007) for the way in which information about others' behavior is put forward in the experiment: we provide subjects with information about the choice of a majority of participants in a previous experimental session. Our departure from existing research is twofold. First, we incorporate cognitive costs in an experimental setting of imitation choices. To the best of our knowledge, this is the first study that investigates experimentally social interactions from a cognitive perspective: we estimate the impact of the reliance on choices made by a majority of subjects in two scenarios that differ in terms of cognitive costs involved in the experimental task.

Second, we implement an original experimental task whose solution is not the result of computations or counting,⁴ but which entails a more complex kind of reasoning that is more similar to modern and dynamic decisional contexts. We show participants a set of abstract figures, and we model the difficulty of such figures into high or low. In these two different scenarios, we test the impact of information available to participants: whether they know or not the choices of a majority of subjects who took part in a previous experimental session. We do not find strong evidence in favor of imitative behavior, but on average participants deem that the majority choice is a winning option when the task to perform presents higher cognitive costs. Therefore, the imitative component is driven by the beliefs about others' performance. We also study the extent to which participants under or over-estimate their skills in completing the experiment, and our results show that

³ See, for an example, Çelen and Kariv (2004) and Sorensen (2007).

⁴ For instance, see Pokorny (2008).

those who rely on a previous majority choice or on a default option are significantly under-confident, whereas the majority of participants is clearly over-confident, especially when the task to be solved is more cognitively demanding.

In addition to investigating the connection between cognitive costs and imitative dynamics, our experimental setting also allows us to examine very precisely the temporal pattern of subjects' decisions. This enables us to provide evidence for an additional heuristic related to the timing of decision making, which lends support to the saliency of more recent memories in scenarios cognitively demanding.

The remainder of the paper is organized as follows: Section 2 outlines the experimental design, the procedures followed during the experiment and the behavioral hypotheses; Section 3 presents the data and results from the experiment; Section 4 discusses and draws conclusions.

2 Design, procedures and behavioral predictions

2.1 Design

The novel feature of our experimental design is that participants do not face a typical experimental game, nor a mechanical counting or computational task.⁵ Rather, on their computer screens, they face covered cards which contain on their inner side abstract figures composed of black squares.⁶ Participants' aim is to select the card containing the figure which most closely resembles the figure appearing on the top left of the screen. In doing so, they can either uncover and inspect the cards that appear covered on the screen, or they can choose a covered card representing a default option, which is placed on the bottom right of the screen. The default card is accompanied by a different amount of information according to treatments. Indeed, in the *DEFAULT* treatments, participants are informed that the default card is one of the best 8 cards (among 16) appearing covered on the screen, while in the *MAJORITY* treatments they are told that the default card (in this case called "majority card") has been chosen in a previous experimental session by the majority

⁵ For an overview of real-effort tasks, see Gill and Prowse (2013).

⁶ We are deeply indebted to Paolo Crosetto for the programming of figures through Python.

of participants, and that, moreover, it represents one of the 8 best cards appearing covered on the screen.

The task is implemented both under a high-cost and a low-cost scenario, both considered in cognitive terms: in the high-cost condition, the visual resolution of the figures (number of “pixels”) is higher than in the low-cost condition. The rationale of the task is to offer subjects an environment in which, while searching the solution, they do not have a reference point, either computational or logical: the assignment is not related to an univocally clear solution, characterized by normative determinants. Our approach aims indeed at resembling a dynamic social environment, where available options must be discovered, but it is not straightforward which is the best one among them.

The implementation of the design entails participants to face a screen with one card on the top left (also known as “target card”), and 16 covered cards below, divided into rows of 4 cards each.⁷ Cards are randomly distributed on the screen for each subject and for each treatment, so that the distribution of the 16 cards is different for each participant. Each card carries an abstract figure on its covered side. For the initial 10 seconds all cards are uncovered, and hence the figures portrayed on them are visible. As previously mentioned, the aim of the game is to choose the card which is the most similar to the target card, within a time constraint of 10 minutes. After the 10 initial seconds, all cards get covered: comparisons between any of the 16 cards and the target card can be carried out by participants by clicking on each card which is chosen for the comparison. After clicking on one of the 16 cards, the selected card is shifted next to the target card and both become uncovered. Participants can keep the selected card beside the target card, for a visual comparison, for as much time as they wish, within the 10 minutes of time constraint of the whole experiment. In order to place the card back to its initial position, it is necessary to click again on the card (see Figure 1).

Moreover, after the 10 initial seconds, a default card appears, covered, on the bottom right of the screen, according to the type of treatment.⁸ If one of the 16 cards is inspected and then placed back in its original position among the other 15

⁷ See Appendix A for the instructions provided to participants.

⁸ In the MAJORITY treatments, the default card is called majority card, since it represents the choice of a majority of participants in a previous experimental session.

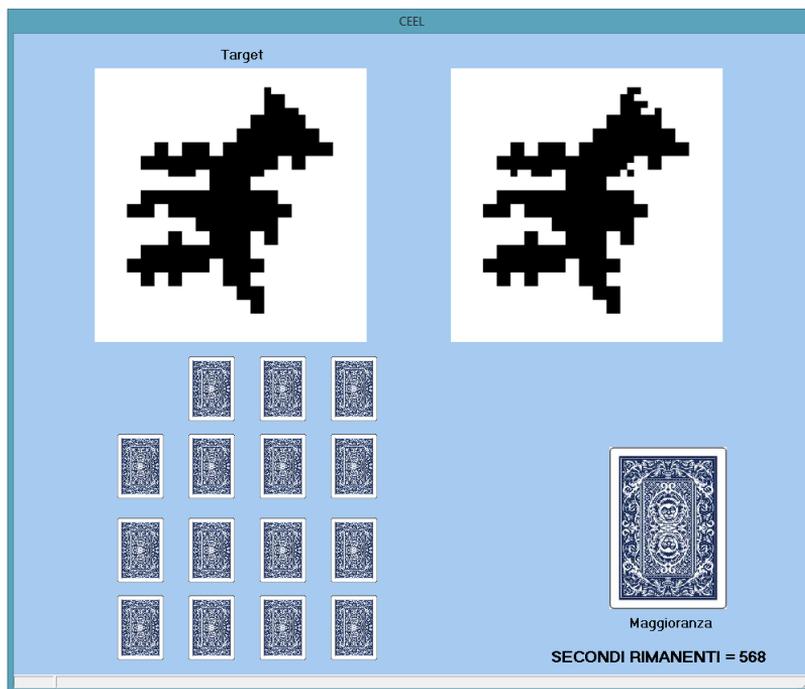


Figure 1: Example of a screenshot during the comparison

Notes: The Figure corresponds to the high-majority treatment.

cards, its boundaries get red. Participants are informed that within the time limit of 10 minutes they have to select either one of the 16 cards or the default/majority card. The 16 cards provide 16 different payoffs, ranging from 0.25 to 4 euros, depending on the fitness of the figure represented on the chosen card with respect to the figure in the target card. For “fitness” we indicate the similarity between the chosen and the target card in terms of black squares constituting the figures portrayed on the cards.

Participants also know that the default/majority card represents one of the 8 best cards in terms of fitness among the 16 cards (hence leading to an earning from 2 to

4 euros). If the default/majority card is selected, it represents the final choice: in this case participants can no longer explore the other cards appearing on the screen and they will receive the outcome attached to the card chosen. Since participants cannot inspect the default/majority card, they can rely only on the information about this card which is provided in the experimental instructions. The rationale for this is to test subjects' beliefs about the goodness of the default/majority card, without subjects being influenced by what they see on this specific card. In case no card is selected within 10 minutes, participants simply receive the show-up fee of 2.5 euros. The countdown of seconds available is displayed at the bottom right of the screen. At the beginning of each session, participants are endowed with a pen and a printed page consisting of the contours of the 16 cards divided into rows of four, as they appear on each computer screen. The idea is to facilitate subjects in remembering the fitness of the cards as they work through them, by allowing participants to take notes related to the cards inspected. An example of the page is provided in Figure 2.

After the completion of the choice, an incentivized belief elicitation task regarding the “fitness” of the default/majority card takes place.

What we mean by “fitness” is the number of black squares composing each figure that differ in terms of position between the selected and the target card. This measure changes across the high-cost and low-cost treatments. In the low-cost scenario, the target figure is made up of 100 black squares. Each of the 16 cards differs from the target not for the total number of squares composing the figure, but for 5 to 20 squares that moved position with respect to the target figure. The idea for the high-cost scenario is that of increasing the number of “pixels” composing the figures, that is the resolution: the target figure in this case consists of 400 black squares and each of the 16 cards differs from the target card for 20 to 80 squares that moved position, while the total amount of squares remains the same. In practice, each black square of a low-cost figure corresponds to 4 black squares in a figure in the high-cost scenario. Therefore, in both scenarios, there isn't the perfect match card among the 16 cards on the screen: the best cards in terms of “fitness” are those with 5 and 20 squares that moved position from the target figure, in the low-cost and high-cost scenarios respectively.

Panel (a) and (b) in Figure 3 provide examples of figures in the low-cost and high-cost treatments, respectively. We expect that the increase in the number of

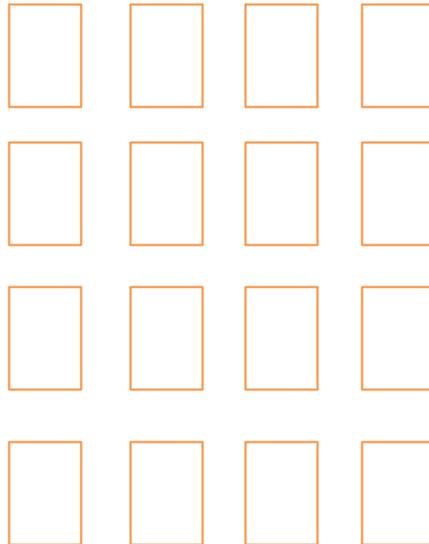


Figure 2: Example of the annotation page

squares composing the figures leads to an increased complexity in the task, and hence to higher cognitive costs.

It follows that in the belief elicitation task participants are asked to choose the number of squares they believe that differ between the target and the default card.⁹ All options, according to the treatment, are presented and participants are asked to tick on one of them in order to end the experiment: in the low-cost treatments, numbers from 5 to 16 were displayed, whereas in the high-cost treatments participants could tick on numbers ranging from 20 to 80, with a distance of 4 to each other (20, 24, 28, 32, and so on, until 80). Hence, subjects were not requested to write the amounts on their own, without any prior information of the possible numbers, but were asked to select one among the numbers displayed.

⁹ Once the experimental task has been completed, we ask subjects the following question “By what amount of squares do you think that the default (or majority) card differs from the target?”

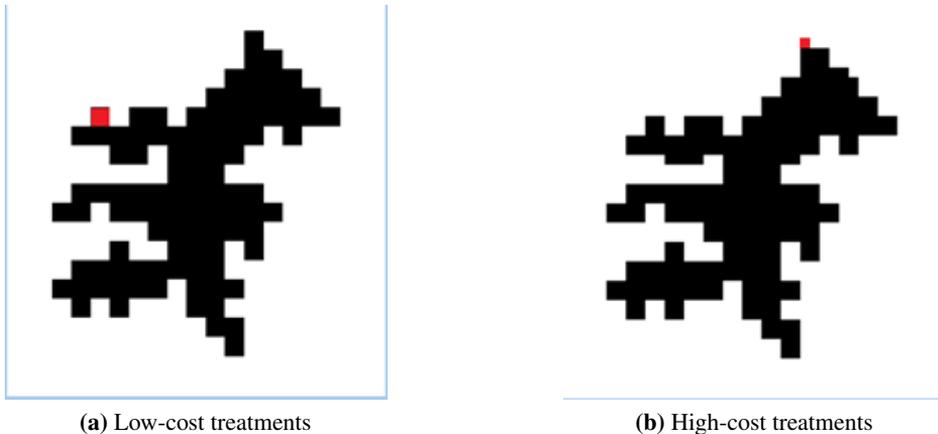


Figure 3: Example of figures

Notes: In red it is represented the dimension of a square.

If the number selected is correct, one additional Euro is earned at the end of the experiment by the participant who made a right guess about the fitness of the default/majority card.

In addition, at the end of each session, participants were asked to fill out an anonymous ex-post questionnaire. This questionnaire aimed at gathering information about the perceived difficulty of the task, impressions about what the figure represented, as well as a series of demographics data such as age, gender and college major (see Appendix B for details).

The experimental design develops across two dimensions of treatments, that allow us to disentangle the role played by the availability of social influence in a high-demanding scenario in cognitive terms versus a low-demanding one. Moreover, the experimental setting also permits us to control for the effect of imitation and the temporal saliency of choices, in conditions of low and high cognitive costs.

A first dimension of the 2 x 2 experimental design is represented by cognitive costs: in the low-cost treatment, the figures are made up of 100 black squares,

		Cognitive costs	
		Low Cost	High Cost
Information on the majority	No	<i>low-default</i>	<i>high-default</i>
	Yes	<i>low-majority</i>	<i>high-majority</i>

Table 1: Treatments

while in the high-cost treatment figures have a higher precision, since they include 400 squares. The second dimension of treatments involves the information given to participants about the default/majority card. In the default treatment, the card on the bottom right is called DEFAULT CARD and participants are informed that it is one of the 8 best cards among the 16, in terms of fitness with the target card. Conversely, in the majority treatment, the card on the bottom right is denominated MAJORITY CARD: participants are notified that it is one of the 8 best cards, and, in addition, that it was chosen by the majority of participants from a previous experimental session. The rationale of providing information also in the default treatment is to preserve comparability between the default and the majority treatment. In this way, the communication of the fitness of the cards, in other words their expected value, remains invariant across treatments. What changes is the additional information of the choice implemented by a previous majority, that in our approach aims at representing a social interaction component which participants may decide to follow when facing a cognitive cost.

Table 1 provides an outline of the four treatments.

2.2 Procedures

A baseline session without the availability of a default/majority card was previously run with 20 participants in order to gather data on the majority card used in the following majority treatments. The low-cost treatment was administered to 10 subjects, and the high-cost treatment was conducted on the remaining 10 subjects.

The best card in terms of fitness with the target card was chosen by the majority of participants in both treatments.

Following this baseline initial session, 154 subjects took part in the experiment, divided into 12 sessions.¹⁰ 40 participants took part in each of the four treatments, except in the low-majority treatment in which the number of participants was 34. The experiment was fully computerized and it was performed through a between-subject design. Participants were volunteer undergraduate and graduate students from different humanities and technical faculties from the University of Trento. Each session lasted around 30 minutes and it was conducted at CEEL (Cognitive and Experimental Economics Laboratory, Trento) between April and June 2013. A show-up fee of 2.5 euros was given. The average earning per participant was of 5.7 euros, and it was paid privately and anonymously in a separate room at the end of the experiment. Subjects could not participate in more than one session.

Upon their arrival at the laboratory, participants were randomly allocated to cubicles inhibiting interaction with other participants. Each participant could read the instructions for the experiment on the computer screen placed in the cubicle, and she was invited to read them privately. Later, a member of staff read the instructions aloud and participants were given the opportunity to privately ask staff members for clarifications. A set of computerized control questions was administered to subjects in all treatments, in order to verify their understanding of the task.

The experiment consisted of four treatments, as depicted in Table 1. Three sessions per treatment were run.

2.3 Behavioral predictions

A first conjecture arises related to the task administered to subjects. We believe that our approach in designing the figures that participants have to visually compare entails a higher cognitive cost in case of a higher amount of resolution of the figure in the card, which consequently involves a lower performance rate. Therefore, we maintain that increasing the cognitive cost implies lower performance choice rates.

¹⁰ Among a total of 166 participants, 154 selected a final choice, while 12 run out of time. For the analysis, we decided to consider only data about the 154 subjects who finalized a choice.

In this study we are interested in analyzing what kind of behavioral shortcuts do individuals put into practice when they face decisions which are cognitively demanding and do not present a clear cut solution.

When facing a new task, which is neither mechanical nor computational, subjects might conform to the norms deriving from the majority of members of their community. According to Bicchieri (2006, p. 5), “to efficiently search our memory and group a new event with previously encountered ones, we use cognitive shortcuts. Cognitive shortcuts play a crucial role in categorization and the subsequent activation of scripts and schemata(..) In the *heuristic* route, behavior is guided by *default rules* stored in memory that are cued by contextual stimuli. Norms are one class of default rules.” Referring to what has already been chosen by other subjects in the same decisional environment may represent a channel through which it is possible to save on cognitive costs. This type of social shortcut may be increasingly relevant when decisions are more demanding in cognitive terms. An underlying behavioral pattern could be driven by the beliefs regarding the behavior of others, and especially of the majority that one decides to conform to: this majority can be considered as acting better than the individual subjects, and hence could be chosen to be a reference point in decisional dynamics.

Moreover, another line of argument refers to the different timing of experiences faced by subjects. Indeed, cognitive processing could be related to memory capacity, so that more recent events become more salient when subjects experience cognitively demanding tasks (see, for instance, Hastie and Dawes (2010)). When facing alternatives, decision makers experience utilities, and they attach such utilities to the alternatives present in their memory. Hence, preferences arise from memory representations (Weber and Johnson, 2006). In this context, a central role is played by the timing of the experience, which leads to different memories of experiences (Miron-Shatz et al., 2009): preferences depend on the memory of past experiences. When confronted with cognitive costs, we expect that individuals attach more relevance to memories of experiences occurred later in time.

This enables us to formulate the following two testable predictions:

Hypothesis 1 People imitate more in contexts highly demanding in cognitive terms.

Hypothesis 2 More recent memories have a greater impact on choices.

3 Results

In this section, we first present descriptive statistics of the sample of participants, with bivariate tests. We then discuss econometric results related to the research hypotheses we aim to test.

3.1 Descriptive statistics and *prima facie* evidence

From the pages for notes collected after each experimental session, it is clear that for most participants they were useful tools for recording their thoughts as they inspected the cards, and hence remembering the goodness of the cards visualized. Overall, 27% of such pages were left blank, 12% presented draws or scrawls, while the remaining 61% exhibited some measures of quality about the cards, as for instance arrows, percentages, pluses and minuses or different numbers of stars written within the contours of some cards. Across treatments we do not observe statistical differences in the types of writings left on the pages for notes, or in the amount of cards left blank.

Table 2 reports means, standard deviations, minimum and maximum values of the main variables used in the analysis.

As can be noted from Table 2, females slightly outnumber males. The variable “payoff” is the total payoff from the experiment for each subject, while for the variable “net payoff” it is subtracted the additional Euro gained in case of a correct guess of the belief about the default/majority card (out of 154 participants, 45 guessed the belief correctly).

On average, subjects complete the task in almost 7 minutes and they uncover 22 cards. The average difficulty declared at the end of the session is approximately 5, on a scale from 1 to 10, ranging from easy to difficult. The variable “belief default” is the ratio between the real difference in black squares between the target and the default card, and the believed corresponding difference in black squares declared at the end of each session by subjects. Since in both low-cost and high-cost treatments the default/majority card is the one with the smallest amount of different squares with respect to the target card (5 over 20 in the low-cost treatment vs 20 over 80 in the high-cost treatment), it follows that the “belief default” values range between 0.25 and 1, in case of the worst and best belief about the default

	Mean	Std. Dev.	Min	Max
Sex (male = 1)	0.46	0.50	0	1
Age	23.42	3.29	19	41
Payoff	5.71	1.00	2.75	7.5
Net payoff	5.42	0.95	2.75	6.5
Seconds	417.23	162.96	0.001	581.67
N. uncovered cards	21.80	8.74	1	49
Difficulty declared	5.38	2.23	1	10
Belief default	0.72	0.23	0.25	1

Table 2: Descriptive statistics

Notes: Difficulty is expressed on a scale from 1 (easy) to 10 (difficult); Belief default is 0.25 if participants believe that the default/majority card is the worst, in terms of similarity with the target card, and it is equal to 1 in case they believe it is the most similar card with respect to the target.

card, respectively. Therefore, a higher value for “belief default” represents higher expectations about the goodness of fit of the default/majority card. In our data set, participants’ average belief falls almost in between the worst and the best belief.

Taking into account these variables across treatments, we can see a more detailed picture of the underlying patterns of the experiment in Table 3.

We find a statistically significant lower payoff in the high-cost treatments compared to the low-cost treatments, both in case of the payoff and net payoff variables (Wilcoxon rank-sum test, $z = 5.492$, $p = 0.00$; $z = 6.859$, $p = 0.00$, respectively¹¹). This is corroborated by looking at Figure 4, in which we can notice that the median payoff values for the high-cost treatments are always below the median payoff values for the low-cost treatments.

This result supports our experimental approach and sheds light in favor of the presence of a higher cognitive difficulty in case of figures with higher resolution, since the net payoff is a clear measure of performance. This preliminary evidence

¹¹ Throughout the paper, treatment means are used for statistical two tailed tests, unless otherwise specified.

	Low-default	High-default	Low-majority	High-majority
Payoff	6.12	5.09	6.14	5.56
Net payoff	5.87	4.84	5.93	5.11
Milliseconds	453786.1	411323.2	405916.3	396203.5
N. uncovered cards	22.75	22.53	20.26	21.43
Difficulty declared	5.33	5.83	5.59	4.8
Belief default	.6967347	.7139879	.6493709	.7947086

Table 3: Means across treatments

Notes: Difficulty is expressed on a scale from 1 (easy) to 10 (difficult); Belief default is 0.25 if participants believe that the default/majority card is the worst, in terms of similarity with the target card, and it is equal to 1 in case they believe it is the most similar card with respect to the target.

is verified in the ordinary least-squares regression analysis which investigates the determinants of the performance reached by the 154 participants in the experiment (see Table 4).

The dependent variable in the model – the payoff in the experiment– is regressed on the following explanatory variables: *info* is a dummy variable equal to 1 if the treatment provides information about the majority (low-majority and high-majority treatments), and it is equal to 0 otherwise (low-default and high-default treatments); *cost* is equal to 1 if the treatment is highly costly in cognitive terms, and it is equal to 0 otherwise; *male* is equal to 1 if the participant is a male, and to 0 in case of females; *difficulty* is the perceived difficulty declared by subjects in a post-experiment questionnaire (it ranges from 1, easy, to 10, difficult); *seconds* is the amount of milliseconds employed for the completion of the task.

The regression output reported in Table 4 confirms what previously discussed, as can be noticed from the negative and highly significant coefficient of *cost*. The positive and statistically significant coefficient of the dummy information might be due to the higher frequency of default choices in the majority treatments: the default choice was attached to the highest payoff, even if participants were not aware of this. Males, *ceteris paribus*, seem to perform slightly better than females.

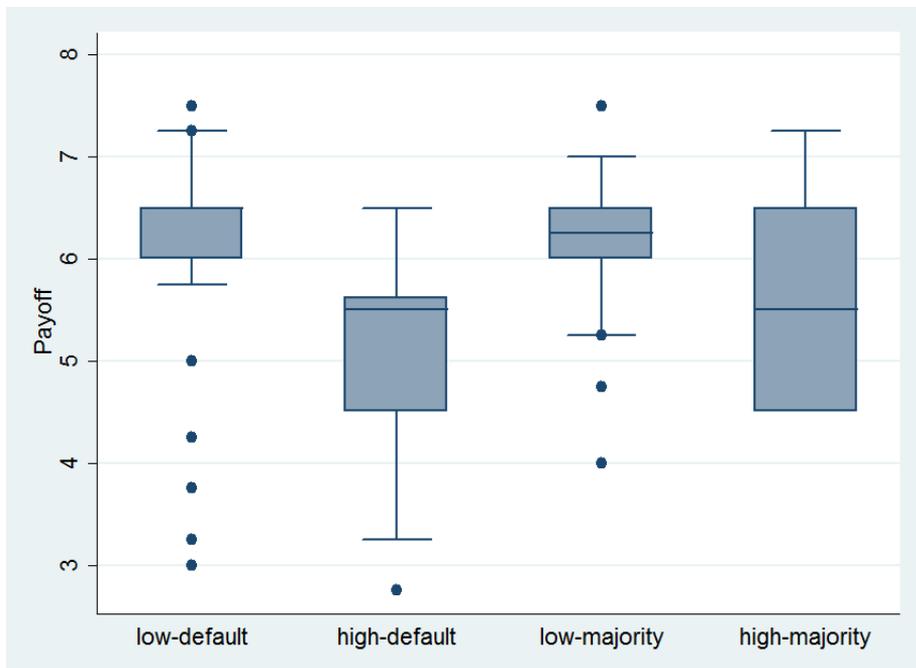


Figure 4: Payoff per each treatment

Notes: Middle bars correspond to median values, the edges of the boxes comprise the inter-quartile range between the 25th and the 75th percentiles, whiskers correspond to 1.5 times this range and circles characterize any other observation.

<i>Payoff</i> ~	Coeff.	Std. Errors
<i>(Intercept)</i>	5.776173	(.2897312)***
<i>info</i>	.2867096	(.1468653) [°]
<i>cost</i>	-.7882182	(.145829)***
<i>male</i>	.2504396	(.146718) [°]
<i>difficulty</i>	-.0510903	(.0330326)
<i>seconds</i>	8.73e-07	(4.55e-07) [°]

Number of Observations: 154

Significance levels: *** 0.001; ** 0.01; * 0.05; [°] 0.1

Table 4: Regression model of the determinants of performance

Notes: The dependent variable in the model is the payoff in the experiment; *info* is a dummy variable equal to 1 if the treatment provides information about the majority, and it is equal to 0 otherwise; *cost* is equal to 1 if the treatment is highly costly in cognitive terms, and it is equal to 0 otherwise; *male* is equal to 1 if the participant is a male, and is equal to 0 in case of females; *difficulty* is the perceived difficulty declared by subjects in a questionnaire at the end of the experiment (it ranges from 1, easy, to 10, difficult); *seconds* is the amount of milliseconds employed for the completion of the task.

Furthermore, the perceived difficulty declared at the end of the experiment does not seem to predict the performance of participants, while the more time spent completing the task is significantly associated to a higher payoff. This evidence leads to the following result, which confirms the validity of our experimental task:

Result 1 *Figures with a higher resolution are associated to higher cognitive costs and hence lower performance rates.*

Following the descriptive statistics of Table 3, in the majority treatments participants employ the smallest amount of seconds, especially in the high-majority one. Moreover, participants finalize their choice faster if the task involves high cognitive costs, seeming hence to give up before if they face a difficult task.

Conversely, the maximum number of uncovered cards is registered in the low-default treatment. This might be due to the low difficulty of the task, so that

	Low-default	High-default	Low-majority	High-majority
Frequencies default	1	2	2	3
N. of subjects	40	40	34	40

Table 5: Frequencies of default options

participants are more willing to inspect the 16 cards on the screen. Nevertheless, and quite surprisingly, the lowest difficulty declared for the task is related to the high-majority treatment. Concerning the beliefs about the goodness of fit of the default/majority card, they improve in the case of the high-cost treatments with respect to the low-cost ones, while the information setting does not seem to influence subjects' beliefs (Wilcoxon rank-sum test, $z = -2.261$, $p = 0.0238$; $z = -0.618$, $p = 0.5369$, respectively). This enables us to derive the following result:

Result 2 *High-cost treatments present higher beliefs about the default/majority card.*

3.2 Imitative dynamics and over-confidence

Table 5 reports the frequencies of subjects who chose either the default or the majority card.

As we can note, the frequencies of subjects choosing the default/majority option are very low: out of 154 individuals, only 8 adopted this decisional shortcut. The highest frequency of default choices is in line with our imitation hypothesis: people decide for a shortcut in case the task is more costly in cognitive terms and a majority outside option is available, but differences across treatments are not statistically significant (two-sample test of proportions with respect to cost and information treatments, respectively: $z = -0.6135$, $p = 0.5395$; $z = -0.8400$, $p = 0.4009$).

We do not find evidence for imitative behavior in choices, but we consider to investigate more deeply the beliefs about the default/majority cards in order to have a better understanding of imitative dynamics. A closer examination of the elicited beliefs across treatments requires us to undertake pairwise comparisons.

If we compare average beliefs between low-default and high-default treatments, we do not find any statistically significant difference (Wilcoxon rank-sum test, $z = -0.483$, $p = 0.6290$). The same result holds when we compare average beliefs between the low-default and the low-majority treatments (Wilcoxon rank-sum test, $z = 1.030$, $p = 0.3029$). A different pattern arises concerning beliefs in the high-majority treatment. Participants in this condition have higher beliefs related to the default card both with respect to the low-majority treatment and the high-default treatment (Wilcoxon rank-sum test, $z = -2.636$, $p = 0.0084$; $z = -1.748$, $p = 0.0804$, respectively). This suggests that a majority, considered as social component in our experiment, is considered to implement better decisions in presence of high cognitive costs. If we investigate participants' beliefs, we can argue that the decisions of the majority represent a focal point in the presence of higher cognitive costs. Hence, we can conclude that:

Result 3 *Beliefs about the choice of a majority of individuals are better with respect to a setting with lower cognitive costs and with respect to a setting in which there is no information about the choice of a majority.*

The higher percentage of participants who believe that the majority card should represent a good choice when the cognitive cost is higher might be supported by two possible explanations. First, people might think that the others are more cognitively efficient than themselves. Second, people could prefer to choose a conformist alternative when the cognitive difficulty is high in order to avoid regret. Since from the data it arises a gap between beliefs and choices, the second explanation seems to be weaker than the first one.

Imitative dynamics are also investigated in connection with under or over-confidence of participants regarding the experimental task. It might indeed be the case that subjects decide to conform to what a majority has already decided because they feel under-confident about their skills in solving the task.

In the ex-post questionnaire there are a few questions addressed to participants that can shed light on this issue.¹² Besides the perceived difficulty of the task (on a scale of 1–10), the questionnaire encompasses a question regarding the perceived accuracy in solving the task, since participants had to indicate the number of

¹² Available in Appendix B. Questions relevant for this subsection are 1)–3).

squares of difference they believe exist between their chosen card and the target card. Moreover, subjects were also asked if they think of having performed better/worse/ or equally with the respect to the majority of other participants.

The difficulty declared about the task is on average 6.75 for those who opted for the default/majority card, whereas it is on average 5.3 for those who chose as final card one inspected by themselves. This is a first clue about referring to the default/majority choice when the task to be completed is perceived as complicated, although the Wilcoxon rank-sum test for the comparison of average difficulty between the two samples is not statistically significant ($z = -1.438$, $p = 0.1506$).

We construct a variable representing under/over-confidence by subtracting from the believed accuracy, resulting from the post-experimental questionnaire, the actual accuracy of participant's final choice. Both accuracy variables are normalized to a 1–16 scale, where 1 indicates the best figure among the 16 on the screen, whereas 16 the worst one. Therefore, the variable takes negative values when subjects are over-confident and hence think of having performed better than they actually did, while it takes positive values when a subject is under-confident (a null value characterizes the exact confidence case). Overall, on average, experimental subjects are over-confident about their skills in the task (the under/over-confidence variable mean value is -2.57). When distinguishing with respect to the difficulty of the task, we find that the over-confidence variable is $-.364$ in the low-cost treatments and -4.625 in the high-cost treatments (the differences in this case are statistically significant: Wilcoxon rank-sum test, $z = 6.477$, $p = 0.0000$). There are no differences in terms of over-confidence when information treatments are compared (value of the over-confidence variable equal to -2.81 and -2.32 in the default treatments and majority treatments, respectively). When comparing subjects who chose the default/majority card and subjects who did not, we find a statistically significant difference in the over-confidence variable: more specifically the former subjects under-evaluate their skills by 1, while the latter over-evaluate their skills on average by 2.77 (Wilcoxon rank-sum test, $z = -2.828$, $p = 0.0047$). That is to say that participants are more over-confident about their capacities when the task is more difficult, and, as expected, are more overconfident when they don't rely on the default/majority card as their final choice. This result is in line with the values declared about the difficulty of the task. With

respect to gender in connection to under/over-confidence, no significant differences result from the dataset. The analysis above leads to the following result:

Result 4 *In the high-cost treatments participants result being more over-confident about their skills. The same applies to participants who don't rely on a default/majority card as their final choice.*

3.3 Temporal patterns of decision

As regards the time distribution of choices, our data set enables us to track the sequence of cards uncovered by subjects, as well as the amount of seconds they spend looking at each card they inspect. On average, subjects uncovered 22 cards each: this means, since the available cards for inspection on their screen are 16, that they often uncover the same card more than once. Therefore, whenever this was the case, we decided to consider as timing of choice the last moment the chosen card was uncovered. For instance, if a subject uncovers the chosen card as the third card in his sequence of uncovered cards, but also as tenth card, we consider that the subject selects the tenth uncovered card as her final choice. Of course, this criterion pushes the distribution of choices towards the right tail, but we believe it is more realistic than other criteria (for example, considering the first time in which the chosen card was uncovered or a point in between the first and the last time in which it was uncovered). Given this assumption, we construct the variable “timing of choice” by normalizing the order in which the chosen card was uncovered by the total number of uncovered cards for each subject. Panel (a) and (b) in Figure 5 display the frequencies for the timing of choices of participants in the high-cost and low-cost treatments, respectively.

From a first visual comparison, we can note that in the high-cost treatment there is a higher frequency of participants choosing among the recently uncovered cards. This is confirmed, although by a marginal significance, when we perform a Wilcoxon rank-sum test on the average timing of decision, comparing low-cost and high-cost scenarios ($z = -1.662$, $p = 0.0965$). When looking at the patterns in which the cards were uncovered and then chosen, we notice that all subjects, except one, uncovered the best card among the 16 present in the screen. That is to say, virtually all participants knew where the best card was. We hence analyze

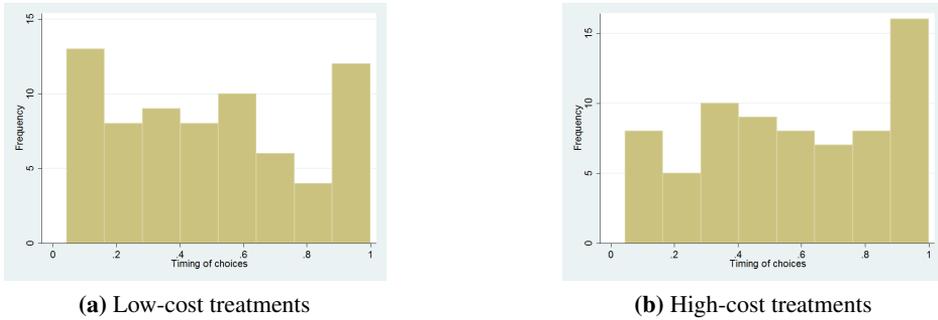


Figure 5: Time distribution of choices

the distance, in terms of number of cards uncovered, between the the best card and the card chosen and we normalize it for the total amount of uncovered cards by each subject. We compare this distance across the high-cost and the low-cost treatments and we find that it is significantly higher in case of high-cost treatments (Wilcoxon rank-sum test, $z = -2.437$, $p = 0.0148$). Also this finding corroborates our Hypothesis 2:

Result 5 *More recent memories have a greater impact on choices, especially in scenarios with higher cognitive costs.*

4 Discussion and Conclusion

Cognitive costs characterize many decisions we frequently face in real life situations. This paper investigates the role played by cognitive costs in contexts involving social interactions and sheds more light on the determinants of imitative behavior. Imitation is a very frequent pattern of real world behaviors but its sources have been limitedly examined by economic theory. Imitation is crucial in the transmission of knowledge and represents one of the main sources of learning.

We analyze the role of imitation in decision making through an experiment on choice behavior, in which we disentangle the potential drivers of imitative dynamics

through the modeling of different levels of cognitive costs. For this purpose, we expressly devise a novel experimental task which is not purely mechanical like for instance counting tasks. Our results show that the experimental task implemented allows us to carry out a scenario with low cognitive costs and one with high cognitive costs, with the latter being characterized by lower performance levels. Moreover, we analyze temporal decisional patterns, since these might uncover relevant heuristics in decision making, especially in scenarios with high cognitive costs.

In our experiment, imitation, conceived as a cognitive shortcut, does not seem to be present in participants' choices, but we find higher expectations concerning the goodness of choice of a majority of subjects from a previous experimental session, if the task to be performed carries high cognitive costs. There is therefore some evidence of imitation of a majority as concerns beliefs. Moreover, we find that participants are statistically more over-confident in the high-cost treatments, and when they do not rely on choosing the default/majority card as their final decision. Therefore, choosing the short-cut of the default/majority card is correlated with under-confidence about own skills in solving the experimental task.

A possible explanation of this finding could be identified in the way in which the imitative component has been introduced in the experiment (i.e., by providing participants information about what a majority did in a previous experimental session): this indeed might have been too weak to generate some imitative behavior. For this reason, further research could be devoted towards an endogenous imitation pattern within the experiment: the majority could be formed during the experiment and feedback about it would be provided during the game. Another stream for future research might be represented by modeling differently the social component, and by strengthening it. This would be achievable by incorporating an in-group/out-group framework in the experimental design.

Moreover, from the heuristics and biases literature introduced by the seminal work of Tversky and Kahneman (1974), we know that people tend to presume they outperform the average, as it is the case of participants in our experiment.¹³

¹³ To the question "Do you think you chose a final card which is better, equal or worse than the card chosen by the majority?", included in the ex-post questionnaire, out of 154 individuals 71 replied they performed better than the majority, 71 equal and just 12 worse, thus confirming the overconfident pattern first proposed by Kahneman.

With this in mind, following what a majority previously did, without knowing the details of what kind of individuals the majority consisted of, might have been too risky for such kind of overconfident agents. In the field of consumer research, the source credibility literature¹⁴ highlights the importance of the origin of an advertising message for strengthening the validity of the message sent. This is the case for instance of celebrity endorsements, in which the celebrity does not sponsor questionable items because of possible negative reputation effects. It is hence conceivable to design future follow-up studies of the current one with different information, across the treatments, about the composition of who has previously chosen the majority card. Telling this kind of information to participants could help increasing the attractiveness of this card and shed more light on the distinction between social and personal heuristics.

In this paper, we also provide empirical evidence that in the presence of high cognitive costs subjects implement decisional heuristics in order to find a solution to the task. More specifically, they attach more relevance to more recent memories. The experiment studies the situation of a new decisional process, which does not have room for learning dynamics since it is played only once by subjects. Nevertheless, for future research, it would be interesting to investigate how participants behave when learning dynamics are taken into account, and hence to see whether practice and repetition have an impact on participants' behavior.

We are aware of the concerns related to the external validity of laboratory experiments, and of the limitations of letting individuals interact with reference groups exogenously imposed. Nevertheless, our approach might mimic those situations in which there is a decisional default option such as in Internet contexts, in which frequently the most popular item is signalled and individuals might choose to opt for it instead of individually looking at its characteristics, which may be very costly in cognitive terms, regarding both effort and knowledge required.

Acknowledgements: We are grateful to conference participants at the 2014 AFSE conference in Lyon and at the 2013 Florence Workshop on Behavioral and Experimental

¹⁴ For a review, see Pornpitakpan (2004).

Economics, and to seminar participants at the University of Trento. We also wish to thank Marco Tecilla for programming the experiment. The usual disclaimers apply.

Appendix A: Instructions for the “high cost-default” treatment (translated from Italian)

Welcome to this experiment!¹⁵

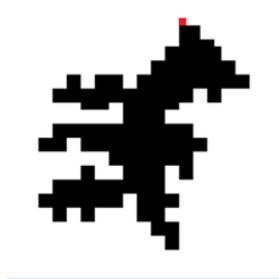
You will receive 2.5 euros as a showing-up fee for coming on time to the experiment. We kindly ask you to read the instructions carefully. During the experiment it is forbidden the communication with other participants. If you have any doubts and you want to ask a question, please raise your hand, an experimenter will come to you and will answer your question. If you don't respect these rules we will have to exclude you from the experiment and you will not be paid.

You are taking part to an economic experiment on decision-making. You can earn an amount of euros depending upon your decisions during the experiment. Both your choices and other participants' choices will remain anonymous and will never be attached to your name.

Initially, for 15 seconds, it will appear on your screen:

- a card, on the top left. This card is called “TARGET” card. It contains a figure composed of many squares. As a square we mean the part represented in red in the figure below:

¹⁵ Instructions for the other treatments are available upon request. In italics are the parts related to the majority treatments, for which DEFAULT card should always be replaced by textscMAJORITY card.



- below the TARGET card you will find 16 cards containing figures. Each figure differs from the figure contained in the TARGET card according to a different amount of squares.

Your aim is to find the card which contains the figure more similar to the figure contained in the TARGET card, **that is with the smallest amount of squares placed in a different position with respect to the figure in the TARGET card.** Your final earning will be as greater as smaller is the amount of squares that differ from the figure of the card you chose and the figure on the TARGET card.

Your earning will be equal to the show-up fee plus:

- 4 euros, in the case in which the card you chose is the best one, that is the most similar to the TARGET figure;
- 0.25 euros, in the case in which you chose is the worst one, that is the least similar to the TARGET figure;
- the intermediate cards, between the best and the worst one, will provide you intermediate earnings between 0.25 and 4 euros, with increases of 0.25 per card.

At the end of the initial 15 seconds, during which both the TARGET card and the 16 cards underneath are visible, all cards will be covered. On the bottom right of the screen will appear a covered card, called DEFAULT card. This card

represents one of the best 8 cards (more similar to the TARGET card among the 16 you find covered). *[Majority treatments only] On the bottom right of the screen will appear a covered card, called MAJORITY card. This card has been chosen by the majority of participants in a preceding session of the experiment and moreover it represents one of the best 8 cards (more similar to the TARGET card among the 16 you find covered. You can choose whether to select this card as your final choice, or to uncover and inspect the other 16 cards that appear on the left of the DEFAULT card.*

In case you choose the DEFAULT card, this represents your conclusive choice and you will not be able to uncover it and then inspect the other 16 cards. Your earning in case you choose the DEFAULT card depends on the similarity of the DEFAULT card with respect to the TARGET card. About the DEFAULT card you are only informed that it is one of the best 8 cards among the 16 present on its left. *[Majority treatments only] About the MAJORITY card you are only informed that it has been chosen by the majority of participants in a preceding session of the experiment and moreover that it represents one of the best 8 cards among the 16 you find covered. **If you choose it, this represents your conclusive choice.***

If you choose to inspect one of the 16 cards, you have to click on it. The card will be enlarged and displayed beside the TARGET card. Once you decide to conclude the visual comparison between the selected card and the TARGET card, in order to bring back to the initial position the selected card you must re-click on it. In this way, the card you have just compared will return to its initial position and and its border will become red, to indicate that the card has already been inspected.

You can decide whether to choose this card as you definitive card, and in this case you must re-click on it, or to proceed with the visualization of another card. In case you want to visualize again a card which has already been inspected (therefore with the red border), you can re-click on it. The computer will ask you whether you want to select this card as your final choice or if you simply want to visualize it again.

If you choose to inspect one or more of the 16 cards, at any moment you will be able to interrupt your search and to choose the DEFAULT card. In this case the experiment ends and your earning will depend on the similarity between the DEFAULT card and the TARGET card.

The time available to you is 10 minutes, starting after the reading of these instructions. **When the time expires, if you haven't selected any card yet, your final earning will be equal to the 2.5 euros of the show-up fee.**

Once you have selected your conclusive card, we ask you to remain seated and silent, while waiting for the other participants to complete the experiment. After you made your choice (either one of the 16 cards or the DEFAULT card), the computer will ask you to specify the amount of squares for which you think the DEFAULT card) (that you cannot visualize) differs with respect to the TARGET card. In case your answer is correct, you will receive an additional 1 Euro to your final earning.

Your final earning will be paid to you in cash and privately, so that the other participants to the experiment will not know your earning.

Appendix B: Ex-post questionnaire

1) How accurate do you think your final choice was? Indicate the number of squares for which you think your final card differed compared to the TARGET card. ¹⁶

(5/ 6/ 7/ 8/ 9/ 10/ 11/ 12/ 13/ 14/ 15/ 16/ 17/ 18/ 19/ 20)

2) On a scale of 1-10, with 1 being the lowest grade of difficulty and 10 the highest, how difficult did you consider the experimental task?

(1/ 2/ 3/ 4/ 5/ 6/ 7/ 8/ 9/ 10)

3) Do you think you chose a final card which is better, equal or worse than the card chosen by the majority?

(better/ equal/ worse)

4) What do you think the TARGET card represented?

5) Among the cards, do you think there was one that will most probably be chosen by the majority? Why?

(yes/ no)

6) Which faculty you are enrolled for?

(Economics/ Humanities/ Engineering/ Law/ Physics and Mathematics/ Sociology or Psychology/ Other Natural Sciences/ Other Social Sciences/ Not a Student)

7) Which year of studies you are enrolled?

(first year/ second/ third/ out-of-course/ master student/ not a student)

¹⁶ In the high-cost treatments, due to the different resolution of figures, the options available to subjects for this question were: (20/ 24/ 28/ 32/ 36/ 40/ 44/ 48/ 52/ 56/ 60/ 64/ 68/ 72/ 76/ 80). The rest of the questionnaire did not present any differences across treatments.

8) What is your gender?

(female/ male)

9) In what year you were born?

References

- Apestequia, J., Huck, S., and Oechssler, J. (2007). Imitation: Theory and experimental evidence. *Journal of Economic Theory*, 136(1): 217–235. URL <http://www.sciencedirect.com/science/article/pii/S0022053106001268>.
- Banerjee, A. V. (1992). A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3): 797–817. DOI 10.2307/2118364. URL <http://qje.oxfordjournals.org/cgi/doi/10.2307/2118364>.
- Bernheim, B. D. (1994). A theory of conformity. *Journal of Political Economy*, pages 841–877. URL <http://www.jstor.org/discover/10.2307/2138650?uid=3738016&uid=2&uid=4&sid=21103126403221>.
- Beugnot, J., Fortin, B., Lacroix, G., and Villeval, M. C. (2013). Social networks and peer effects at work. IZA Discussion Paper 7521, Institute for the Study of Labor (IZA). URL <http://ideas.repec.org/p/iza/izadps/dp7521.html>.
- Bicchieri, C. (2006). *The grammar of society: The nature and dynamics of social norms*. Cambridge University Press.
- Calvo-Armengol, A., Patacchini, E., and Zenou, Y. (2009). Peer effects and social networks in education. *The Review of Economic Studies*, 76(4): 1239–1267. URL <http://restud.oxfordjournals.org/content/76/4/1239.short>.
- Camerer, C. (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press.
- Çelen, B., and Kariv, S. (2004). Distinguishing informational cascades from herd behavior in the laboratory. *American Economic Review*, 94(3): 484–498. DOI 10.1257/0002828041464461. URL <http://www.aeaweb.org.ezp.biblio.unitn.it/articles.php?doi=10.1257/0002828041464461>.
- Conte, A., Di Cagno, D. T., and Sciubba, E. (2009). Strategies in social network formation. Discussion paper, Jena economic research papers. URL <http://www.econstor.eu/handle/10419/32636>.

- Earl, P. E., Peng, T.-C., and Potts, J. (2007). Decision-rule cascades and the dynamics of speculative bubbles. *Journal of Economic Psychology*, 28(3): 351–364. DOI 10.1016/j.joep.2006.08.001. URL <http://www.sciencedirect.com/science/article/pii/S0167487006000845>.
- Earl, P. E., and Potts, J. (2004). The market for preferences. *Cambridge Journal of Economics*, 28(4): 619–633. URL <http://ideas.repec.org/a/oup/cambje/v28y2004i4p619-633.html>.
- Fortin, B., Lacroix, G., and Villeval, M.-C. (2007). Tax evasion and social interactions. *Journal of Public Economics*, 91(11–12): 2089–2112. DOI 10.1016/j.jpubeco.2007.03.005. URL <http://www.sciencedirect.com/science/article/pii/S0047272707000497>.
- Gigerenzer, G., and Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, 62(1): 451–482. DOI 10.1146/annurev-psych-120709-145346. URL <http://www.annualreviews.org/doi/abs/10.1146/annurev-psych-120709-145346>.
- Gill, D., and Prowse, V. (2013). A novel computerized real effort task based on sliders. MPRA Paper 48081, University Library of Munich, Germany. URL <http://ideas.repec.org/p/prapa/mprapa/48081.html>.
- Glaeser, E., and Scheinkman, J. (2001). Measuring social interactions. In S. Durlauf, and P. Young (Eds.), *Social dynamics*, pages 83–132. Boston, MA: MIT Press.
- Hastie, R., and Dawes, R. M. (2010). *Rational choice in an uncertain world: The psychology of judgment and decision making*. SAGE.
- Hayakawa, H. (2000). Bounded rationality, social and cultural norms, and interdependence via reference groups. *Journal of Economic Behavior & Organization*, 43(1): 1–34. DOI 10.1016/S0167-2681(00)00106-2. URL <http://www.sciencedirect.com/science/article/pii/S0167268100001062>.

- Manski, C. F. (1993). Identification of endogenous social effects: The reflection problem. *The Review of Economic Studies*, 60(3): 531–542. URL <http://restud.oxfordjournals.org/content/60/3/531.short>.
- Miron-Shatz, T., Stone, A., and Kahneman, D. (2009). Memories of yesterday's emotions: Does the valence of experience affect the memory-experience gap? *Emotion*, 9(6): 885. URL <http://psycnet.apa.org/journals/emo/9/6/885/>.
- Offerman, T., and Sonnemans, J. (1998). Learning by experience and learning by imitating successful others. *Journal of Economic Behavior & Organization*, 34(4): 559–575. URL <http://www.sciencedirect.com/science/article/pii/S0167268197001091>.
- Pokorny, K. (2008). Pay – but do not pay too much: An experimental study on the impact of incentives. *Journal of Economic Behavior & Organization*, 66(2): 251–264. URL <http://www.sciencedirect.com/science/article/pii/S0167268106002289>.
- Pornpitakpan, C. (2004). The persuasiveness of source credibility: A critical review of five decades' evidence. *Journal of Applied Social Psychology*, 34(2): 243–281.
- Sorensen, A. T. (2007). Bestseller lists and product variety. *The Journal of Industrial Economics*, 55(4): 715–738. DOI 10.1111/j.1467-6451.2007.00327.x. URL <http://onlinelibrary.wiley.com/doi/10.1111/j.1467-6451.2007.00327.x/abstract>.
- Tversky, A., and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157): 1124–1131. ISSN 0036-8075, 1095-9203. DOI 10.1126/science.185.4157.1124. URL <http://www.sciencemag.org/content/185/4157/1124>.
- Weber, E., and Johnson, E. (2006). Constructing preferences from memory. *The Construction of Preference*, Lichtenstein, S. & Slovic, P.,(eds.), pages 397–410. URL http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1301075.

Please note:

You are most sincerely encouraged to participate in the open assessment of this article. You can do so by either recommending the article or by posting your comments.

Please go to:

<http://dx.doi.org/10.5018/economics-ejournal.ja.2014-36>

The Editor