Pandora's Rules in the Laboratory*

Efthymios Lykopoulos† Georgios Voucharas‡ Dimitrios Xefteris§

Abstract

In theory, search conditions are a key determinant of Pandora's rule (i.e. of the optimal search pattern), and of the eventual payoffs (Weitzman, 1979; Doval, 2018). We compare different search conditions in the laboratory and find strong evidence that they affect subjects' order of inspection and payoffs greatly. Subjects are more conservative at the beginning of the search, and earn less if they are constrained to choose only among the inspected alternatives, than if they are not. These findings reinforce the empirical pertinence of the formal search literature, and generate novel insights regarding relevant settings of applied interest (e.g. the impact of market digitization on consumer behavior).

Keywords: Pandora's rule; laboratory experiment; Weitzman; search.
JEL classification: D83; D81; C91

*We thank Laura Doval, Philippos Louis, two anonymous reviewers and the editor, John Duffy, for their valuable feedback and advice. This project has received funding from the University of Cyprus (Dimitrios Xefteris) under the IMVASY research grant. The replication material for the study is available at https://doi.org/10.7910/DVN/RE8SJN.

†Department of Economics, University of Cyprus, PO Box 20537, 1678 Nicosia, Cyprus, e-mail: lykopoulos.efthymios@ucy.ac.cy

‡Department of Economics, University of Macedonia, Thessaloniki, Greece, e-mail: g.voucharas@uom.edu.gr

§Department of Economics, University of Cyprus, PO Box 20537, 1678 Nicosia, Cyprus, e-mail: xefteris.dimitrios@ucy.ac.cy, corresponding author.

1 Introduction

Weitzman (1979), in his seminal work, provided a search model with the name of "Pandora's box". The design of the problem was simple, yet, its solution, up to that point, seemed complex and computationally tedious. Weitzman elegantly provided a solution to this perplexing situation with an optimal stopping algorithm, which was forward-looking, easy to follow, and intuitive.[1] Optimal search algorithms in such problems since then have been commonly referred to as Pandora's rules.

The "Pandora's box" problem in its original form is as follows: An agent is confronted with N different boxes and may choose only one of them. Each box contains some reward that is randomly drawn from a box-specific distribution. The agent cannot observe the content of the box unless she decides to inspect it, but knows the distribution from which the reward is drawn. In each round, the agent is allowed to inspect only one box by paying a cost, while standard discounting is applied. The agent can stop the search at any point and, while she is not bound to keep the box which she last opened, she is constrained to choose one of the previously inspected boxes. Since information in this setup is costly, it is hardly ever optimal to open all of the boxes and simply pick the one with the highest prize. At the same time, following a search rule based on the expected reward of each box can lead to a sub-optimal outcome.

Weitzman (1979) showed that, in such cases, what the agent needs to do is to simply compute the reservation value of each box, and apply the following algorithm: First, sort the boxes according to their reservation value in descending order. Then, open the box with the highest reservation value and compare the realized prize of that box with the reservation value of the next unopened box. If this value is greater or equal to the reservation value of the next-to-be-inspected box, the agent should keep the open box and terminate any further search.[2] Otherwise, the search continues.

Pandora's rule is easy to be identified in the setting of Weitzman (1979) but at the same time it is fairly sensitive to changes in the search conditions. Importantly, Doval (2018) demonstrated that by adding some flexibility (i.e. by allowing the agent to simply choose an uninspected box) Pandora's rule can change dramatically. Indeed, in a simple setup with a safe (high probability of a moderate payoff, and low probability of a zero payoff) and a risky box (low probability of a high payoff, and high probability of a zero payoff), inspection tends to start more often from the risky box when the agent is additionally allowed to choose an uninspected box, while search starts more often from the safe alternative when she is constrained to choose only from previously inspected boxes. Notice that without the flexibility to choose an uninspected box, inspecting a box both resolves uncertainty and, at the same time, expands the agent's choice set. However, when this

[1] See also Gittins (1979).

[2] Intuitively, as it is described in Weitzman (1979), the reservation value can be seen as the internal rate of return of each box.

type of flexibility is provided, the second channel is such that agents have more incentives to begin the search with more uncertain options.

While pinning down the theoretical properties of providing extra flexibility in search problems is of paramount importance, one would further like to assess the empirical relevance of the derived Pandora's rules. Unfortunately, this is not possible by analyzing observational data: Real search problems are characterized, not only by different search conditions, but also by heterogeneous agents' backgrounds, behavioral traits, education, etc. Therefore, to identify the effect of a change in search conditions on the employed search pattern and the corresponding payoffs, one has to turn to more controlled environments.

In this paper, we report results from a laboratory experiment suitably tailored to answer the questions above. We focus on a simplified search problem according to which a subject faces a risky and a safe option. The risky option has a 25 percent chance of containing 100 coins, and a 75 percent chance of containing nothing, while the safe option has a 50 percent chance of containing $X \in \{41, 42, ..., 99\}$ coins, and is empty with the remaining probability. That is, the maximum potential reward of the risky box is higher than that of the safe box, while the probability of a box being non-empty is higher for the safe box. The potential coins contained in the safe box, X, is known to subjects, but we allow them to vary in each round in order to extract more information regarding the employed search rule.

We investigate two alternative types of search conditions: the strict and the flexible one. Under strict search conditions, the search is conducted a la Weitzman (1979); that is, a subject can only keep a previously inspected box. Under flexible search conditions, the search is conducted a la Doval (2018); a subject can keep an uninspected box, provided that she has inspected at least one.[3] To inspect a box a subject must sacrifice 20 coins and throughout the experiment, each subject decides according to the same set of search conditions (across subject design).

Under strict search conditions, Pandora's rule dictates that search should be initiated with the risky option if the potential content of the safe box X is below 60 coins, and with the safe box if X is above 60 coins. In contrast, Pandora's rule under flexible search conditions is such that the search starts from the risky box independently of X.[4] What we find is that under strict search conditions search starts from the risky box when the payoff-to-riskiness index of the risky box exceeds the inspection cost. This index can be seen as the excess maximum payoff of the risky box (i.e. 100-X) over a measure of its relative riskiness. Meanwhile, under flexible search conditions

[3]While Doval (2018) does not use this restriction, employing it in the lab is very useful. This is because it imposes the same set of actions in the first step of the search process, under both sets of search conditions, while still allowing for different Pandora's rules.

[4]The optimal stopping rule under strict conditions suggests that subjects should stop searching whenever they encounter a positive reward. Under flexible search conditions, the optimal stopping rule is to inspect the risky box first and keep it if it is non-empty. In case it is empty, one should take the content of the safe box without first inspecting it.

search starts from the risky box when payoff-to-riskiness is positive, which by design is always true.

These theoretical predictions lead to our main empirically testable hypotheses: a) under strict search conditions, the search process starts more often from the safe box than under flexible conditions, and b) under strict search conditions the first inspection choice will be more reactive to changes in X than under flexible conditions.

The experimental results support, by and large, these theoretical predictions. Subjects, indeed, react to the set of search conditions in the predicted way. That is, even if they are boundedly rational, and their choices are affected by salient features of the options (e.g. expected reward of each box, etc.), they also react to the different sets of search conditions and exhibit a behavior that is comparatively similar to the corresponding Pandora's rules. Additional to our main hypothesis, we test for differences in payoffs across treatments since flexible search conditions allow the agent to save on the inspection cost. Our results indicate that subjects operating under flexible search conditions enjoy higher payoffs than subjects operating under strict search conditions. Moreover, since the optimal search rule remains constant under flexible search conditions we test for the ease of identification of the correct order of inspection across treatments. To this end, we find empirical evidence supporting that Pandora's rule is more often identified under flexible conditions than under strict conditions. Such findings confirm that the mathematical modeling of search problems can inform competently empirical research and policy design. In turn, this reaffirms the need for further formal analysis of search problems, since such studies do not only produce elegant results of theoretical interest, but also insights that are pertinent to search conducted by real subjects.

Finally, it is worth noting that the experimental testing of Pandora's rules –beyond its aforementioned general interest– also admits another –more applied– motivation. Traditionally, shopping has been an exploratory process where a potential buyer had to visit several stores before buying a good of uncertain quality. Nowadays, most stores (if not all) also provide the flexibility to the buyer to acquire a good online. That is, traditional shopping in physical stores is a search problem with little flexibility (one can only buy a previously inspected good), while contemporary shopping is a search problem with more flexibility (one can either inspect the good by visiting the physical store, or acquire the good online, saving the inspection costs at the expense of not resolving/reducing the uncertainty regarding product quality). Therefore, the empirical testing of adding flexibility in search problems generates novel insights regarding the non-trivial impact of market digitization on consumer behavior. For example, when two outlets provide both the option to shop online and in their physical stores, consumers are predicted to visit only the physical store of the riskier outlet; while when shopping online is not an option, starting the search from the safer outlet can be optimal.

The rest of the paper is organized as follows: In Section 2 we review the relevant literature, in Section 3 we present the theoretical benchmark, in Section 4 we describe the experimental design and state the testable hypotheses, in Section 5 we develop our results, and in Section 6 we conclude.

2 Relevant literature

The theoretical underpinning of our work can be seen as a subset of the broad family of "multi-armed bandit problems" as first described by Robbins (1952).[5] A typical setup is one where a risk-neutral agent faces a set of potentially different and independent arms of a slot machine which entail some uncertainty concerning the monetary prize they deliver. For the agent to gain a better perspective about the distribution of each hand, some exploration of the available options is required, which as a process, is assumed to be costly. Hence, a natural trade-off emerges between exploration and exploitation. Every time the agent inspects one option gains additional information which, in turn, affects the decision of which arm should be chosen next.

Naturally, several variations of the bandit setup have been adapted to answer questions related to market uncertainty. Regarding price dispersion, it is a well-established theoretical result that it emerges, in equilibrium, due to the combination of price uncertainty and search costs within a market. The most prominent example of this case comes from Stigler (1961) who, motivated by a consumer search problem, offers the foundation of what is commonly known in the economics literature as search theory.[6] A key characteristic of this framework is that the agent searches by deciding upon the number of quotations to be obtained, i.e. the sample size, which is assumed not to change as the search process unveils new information. McCall (1970) expands on Stigler's original idea with a model of a dynamic analysis of job search where the decision is based on a sequential rather than a fixed-sample approach. As Mortensen (1986) points out, in terms of expected future profits, a sequential process of sampling dominates that proposed by Stigler.[7] Consequently, this observation has shifted the literature to formulate search problems such that agents operate following a sequential search strategy rather than a fixed-sample one.[8]

More broadly, a sequential search strategy highlights the importance of the order of inspection. Simply put, in a sequential search setting the choice regarding the first or the next-to-be-inspected option is conditional on the current status of the search process. This implies that the decision process in these types of problems is state-dependent. Gittins (1979) exploits this Markovian structure and provides a dynamic index dictating the optimal search order and the optimal stopping rule. Concurrently, Weitzman (1979) characterizes the optimal search and stopping rule in a framework of deciding among different R&D projects. Doval (2018) considers a variant of Weitzman (1979) and highlights the importance of the inspection and choice rules. Namely, altering the flexibility of the original sampling rules, that is, allowing the agent to acquire an option without

[5] See Bergemann and Valimaki (2006) for an informative review of bandit problems.

[6] See also Stigler (1962) for a case relevant to the job market.

[7] More examples of the sequential search approach in the job market can be found in Mortensen (1970) and Gronau (1971).

[8] From the producer's point of view, an example of sequential market exploration comes from Rothschild (1974) who also finds that price dispersion can emerge in markets where demand is unknown. In such instances, a firm may engage in a costly process of price setting to get a better gauge of consumers' valuation.

having first to suffer the inspection cost, potentially yields a completely different optimal search and stopping pattern. Our main goal in this paper is to empirically verify the effect on the order of inspection of a varying level of flexibility in search conditions. To our knowledge, no paper undertakes this task and thus, our work aims to fill this gap in the literature.

Moreover, this paper also is intended to enhance our understanding of limited information acquisition and highlight its implications in terms of consumer welfare. Theoretically, bypassing, partly or completely, a costly research effort can have profound implications for an agent's welfare under an extended set of circumstances. Relevant to this notion, a result that frequently emerges is that partial evaluation of a product can be beneficial to consumers. As it is shown in Liu and Dukes (2016), it may not be optimal for consumers to fully evaluate a product before making a purchase. Abstaining from a thorough assessment may allow an agent to evaluate more firms which, in turn, may lead to lower prices and increased welfare. Similarly, Jain and Whitmeyer (2021) show that increased search costs, related to visiting a location, unambiguously hurt the consumers, but increased information costs, related to product investigation, benefit consumers.

Relatively low search activity compared to what is suggested by conventional consumer search theory –as in Weitzman (1979)– is a well-observed empirical outcome in the consumer search literature. To this end, a large body of field experiments has been developed to explain this phenomenon in terms of how information seeking in consumer search problems is affected by the core parameters linked with consumer search such as cost, relative and absolute uncertainty, risk aversion, etc. Most notably, Moorthy, Ratchford, and Talukdar (1997) explore the effect of prior beliefs and relative uncertainty on consumer search behavior and demonstrate that prior beliefs of a brand greatly shape consumers' search strategy. Relevant to the effect of search costs, Kim, Albuquerque, and Bronnenberg (2010) find that online tools increase consumer welfare while at the same time generate concentration around popular products. Similarly, Fong (2017) inspects how targeted advertisements may curtail customer incentives to search for unadvertised products. Fox and Hoch (2005) highlight the benefit of sequential exploration on consumer welfare in the context of grocery shopping. Nevertheless, it might be challenging for field studies to address differences in more fundamental aspects of the search environment, like heterogeneous search flexibility. Apart from the fact that collecting the necessary data for this purpose can be a difficult task, in terms of analysis, confounding factors affecting behavior may render causality hard to establish.

Parallel to field studies related to sequential search problems, a branch of the literature uses data generated by controlled experiments to assess the empirical relevance of theoretical models. Mainly, the focus of this type of research relates to the amount of search that is conducted by subjects and the extent to which search patterns are in line with an implied reservation value. For example, from a consumer's search perspective, Urbany (1986) testing Stigler's predictions reveals the negative effect that more certain prior beliefs about price dispersion have on the amount of search that is conducted by consumers. Concerning the labor market and job search, Schotter and Braunstein (1981) and Braunstein and Schotter (1982) test several theoretical predictions related to optimal search and find evidence in support of the reservation wage hypothesis. In total, subjects' observed behavior appears to be a combination of Stigler's prediction (1961; 1962) and an optimal search strategy, which as a consequence generates a decreasing reservation wage. Additionally, Cox and Oaxaca (1989) argue that subjects' decreasing reservation wage can be attributed to operating under a finite horizon, while Brown, Flinn, and Schotter (2011) find that this puzzle can be primarily explained by behavioral elements operating in conjunction with market factors. Relevantly, Casner (2021) shows that gaining more information on a previously unknown distribution increases the rate of decline in reservation value attributed to each choice compared to cases of full information. As we mentioned earlier and despite the useful insights about search theory generated by these experimental studies, the existing literature focuses on the amount of search that is conducted rather than the order of inspection of the available options.

Focusing on the behavioral aspect of this topic it is not unusual for individuals to follow some heuristic strategy when being confronted by such a complex and computationally challenging task as sequential search. Divergence from the optimal strategy as proposed by theory has been demonstrated in several instances, whether that is due to following a partially myopic algorithm when facing complex problems (Gabaix, Laibson, Moloche, & Weinberg, 2006), underestimating the time cost of search (Botti & Hsee, 2010), reaching a satisfactory level of utility (Caplin, Dean, & Martin, 2011), or due to an altered set of available choices (Ge, Brigden, & Häubl, 2015). In this paper we show that, when search conditions change, real agents adapt their search rule in the direction predicted by rational choice theory –but do not fully align with it– validating that both rational-choice factors and behavioral ones jointly determine the observed behavior.[9]

3 Theory

In this section we describe the two search environments that differ solely in the strictness of the search conditions, captured by the ability of the agent (or lack thereof) to select uninspected objects. We consider a minimal setup with only two options, and describe the optimal search rule for each set of search conditions.

A risk-neutral agent is faced with two boxes, r and s. Box $i \in \{r, s\}$ contains a prize X_i with a known to the agent probability p_i or is empty with probability $1 - p_i$. We assume that box r is the "risky" option and box s is the "safe" one. That is, $p_r < p_s$ and $X_r > X_s$. Whenever the agent decides to open a box has to pay a common inspection cost $c > 0$. We also assume that the expected value of opening every box i is strictly positive, that is, $p_i X_i > c$, and that there is no time discounting.

[9]This is similar in spirit to Payne, Bettman, and Johnson (1988) who show that, from a consumer search perspective, individuals exhibit search strategies that adapt to the search environment.

Under **strict search conditions**, an agent can inspect the boxes in any order, and can stop inspecting at any point in time by keeping the content of one of the previously inspected box. Moreover, it is assumed that to acquire information about the content of each box, the agent must always pay upfront the inspection cost. Weitzman (1979) was the first to characterize the Pandora's rule for this problem. As he pointed out, in a sequential search problem similar to this, one cannot simply decide based on the expected value of each box since further search might improve the agent's payoff in case the previously inspected box is empty. It is obvious that when a box is empty, given the fact that the expected payoff of inspecting the remaining box is strictly positive, the agent continues with the unopened box.

Accordingly, to highlight the solution to this problem we need to study the ex-ante expected payoffs considering all plausible eventualities. The expected payoff of inspecting the risky box first, and keeping its content (and stopping the search) if it is not empty, and continuing the search with the safe box in case it is empty is

$$E_{rs} = -c + p_r X_r + (1 - p_r)(-c + p_s X_s) \tag{1}$$

while the expected payoff of inspecting the safe box first, and continuing the search only in case it is empty, is

$$E_{sr} = -c + p_s X_s + (1 - p_s)(-c + p_r X_r). \tag{2}$$

Notice that there is no other search rule that gives a higher expected payoff than these two. Indeed, a) stopping the search when the first inspected box is empty is not rational given that the expected value of opening each box is strictly positive, b) continuing the search after opening the risky box first and finding out that it contains a prize is dominated by stopping and keeping the content, and c) inspecting the safe box first and continuing the search even if it contains the prize delivers a lower expected payoff than E_{rs}.

Hence, the difference between the above expected values determines which box should be opened first. When $E_{rs} > E_{sr}$ holds, implies that

$$c(p_r - p_s) + p_r p_s (X_r - X_s) > 0 \tag{3}$$

or

$$X_r - X_s > c(\frac{1}{p_r} - \frac{1}{p_s}). \tag{4}$$

For notational convenience we define $\lambda_i = \frac{1}{p_i}$ for $i \in \{r, s\}$ to describe the level of riskiness of each box. Notice that an increased probability of a box i being non-empty, p_i, translates to a lower value of λ_i which implies that this box is less risky. Then, if we write $\Delta X = X_r - X_s$ and $\Delta\lambda = \lambda_r - \lambda_s = \frac{1}{p_r} - \frac{1}{p_s}$ we can define the payoff-to-riskiness index of the risky box as $\frac{\Delta X}{\Delta \lambda}$ and we can write the above expression as:[10]

$$\frac{\Delta X}{\Delta \lambda} > c. \tag{5}$$

Intuitively, in order for opening first the risky box to be optimal, its payoff-to-riskiness index should exceed the inspection cost; otherwise, the optimal search conditions require opening the safe box first.

Under **flexible search conditions**, an agent is allowed to choose a search rule from a wider class. That is, after the agent opens the first box she can either keep its content and stop, or continue the search by inspecting the second box as in the previous problem; however, now she can additionally stop the search by keeping the content of the remaining uninspected box. Search problems with additional flexibility have been studied by Doval (2018) and, as it turns out, they admit quite different Pandora's rules compared to the ones characterized by strict search conditions.

Note that Doval (2018) does not require the agent to inspect necessarily at least one box as we do, but for the applied purposes of the current study, it is important to add this extra assumption. Indeed, by doing so the agent has the same options in the first stage of the search under both sets of search conditions, thus allowing one to cleanly identify the effects of the different search conditions on the agent's first stage decisions. Undoubtedly, this restriction also imposes limitations on the external validity of the experiment, since several search problems that allow choice without costly inspection (e.g. online shopping) do not impose a minimum number of inspections. However, as it becomes clear below, this restriction cannot upset the optimal search order. By removing it, either Pandora's rule remains exactly the same (when search costs are low), or search does not take place at all (when search costs are high).[11]

For a detailed description of the agent's optimal search rule, suppose that she has decided to open the safe box first and that this inspected box contained a prize. Should she proceed with inspecting the risky one? It would make sense to inspect it if and only if

$$X_s < -c + p_r X_r + (1 - p_r) X_s \tag{6}$$

which simplifies to $\Delta X > c\lambda_r$. Nevertheless, it is easy to see that even when $\Delta X > c\lambda_r$ is true, the agent would still be better off by inspecting the risky box first, and if it is empty to take the safe box without inspection since

$$-c + p_r X_r + (1 - p_r) p_s X_s > -c + (1 - p_s) p_r X_r - p_s c + p_s p_r X_r + (1 - p_r) p_s X_s \tag{7}$$

[10] Intuitively, an increased value of $\Delta\lambda$ implies an increased relative riskiness of the risky box since $\frac{\partial \Delta\lambda}{\partial p_s} > 0$ and $\frac{\partial \Delta\lambda}{\partial p_r} < 0$.

[11] Without this restriction in place, a relatively high inspection cost induces an agent to forego the inspection of any of the available options and just choose the box with the highest expected value.

which can be reduced to $\frac{c}{\lambda_s} > 0$, which given the parameters we have chosen is always true. Moreover, it is straightforward that if the search process begins from the risky box, and that box is found to contain the prize, the optimal move is to keep the identified prize. Thus, we have demonstrated that whenever a non-zero prize is sampled it is never optimal to inspect another box. Hence, what is left to examine is what determines which box is to be inspected first. The expected payoff of inspecting the risky box first and if it is empty to take the safe box without inspection is

$$E'_{rs} = -c + p_r X_r + (1 - p_r) p_s X_s \tag{8}$$

and, symmetrically, the expected payoff of inspecting first the safe box and if it is empty to take the risky box without inspection is

$$E'_{sr} = -c + p_s X_s + (1 - p_s) p_r X_r. \tag{9}$$

The agent should open the risky box first if and only if $E'_{rs} > E'_{sr}$ which conveniently simplifies to $\Delta X > 0$, which by assumption is true. That is, by giving the agent this additional option to keep an unopened box in the second stage crucially affects her optimal first-stage behavior. It is this non-trivial consequence of search conditions on optimal search rules that we are primarily interested in testing empirically in the next section.

At this point we summarize Pandora's rule for each of the two problems, by the means of the following proposition.

Proposition 1 *The flexible search Pandora's rule is different from the strict search Pandora's rule. In specific:*

1. *under strict search conditions it dictates that a) search begins from the risky box if and only if the payoff-to-riskiness index of the risky box exceeds the inspection cost, and from the safe box otherwise, and b) agents keep a prize if identified, and continue the search otherwise.*
2. *under flexible search conditions it dictates that a) search always begins from the risky box, and b) agents keep a prize if realized, and keep the content of the uninspected box otherwise.*

4 Experimental design

The experiment was carried out at the Laboratory for Experimental Economics at the University of Cyprus (UCY LExEcon) and was designed with the use of *z-Tree* (Fischbacher, 2007). In total, sixty subjects were equally recruited in four sessions (that is, fifteen participants per session). The experiment comprised of two different treatments with two sessions per treatment, as Table 1, Panel A shows. On average, the duration of the experiment was fifty minutes and the total payment per subject was approximately 10.45 euros.[12]

[12]This includes five euros as a show-up fee. Each payment was made privately, in cash.

At the beginning of each session, subjects were given instructions which were read out.[13] Each of the experimental sessions lasted one hundred rounds, prior to which there were three trial rounds to help subjects become more familiar with the environment. All rounds were independent, while, in every session, subjects could not interact with each other.

At the beginning of each round, two closed boxes were displayed on the computer screen of each subject, a risky box (displayed as Box A) and a safe box (displayed as Box B). Each box could contain a positive amount of coins or could be empty. In particular, the risky box had a 25 percent chance of containing 100 coins and the safe box had a 50 percent chance of containing X coins. Specifically, X describes an integer that took values from 41 to 99, and was revealed to each subject at the beginning of every round, with each value in this set having the same probability of being selected. That is, the values of X were randomly drawn from a uniform distribution.[14]

The realized values of X, as well as the probability realizations, were constant across treatments, but differed across subjects within the same treatment, and across rounds. For instance, the first subject in the strict treatment faced the same values of X, both in terms of magnitude and order, with the first object in the flexible treatment. This allows us to compare different search conditions under equal terms. Table 1, Panel B reports the descriptive statistics for the values of X. Figure 1 presents the frequency density and Figure 2 the distribution of the X values across rounds.

At the beginning of each round, subjects did not know the actual content of each box, except for the total amount of coins each box might have contained along with the corresponding probability. In the first stage of each round, each subject was asked to decide which box to open first, the risky or the safe box. Once a box was opened, its content was revealed and then the subject could proceed to the second stage. Depending on the treatment, the second stage either consisted of two (strict) or three options (flexible).

More specifically, under strict search conditions, each subject had the following options: a) to keep the box that had been opened and receive its content, or b) to open the remaining box and keep either of the two, receiving the content of the selected box. Under flexible search conditions each participant had one more option in addition to the aforementioned ones. That is, she could choose to take the remaining uninspected box without first opening it. At the beginning of each round, under both treatments, each subject had an endowment of 40 coins. To open a box, each subject had to pay a fixed cost of 20 coins per inspection, which was deducted from the initial coins of each round.[15] Consequently, under strict search conditions, subjects who chose to keep the uninspected box did not have to pay an inspection cost. At the end of each round, each subject was informed about the content of the box that had been selected, the number of inspections that

[13]The instructions were originally written in Greek. A translated version of the instructions in English is presented in Appendix B. The Greek version is available upon request.

[14]The corresponding sequences of draws were generated with the use of *MATLAB*.

[15]Employing a relatively high inspection cost raises the salience of the order of inspection, discouraging subjects from opening both boxes in no particular order.

had been conducted, and the total payoff for this round. Payoffs were calculated as the sum of the coins included in the selected box plus the initial coins minus the cost of inspection.[16] In Appendix C, we report screenshots from all stages of the laboratory experiment.

Table 1: Summary

Panel A: Participation details				Panel B: Descriptive statistics of X values			
Session	Treatment	Subjects	Obs.	Mean	Std. dev.	Min	Max
1	Strict	15	1500	69.44	17.05	41	99
1	Flexible	15	1500	69.44	17.05	41	99
2	Strict	15	1500	69.31	17.00	41	99
2	Flexible	15	1500	69.31	17.00	41	99

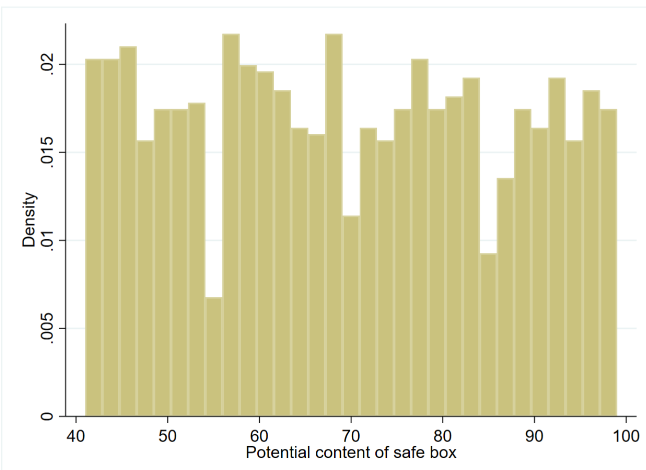

(a) Session 1

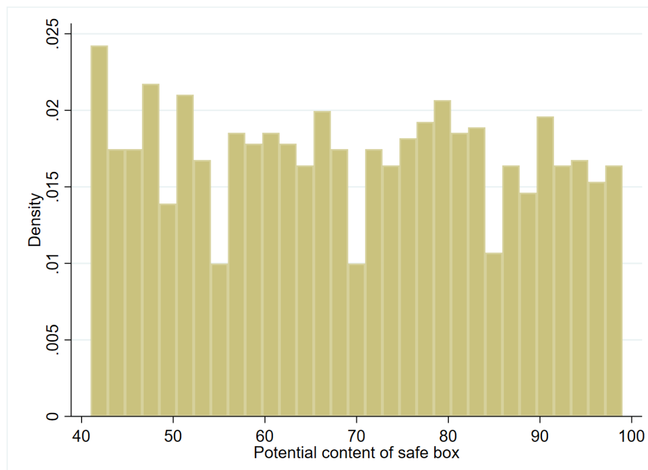

(b) Session 2

Figure 1: Histograms of the potential content of the safe box

[16]Following the completion of the experiment, five out of one hundred rounds were selected randomly and the final earnings of the participants were based on each subject's collected coins (i.e. payoffs) in these rounds. By doing so, we address considerations mainly with respect to the wealth effect (see also, Louis, Troumpounis, Tsakas, & Xefteris, 2022). The coins they have collected were then converted into euros at the rate of 1 euro for every 60.

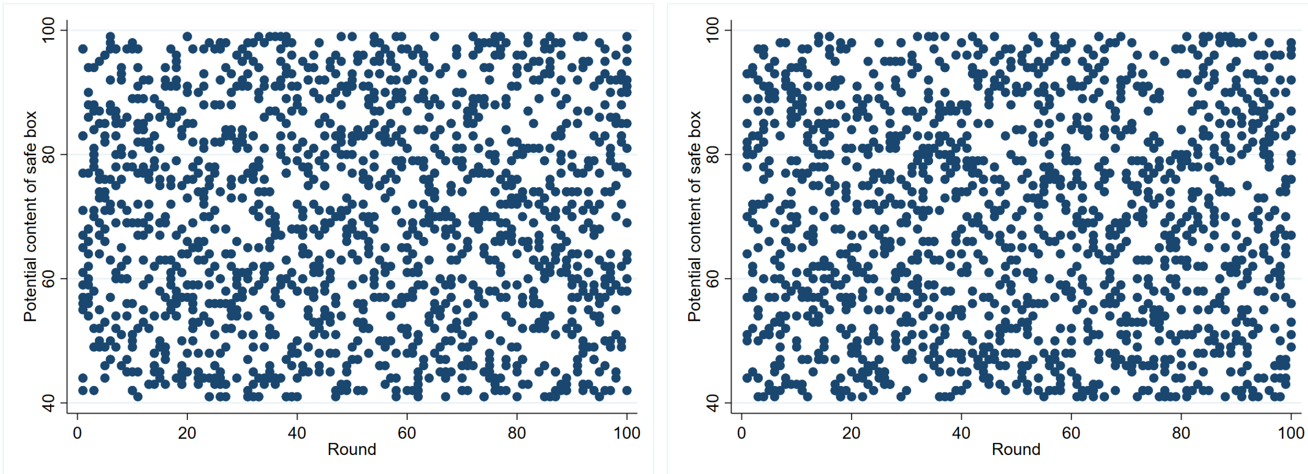

(a) Session 1 (b) Session 2

Figure 2: Distribution of the potential content of the safe box across rounds

4.1 Main Hypotheses

Given the theoretical arguments presented in Section 3 and the particular parameterization employed, for each of the two search conditions, Pandora's rules can be fully characterized. Accordingly, we present the underlined rules in Figure 3.

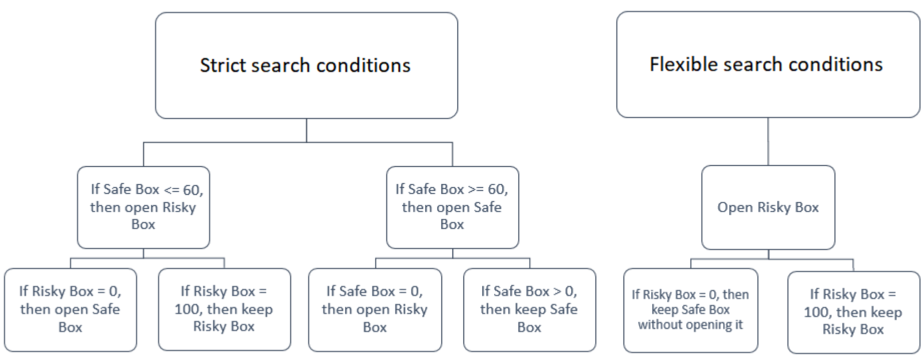

Figure 3: Optimal strategies across treatments

These formal results allow us to state a number of empirically testable hypotheses. The first thing we observe is that under strict conditions there are cases in which initiating the search from the safe box is optimal, while this is never the case under flexible search conditions. Hence, search should start more often from the risky box under flexible search conditions.

Hypothesis 1: Riskiness of first inspection. *Under flexible search conditions subjects open first the risky box more frequently than under strict search conditions.*

Moreover, the fact that under flexible search conditions the optimal search sequence always begins from the safe box, while under strict conditions the first inspection depends on the potential content of the safe box, indicates that we should observe subjects' first-stage decisions to be more reactive to changes in the potential prize of the safe box under strict search conditions, compared to the case of flexible search conditions.

Hypothesis 2: Reactivity to payoff-to-riskiness index. *Under flexible search conditions subjects are less reactive to changes in the payoff-to-riskiness index of the risky box, as far as their first inspection decision is concerned, than under strict search conditions.*

As previously stated, depending on the potential content of the safe box, the optimal search sequence changes under strict search conditions, while, it remains constant under flexible search conditions. This implies that it should be easier for subjects to identify the optimal search sequence in the latter case as opposed to the former.

Hypothesis 3: Success rate. *Under flexible search conditions subjects should be able to identify Pandora's rule more frequently than under strict search conditions.*

The flexible search conditions contain all search rules that are available under strict search conditions, and additional ones that allow the agent to follow any search order while saving on inspecting costs (by choosing the unopened box if the initially-inspected box turns out to be empty). So, it should be the case that agents manage to secure higher payoffs when they have a larger number of available rules at their disposal. Moreover, if Hypothesis 3 is true, and it is harder for subjects to identify Pandora's rule under strict search conditions, the relative payoff gain from more flexible search conditions should be even higher.

Hypothesis 4: Payoffs. *Under flexible search conditions subjects enjoy higher payoffs than under strict search conditions.*

While the above reasoning allows us to state this hypothesis, we need to stress that its confirmation by empirical/experimental data should not be taken for granted. For instance, as shown by Shah and Wolford (2007), the expansion of the choice set in a standard consumer problem has non-monotonic effects on the likelihood that the consumer will actually purchase a product, and thus on consumer's surplus.

5 Results

In this section, we test the hypotheses described in Section 4.1 and present some additional observations arising from our experiment.

We begin our analysis with the first two hypotheses regarding which box should be inspected first under the two search conditions of interest. By construction, Pandora's rules under both search conditions can be identical, or not, depending on the content of the safe box. Furthermore, whenever the two rules diverge, our setup conveniently postulates that subjects will initiate the search from different boxes. For this reason, we employ the variable *safeboxfirst*, which captures which box was inspected first during each round, taking the value of one whenever the safe box is opened first and zero otherwise.

Based on our theoretical results, we expect subjects to inspect the risky box more frequently under flexible search conditions.[17] On a preliminary basis, this prediction is validated by Table 2 which reports the frequency of opening the risky box first in each session. On average, under strict search conditions, the risky box was inspected first in 31.13% of the cases, whereas, under flexible search conditions, this happened in 76.27% of the cases.

On top of this, we employ three tests (i.e. t-test, Wilcoxon signed-rank test and sign test) to assess whether the aforementioned differences are statistically significant. Whenever we use these tests we conduct them on a subject level (i.e. we compute the average value for each variable of interest per subject, and then we use these averages as units of observation), to ensure that the assumption of independence is satisfied. In addition, given the paired nature of our observations (that is, subject pairs across treatments face the same sequences of draws) we focus on matched-pairs tests. The results are presented in Table 3 and clearly indicate that subjects choose to open the risky box more often when the flexible search conditions are in effect.

To fully exploit the richness of our data, we also perform probit regressions using *safeboxfirst* as the dependent variable on a treatment dummy variable, taking the value of zero in case of the strict search conditions and the value one under flexible search conditions. In all our regressions the observations are on a round-subject level (i.e. we have 100 observations per subject). Then, we estimate and present the average marginal effects. As can be seen in Table 4, column (1), we find strong evidence of a treatment effect.[18] Specifically, we find that under strict search conditions, subjects inspect the risky box first less frequently in comparison to the case of flexible search conditions. Moreover, when we include information about the content of the safe box in our regression as shown in Table 3, column (2), we see that as the potential content of the safe box increases, subjects tend to inspect the risky box first even less frequently, thus highlighting the switch that takes place under strict search conditions.

Based on these observations we can state our first result:

Result 1: *Under strict search conditions, subjects inspect the risky box first less frequently.*

Since under flexible search conditions subjects should always initiate the search from the risky box, the potential content of the safe box should not affect their decision. On the contrary, under strict search conditions, the optimal initial inspection depends on the potential content of the safe box, hence subjects' behavior should be more sensitive to such changes in the corresponding treatment.

[17]This hypothesis holds even when subjects are characterized by some degree of risk aversion. Refer to Appendix A for a numerical demonstration of this result.

[18]For comparison reasons, we have replicated the analysis by employing an OLS regression. Both approaches provide similar results. The OLS results are available upon request.

Table 2: Frequency of opening the risky box first

Time	Rounds 1-100		Rounds 51-100	
Treatment	Strict	Flexible	Strict	Flexible
Session 1	32.27%	76.33%	29.73%	83.87%
Session 2	30.00%	76.20%	27.73%	81.60%
Average	31.13%	76.27%	28.73%	82.73%

Notes: Percentages reported are based on the average observed frequencies.

Table 3: Subject-level tests: Frequency of opening the risky box first

Tests	t-test	Wilcoxon	sign test
Rounds 1-100	-7.373 (0.000)	-4.330 (0.000)	(0.000)
Rounds 51-100	-8.698 (0.000)	-4.588 (0.000)	(0.000)

Notes: For the one-sided t-test we report t-statistics, for the Wilcoxon signed-rank test we report z-statistics and for the one-sided sign test we report only p-values. P-values are in parentheses. All tests are based on paired data on a subject level.

Table 4: Marginal Effects

Dependent Variable: safeboxfirst	(1)	(2)	(3)
Treatment	-0.3998***	-0.3999***	-0.0004
	(0.4001)	(0.0404)	(0.1352)
SafeBox		0.0085***	0.0113***
		(0.0009)	(0.0014)
Treatment*SafeBox			-0.0057***
			(0.0018)
Observations	6,000	6,000	6,000

Notes: Table 4 reports the average marginal effects of the variables of interest after estimating a probit model with *safeboxfirst* as the dependent variable and *Treatment, SafeBox, Treatment*SafeBox* and round dummies as covariates. Delta-method robust standard errors are in parentheses. *** denote statistical significance at the 1% level. *safeboxfirst*=0 if the risky box is opened first and *safeboxfirst*=1 if the safe box is opened first. *SafeBox* refers to the potential content of the safe box.

To test this proposition, we incorporate an interaction term between the search rule and the potential content of the safe box. What we find is that when the potential content of the safe box and the interaction term are included, as presented in Table 4, column (3), the treatment effect enters only through the interaction term. Hence, one could argue that subjects are more reactive to changes in the content of the safe box when operating under the strict search conditions rather than under the flexible ones.

To shed further light on this observation, we create a variable called *reactivity* on a subject level, which captures the effect of the content of the safe box on the first inspection decision for every single subject. Reactivity is calculated by extracting the beta coefficients for every subject from regressing *safeboxfirst* on the potential content of the safe box. That is, we perform an OLS regression for each subject using *safeboxfirst* as the dependent variable, and the potential content of the safe box as the independent variable using round-subject level information (i.e. 100 observations per subject); we then define the relevant coefficient as the *reactivity* of the corresponding subject. Next, we test whether subjects operating under strict search conditions are, on average, more reactive to changes in the content of the safe box. We report the results of the relevant statistical tests in Table 5 which confidently point toward a specific direction. More concisely,

Result 2.1: *Under strict search conditions the content of the safe box affects subjects' decisions more than under flexible search conditions.*

Result 2.2: *Under strict search conditions and by taking into account the first inspection decision, subjects are more reactive to changes in the payoff-to-riskiness index.*

Table 5: Subject-level tests: Differences in reactivity parameter across treatments.

Tests	t-test	Wilcoxon	sign test
Rounds 1-100	2.795 (0.005)	2.520 (0.012)	(0.049)
Rounds 51-100	2.697 (0.006)	2.469 (0.014)	(0.018)

Notes: For the one-sided t-test we report t-statistics, for the Wilcoxon signed-rank test we report z-statistics and for the one-sided sign test we report only p-values. P-values are in parentheses. All tests are based on paired data on a subject level.

We proceed with the next hypothesis which describes how well individuals identify what is best for them, i.e. the optimal search sequence as visually presented in Figure 3. As stated above, under flexible search conditions the optimal sequence remains constant throughout all rounds and thus, easier for subjects to eventually identify. To test this hypothesis, we have created a dummy variable that accounts for cases where individuals have followed the rule successfully. The observed difference across treatments in the success rate, which captures the percentage of rounds in which a subject follows Pandora's rule, is statistically significant and the detailed results of the relevant subject-level tests are presented in Table 6.

Having established this difference between the two treatments, we delve a bit deeper into the success rate of individuals and explore how it evolves along with the rounds of the experiment. To some degree it is natural to expect that subjects need some time to figure out how to initiate the search properly. Thus, as the experiment progresses one would expect individuals to get more accustomed to the process and become more consistent in terms of choosing optimally. On a preliminary basis regarding this matter, the descriptive statistics presented in Table 7 point toward this direction. In this table we present the average per-subject success rate under both search conditions, first, with respect to the experiment in its entirety, and second, with respect to only the last fifty rounds of the experiment. [19] What we observe is that, on average, subjects operating under flexible search conditions tend to improve their success rate over the last fifty rounds of the experiment while, those operating under strict conditions do not seem to exhibit any difference in performance compared with their overall success rate. However, to properly establish the existence of a learning process that presumably takes place under flexible search conditions we use regression analysis. Using the success rate as the dependent variable on the rounds of the experiment unveils an on-average improvement in performance as reported in Table 8. Even so, when this effect is decomposed between treatments, we find that it is predominantly driven by the flexible search conditions where the optimal search sequence is more apparent. This result is also visually illustrated in Figures 4.a and 4.b where the red line corresponds to the fitted values at 95% confidence interval.[20]

Based on the aforementioned results we can state our findings with respect to Pandora's rule identification frequency across treatments. There is ample evidence of a statistically significant difference in the success rates across treatments, while, there is an indication of learning under flexible search conditions.

Result 3.1: *Subjects under flexible search conditions are more successful in identifying the optimal search sequence.*

Result 3.2: *Subjects under flexible search conditions exhibit improvement as they get more accustomed to the game.*

Table 6: Subject-level tests: Success rate across treatments

Tests	t-test	Wilcoxon	sign test
		First move	
Rounds 1-100	-0.402 (0.346)	-1.275 (0.202)	(0.021)
Rounds 51-100	-1.417 (0.084)	-1.945 (0.052)	(0.003)
		All moves	
Rounds 1-100	-3.595 (0.001)	-3.024 (0.003)	(0.001)
Rounds 51-100	-4.514 (0.000)	-3.385 (0.001)	(0.000)

Notes: For the one-sided t-test we report t-statistics, for the Wilcoxon signed-rank test we report z-statistics and for the one-sided sign test we report only p-values. P-values are in parentheses. All tests are based on paired data on a subject level.

We conclude our analysis with the realized payoffs in each treatment. We have already empirically demonstrated that under flexible search conditions subjects become better in identifying

[19] Examining separately the last fifty rounds of the experiment can unveil potential changes of subjects' behavior across time.

[20] We also compare the performance between the two treatments over the first 50 rounds of the experiment and find that the average performance does not differ statistically across treatments. This observation may be an indication that the observed better performance in the flexible search conditions is due to individuals learning as the experiment progresses.

Table 7: Average success rate per subject

Time	Rounds 1-100		Rounds 51-100	
Treatment	Strict	Flexible	Strict	Flexible
		First move		
Session 1	73.27%	76.33%	75.07%	83.87%
Session 2	75.60%	76.20%	76.13%	81.60%
Average	74.43%	76.27%	75.60%	82.73%
		All moves		
Session 1	56.53%	75.60%	56.93%	82.80%
Session 2	59.73%	74.53%	60.80%	81.60%
Average	58.13%	75.07%	58.87%	82.20%

Notes: Success rate is based on the optimal strategies depicted in Figure 3.

Table 8: Success rate across treatments

Dep. Var:	First move			All moves		
	Strict & Flexible	Strict	Flexible	Strict & Flexible	Strict	Flexible
Round	0.0014***	0.0001	0.0027***	0.0015***	-0.0001	0.0031***
	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)	(0.001)
Observations	6,000	3,000	3,000	6,000	3,000	3,000
R-squared	0.009	0.000	0.034	0.008	0.000	0.042

Notes: Standard errors clustered on a subject level are in parentheses. *** denote statistical significance at the 1% level. A constant term is included in all specifications.

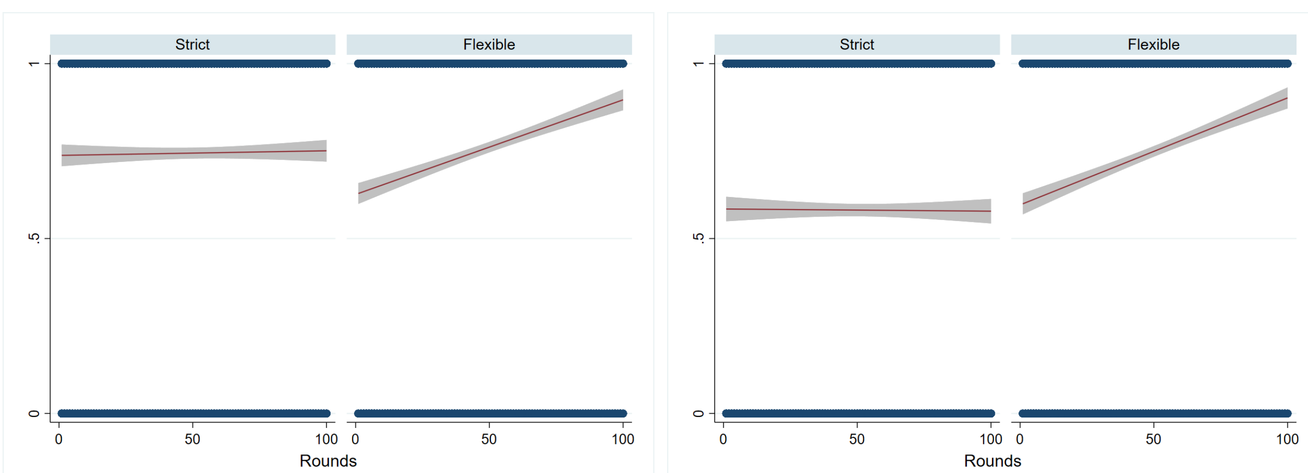

(a) Success rate restricted to first move

(b) Success rate based on all moves

Figure 4: Success rate across treatments

Pandora's rule with each round passing. At the same time, theory suggests that under flexible search conditions individuals –whenever the inspected box is empty– can save on the inspection costs by taking the remaining uninspected box. This observation is highlighted in Figure 5, which

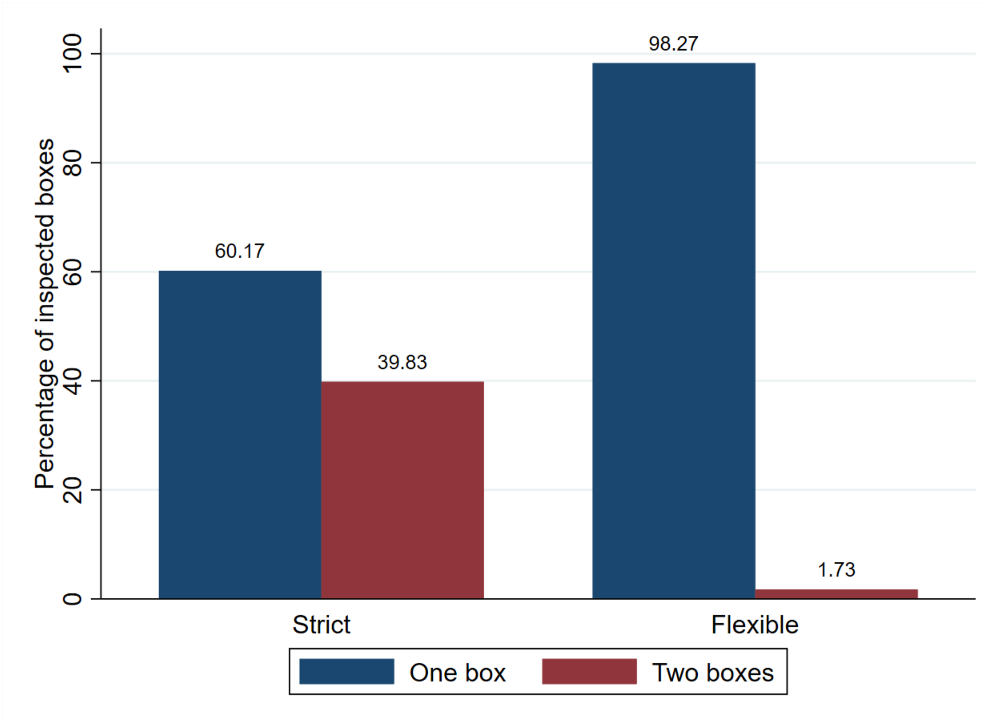

Figure 5: Opened boxes across treatments

shows that, under strict search conditions, subjects indeed inspect on average more boxes.[21]

Given the above, it comes as no surprise the result that payoffs are on average higher under flexible search conditions, both in the experiment in its entirety, and when considering only the last fifty rounds.[22] We conduct statistical tests on a subject level to bulletproof the aforementioned difference in payoffs across treatments in Table 9. All tests perform similarly, indicating that this difference is statistically significant at the 1% level. Nevertheless, to provide more concrete evidence we perform a linear regression as well. Our regression results provided in Table 10, strongly indicate that subjects, when given the option to save up on inspection costs, enjoy higher payoffs compared to those that do not have this option. This leads us into stating our final result.

Result 4: *Payoffs under flexible search conditions are higher.*

While we find ample evidence in favour of higher payoffs and stronger learning under the flexible search conditions the underlying mechanism driving these differences is not clear. For instance, is this discrepancy in learning dynamics due to the fact that one treatment provides more options than the other (hence, learning the optimal way to play is harder), or is it because the optimal strategy in one treatment changes with the potential content of the safe box, while in the other it remains constant (again, making it harder to learn the optimal strategy)? To address these concerns we conduct four additional experimental treatments in which we keep the

[21] This finding is strongly supported by performing linear regression analysis. See also Table A1 in Appendix D.
[22] We provide descriptive statistics in Table A2 in Appendix D.

Table 9: Subject-level tests: Payoffs across treatments

Rounds	t-test	Wilcoxon	sign test
Full (1-100)	-37.28 (0.000)	-4.782 (0.000)	(0.000)
Half (51-100)	-39.21 (0.000)	-4.782 (0.000)	(0.000)

Notes: For the one-sided t-test we report t-statistics, for the Wilcoxon signed-rank test we report z-statistics and for the one-sided sign test we report only p-values. P-values are in parentheses. All tests are based on paired data on a subject level.

Table 10: Payoffs across treatments

Dependent Variable: Payoff	(1)	(2)
Treatment	14.6520***	14.6520***
	(0.779)	(0.719)
SafeBox		0.4691***
		(0.030)
Observations	6.000	6.000
R-squared	0.054	0.085

Notes: Standard errors clustered on a subject level are in parentheses. *** denote statistical significance at the 1% level. *Treatment*=0 for strict search conditions and *Treatment*=1 for flexible search conditions. *SafeBox* refers to the potential content of the safe box. Round dummies and a constant term are included in all specifications.

potential content of the safe box constant throughout the session.[23] More specifically, we follow a 2x2 design, where in each treatment we either have strict or flexible search conditions while, the potential content of the safe box is either 50 or 90 in all rounds. That is, we select one value for the potential content of the safe box to be below and one value above 60: the threshold at which, in theory, Pandora's rule is expected to change under the strict search conditions.[24]

By keeping the optimal search rule constant in each of these new four treatments, the differences in learning patterns across search conditions disappear. That is, when the potential content of the safe box is 50 we find that no learning takes place neither under flexible, nor under strict conditions, while when this is 90 we find evidence of learning under both sets of search conditions. Therefore, one can attribute the differences in learning observed between our two main treatments, mainly to the fact that in one of them Pandora's rule varies with the potential content of the safe box while in the other it does not. However, we still find that the payoffs are significantly higher under the flexible search conditions than under strict search conditions. Hence, the discrepancy in payoffs between the two main treatments can be attributed with confidence to the fact that the

[23] We would like to thank the Editor and one anonymous Reviewer for urging us to work in this direction.

[24] We also conduct a threshold test proposed by Hansen (2000) using data from our main treatment with strict search conditions (Table A3 in Appendix D). The estimated threshold we get is 67, which is compatible with a mildly risk-seeking attitude (see Figure A4 in Appendix A).

flexible search conditions enlarge the set of available strategies, thus allowing agents to achieve better outcomes.[25]

6 Conclusion

We have conducted the first laboratory experiment that comparatively tests the empirical relevance of Pandora's rules corresponding to different search environments. Our results align with the theoretical predictions and provide strong evidence of the different search patterns that emerge conditional on the search rules. Subjects initiate the search process more frequently from the safe alternative whenever they are constrained to always pay the inspection cost, while, subjects not facing this restriction are more likely to begin the search from the risky option. Moreover, we discuss the implications with respect to payoffs and learning in this context. We demonstrate that when subjects are not constrained to choose only among inspected options, they enjoy higher payoffs compared to when such constraints are present, while –once we control for Pandora's rule variability– learning patterns seem broadly symmetric across search environments.

Of course, our work does not address all the aspects of the issue at hand. Future research could focus on other elements in the search environment that may interact with the strictness of search conditions in determining subjects' behavior. Considering asymmetric (i.e. box-specific) search costs and/or a larger number of available alternatives appears particularly promising. The reason is that the inspection cost is of paramount importance in determining the optimal order of search. Hence, a heterogeneous cost structure may exaggerate or dampen the observed differences across search conditions identified by this paper. Moreover, a larger number of alternatives might uncover additional factors that determine optimal behavior beyond the payoff-to-riskiness index of the riskiest box. While dealing with such extensions falls out of the scope of the current paper, we hope that our work informs subsequent studies and helps them design experiments that can assess the pertinence of additional relevant factors.

References

Bergemann, D., & Valimaki, J. (2006). Bandit problems. Cowles Foundation Discussion Paper, No.1551.

[25] One is referred to Appendix E for the quantitative analysis supporting these claims, as well as for additional information regarding these four new treatments. Table A4 presents participation details, Tables A5, A6, and A7 deal with payoff comparisons, and Table A8 addresses learning. Using these new data we also find that, ceteris paribus, more flexible search conditions increase the likelihood that search starts from the risky box, providing further evidence in support of Hypothesis 1 (Table A9). Given the fixed nature of the potential content of the safe box, these additional treatments are less well-suited to address Hypothesis 2 –and any question related to the sensitivity of the subjects' choices to changes in X– compared to our original treatments.

Botti, S., & Hsee, C. K. (2010). Dazed and confused by choice: How the temporal costs of choice freedom lead to undesirable outcomes. *Organizational Behavior and Human Decision Processes*, *112*(2), 161–171.

Braunstein, Y. M., & Schotter, A. (1982). Labor market search: An experimental study. *Economic Inquiry*, *20*(1), 133–144.

Brown, M., Flinn, C. J., & Schotter, A. (2011). Real-time search in the laboratory and the market. *American Economic Review*, *101*(2), 948–74.

Caplin, A., Dean, M., & Martin, D. (2011). Search and satisficing. *American Economic Review*, *101*(7), 2899–2922.

Casner, B. (2021). Learning while shopping: an experimental investigation into the effect of learning on consumer search. *Experimental Economics*, *24*, 238–273.

Cox, J. C., & Oaxaca, R. L. (1989). Laboratory experiments with a finite-horizon job-search model. *Journal of Risk and Uncertainty*, *2*(3), 301–329.

Doval, L. (2018). Whether or not to open pandora's box. *Journal of Economic Theory*, *175*, 127–158.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, *10*(2), 171–178.

Fong, N. M. (2017). How targeting affects customer search: A field experiment. *Management Science*, *63*(7), 2353–2364.

Fox, E. J., & Hoch, S. J. (2005). Cherry-picking. *Journal of Marketing*, *69*(1), 46–62.

Gabaix, X., Laibson, D., Moloche, G., & Weinberg, S. (2006). Costly information acquisition: Experimental analysis of a boundedly rational model. *American Economic Review*, *96*(4), 1043–1068.

Ge, X., Brigden, N., & Häubl, G. (2015). The preference-signaling effect of search. *Journal of Consumer Psychology*, *25*(2), 245–256.

Gittins, J. C. (1979). Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society: Series B (Methodological)*, *41*(2), 148–164.

Gronau, R. (1971). Information and frictional unemployment. *The American Economic Review*, *61*(3), 290–301.

Hansen, B. E. (2000). Sample splitting and threshold estimation. *Econometrica*, *68*(3), 575–603.

Harrison, G. W., & Rutstrom, E. (2008). *Risk aversion in the laboratory, risk aversion in experiments, vol. 12, 41-196.* Emerald Group Publishing Ltd.

Jain, V., & Whitmeyer, M. (2021). Search and competition with flexible investigations. *arXiv:2104.13159*.

Kim, J. B., Albuquerque, P., & Bronnenberg, B. J. (2010). Online demand under limited consumer search. *Marketing Science*, *29*(6), 1001–1023.

Liu, L., & Dukes, A. (2016). Consumer search with limited product evaluation. *Journal of Economics & Management Strategy*, *25*(1), 32–55.

Louis, P., Troumpounis, O., Tsakas, N., & Xefteris, D. (2022). Coordination with preferences over the coalition size. *Journal of Economic Behavior & Organization*, *194*, 105–123.

McCall, J. J. (1970). Economics of information and job search. *The Quarterly Journal of Economics*, 113–126.

Moorthy, S., Ratchford, B. T., & Talukdar, D. (1997). Consumer information search revisited: Theory and empirical analysis. *Journal of Consumer Research*, *23*(4), 263–277.

Mortensen, D. T. (1970). Job search, the duration of unemployment, and the phillips curve. *The American Economic Review*, *60*(5), 847–862.

Mortensen, D. T. (1986). Job search and labor market analysis. *Handbook of Labor Economics*, *2*, 849–919.

Pan, X., Ratchford, B. T., & Shankar, V. (2004). Price dispersion on the internet: a review and directions for future research. *Journal of Interactive Marketing*, *18*(4), 116–135.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1988). Adaptive strategy selection in decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *14*(3), 534.

Robbins, H. (1952). Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, *58*(5), 527–535.

Rothschild, M. (1974). A two-armed bandit theory of market pricing. *Journal of Economic Theory*, *9*(2), 185–202.

Schotter, A., & Braunstein, Y. M. (1981). Economic search: an experimental study. *Economic Inquiry*, *19*(1), 1–25.

Shah, A. M., & Wolford, G. (2007). Buying behavior as a function of parametric variation of number of choices. *Psychological Science*, *18*(5), 369-370.

Slonim, R. (1994). Learning in a search-for-the-best-alternative experiment. *Journal of Economic Behavior & Organization*, *25*(2), 141–165.

Stigler, G. J. (1961). The economics of information. *Journal of Political Economy*, *69*(3), 213–225.

Stigler, G. J. (1962). Information in the labor market. *Journal of Political Economy*, *70*(5, Part 2), 94–105.

Urbany, J. E. (1986). An experimental examination of the economics of information. *Journal of Consumer Research*, *13*(2), 257–271.

Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, *47*, 641–654.

Appendix

A Risk Aversion

In this section we consider a more general case of the problem where the agent is characterized by some degree of risk aversion. Note, that our goal here is not to demonstrate that the properties of the model in the main body of our work extend to cases where risk aversion is present. Rather, we provide numerical examples of the robustness of this framework conditional on the parameters that have been used during the experiment.

To elaborate that our main results hold under risk aversion, we take a case where an agent's preferences are best described by a constant relative risk-aversion –henceforth CRRA– utility function which has the form of $U(W) = \frac{W^{1-\sigma}-1}{1-\sigma}$, where W is the total wealth of the agent and σ is the risk aversion parameter with $\sigma \in [0,1) \cup (1,\infty)$. Note that the model discussed previously is just a special case of CRRA utility where $\sigma = 0$.

A.1 Strict Search Conditions

Consider the case where the agent has already inspected either of the two boxes and to her dismay found out that the box was empty. Having reached this outcome and in conjunction with the presence of risk aversion, proceeding with the next uninspected box might not be the most preferred move anymore. More intuitively, a risk averse individual might prefer abstaining from a gamble in order to save on the inspection cost regardless of the fact that the expected payoff of inspecting the remaining box is positive, thus ending the search process. Accordingly, in our numerical analysis we consider two sets of cases for the agent: a) continuing with the search when the first inspected box is empty, and b) stopping whenever the inspected box is empty. Again, under both search environments if opening a non-empty box induces more search, then the agent would rather inspect the second box first, hence, with respect to our numerical exercise we leave these cases out. The expected payoffs that the agent needs to consider are

$$E^{\sigma}_{rs} = \frac{p_r\Big((c+X_r)^{1-\sigma}-1\Big)}{1-\sigma} + \frac{(1-p_r)p_s\Big(X_s^{1-\sigma}-1\Big)}{1-\sigma} \tag{10}$$

$$E^{\sigma}_{rstop} = \frac{p_r\Big((c+X_r)^{1-\sigma}-1\Big)}{1-\sigma} + \frac{(1-p_r)\Big(c^{1-\sigma}-1\Big)}{1-\sigma} \tag{11}$$

$$E^{\sigma}_{sr} = \frac{p_s\Big((c+X_s)^{1-\sigma}-1\Big)}{1-\sigma} + \frac{(1-p_s)p_r\Big(X_r^{1-\sigma}-1\Big)}{1-\sigma} \tag{12}$$

$$E^{\sigma}_{sstop} = \frac{p_s\Big((c+X_s)^{1-\sigma}-1\Big)}{1-\sigma} + \frac{(1-p_s)\Big(c^{1-\sigma}-1\Big)}{1-\sigma} \tag{13}$$

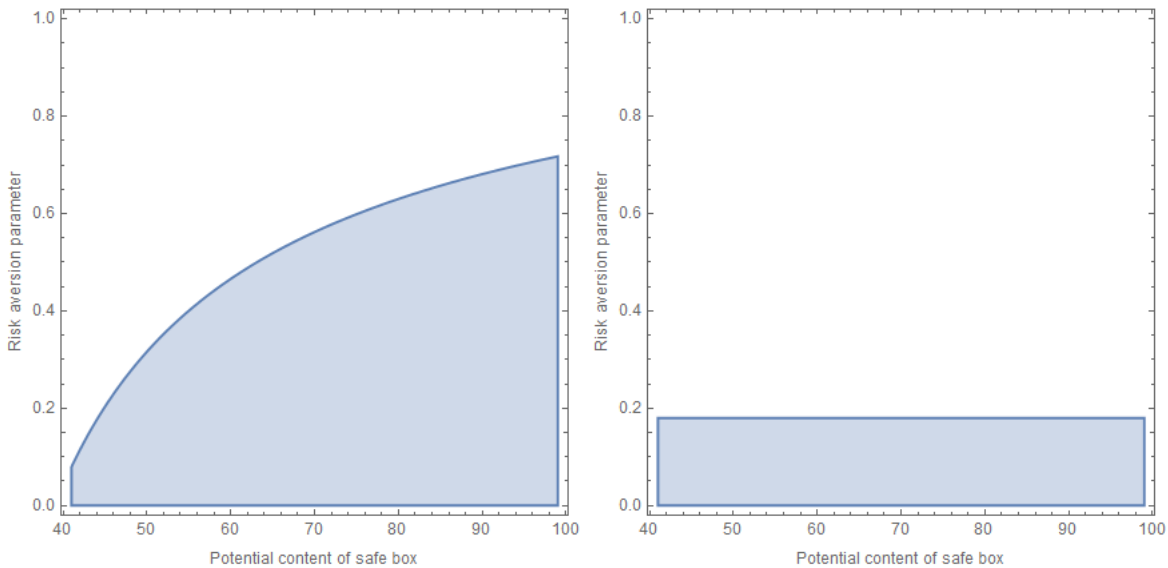

Figure A1: Left panel: Region where $E^{\sigma}_{rs} > E^{\sigma}_{rstop}$. Right panel: Region where $E^{\sigma}_{sr} > E^{\sigma}_{sstop}$. Both panels refer to strict search conditions

which correspond to inspecting the risky box and then the safe box, inspecting the risky box and then stopping, inspecting the safe box and then the risky box, and inspecting the safe box and then stopping respectively.

Given the parameters in our experiment, that is $X_r = 100$, $p_r = 0.25$, $X_s \in [41, 99]$, $p_s = 0.5$ and for $\sigma \in [0, 1) \cup (1, \infty)$ we compare the above payoffs. A natural point to begin with is by exploring how risk aversion might discourage an agent from further inspection when the first box in the search process turns out to be empty. In Figure A1 we see that this is exactly the case. More specifically, it becomes evident that given a sufficiently large degree of risk aversion, investigating a box and then stopping dominates investigating the same box and then continuing with the next uninspected box. Notice that the upper bound in both shaded regions plotted above corresponds to the combinations of σ and X_s for which the individual is indifferent between continuing with the search and stopping when the already-inspected box is empty. Intuitively, a sufficiently risk averse individual prefers retaining the inspection cost rather than participating in another gamble.

In the same spirit, in Figure A2 we highlight the area where inspecting the safe box first, regardless of whether the agent continues or not. It now becomes clear that the larger the degree of risk aversion, the more probable it is for the agent to turn to a safer option with respect to which box should be inspected first regardless of whether she should stop after the first inspection. Thus, we have demonstrated numerically that under strict search conditions, the presence of risk aversion enhances the main result of our model in Section 3.

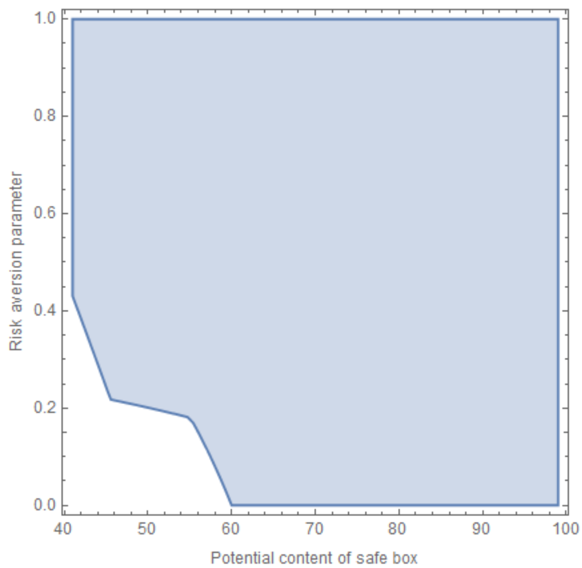

Figure A2: Region where inspecting the safe box first dominates inspecting the risky box first under strict search conditions

A.2 Flexible Search Conditions

As previously, in this type of framework, i.e. when the agent has a larger set of options, incentives become more clear-cut. As in the risk-neutral case, whenever the first box that has been inspected is empty it is not optimal to stop because, once again, taking the uninspected box without accruing the inspection cost yields a higher payoff than leaving empty-handed. Accordingly, the payoffs that the agent needs to consider are

$$E^{\sigma\,\prime}_{rs} = \frac{p_r\Big((c+X_r)^{1-\sigma}-1\Big)}{1-\sigma} + \frac{(1-p_r)p_s\Big((c+X_s)^{1-\sigma}-1\Big)}{1-\sigma} \tag{14}$$

and

$$E^{\sigma\,\prime}_{sr} = \frac{p_s\Big((c+X_s)^{1-\sigma}-1\Big)}{1-\sigma} + \frac{(1-p_s)p_r\Big((c+X_r)^{1-\sigma}-1\Big)}{1-\sigma} \tag{15}$$

which correspond to inspecting the risky box first and then taking the remaining safe one without inspection and vice versa. It is apparent that for any degree of risk aversion it remains non-optimal inspecting both boxes compared with inspecting just one. In Figure A3 we present our results concerning which box should be inspected first in the presence of risk aversion. As it would be expected, introducing risk aversion to our setup does not leave the agent unaffected. More specifically, as the level of risk aversion increases the agent becomes more prone to inspecting the safe box first instead of the risky one. Nevertheless, our numerical extension shows that this happens for relatively extreme levels of risk aversion.[26] The results from our numerical analysis can be seen as confidently demonstrating that risk aversion should amplify the expected outcome

[26] For a reference of what a typical level of risk aversion would be, see Harrison and Rutstrom (2008).

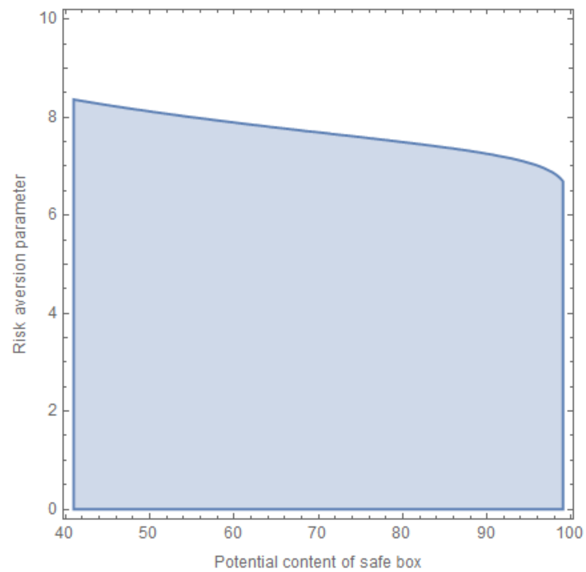

Figure A3: Region where inspecting the risky box first dominates inspecting the safe box first under flexible search conditions

under strict search conditions while not affecting it considerably under flexible search conditions. Hence, our predictions remain robust to risk aversion.

A.3 Risk-seeking

For completeness we assess the case of risk-seeking individuals in this context, which refers to cases where $\sigma \in (-\infty, 0)$. With regards to the first move under flexible search conditions, the prediction of our model trivially remains the same, as a more risk-seeking agent is even the more probable to initiate the search process from the risky box compared to a risk-neutral agent. On the other hand, this is not the case under strict search conditions. Briefly, it is intuitive to think –and easy to verify– that a risk-seeking individual would never stop the search process after inspecting either box and finding it empty. This translates to $E^{\sigma}_{rs} > E^{\sigma}_{rstop}$ and $E^{\sigma}_{sr} > E^{\sigma}_{sstop}$. This implies that, as in the risk-neutral case, the decision regarding whether the risky box is inspected first depends on whether this is true $E^{\sigma}_{rs} > E^{\sigma}_{sr}$.

As can be seen in Figure A4a, as the degree of risk-seeking increases, an individual requires a larger potential amount from the safe box in order to be deterred from beginning the search from the risky box. Finally, in Figure A4b we present the area where inspecting the safe box first is preferred under strict search conditions, where this time we also include negative values of σ.

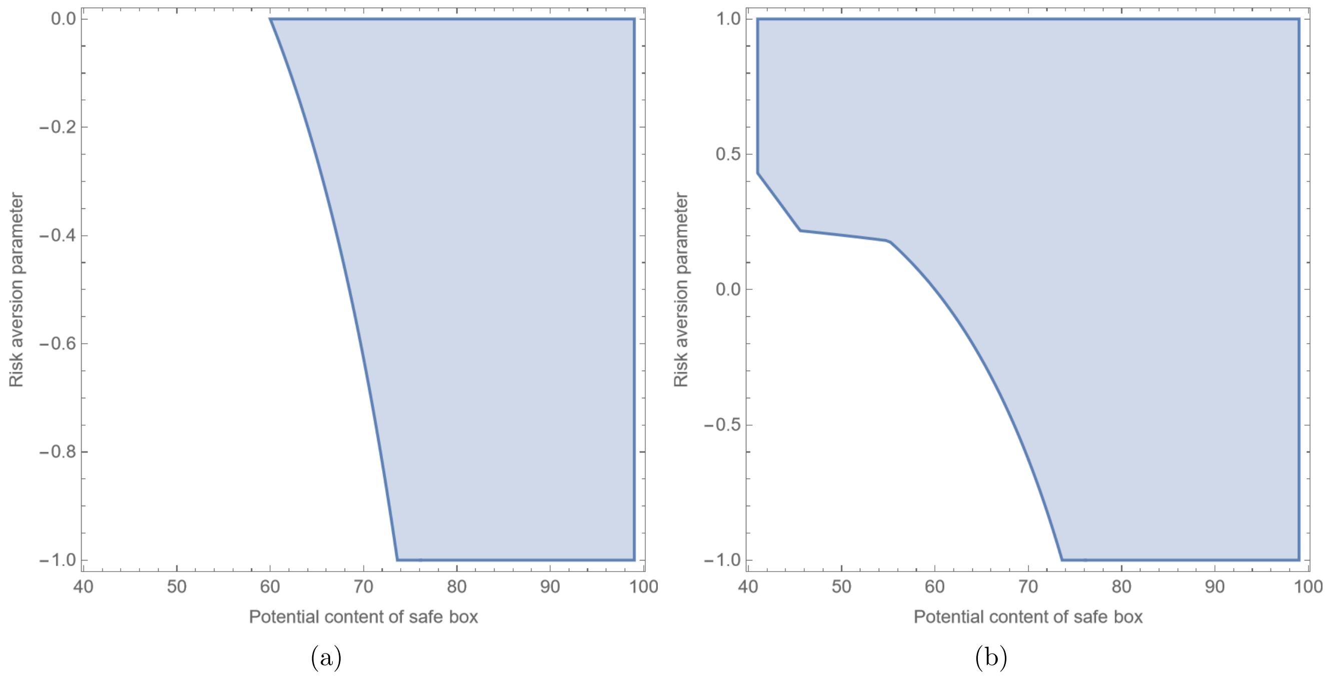

Figure A4: Region where inspecting the safe box first dominates inspecting the risky box first under strict search conditions

B Experimental instructions

The experiment was run in Greek. A translated version of the instructions in English is presented below for each treatment. The Greek version is available upon request.

B.1 Treatment: Strict Search Conditions

Thank you for participating in this session. The experimental session will be run using a computer and all answers will be given through it. Please do not talk to each other and keep quiet during the session. Please note that the use of mobile phones and other electronic devices is not permitted. Please read the instructions carefully, and if you have any questions, raise your hand. The answer that will be given will be announced to everyone.

The experiment

The experiment consists of one hundred rounds and it is individual. That is, each of the participants will not be able to interact with other participants. The rules are the same throughout the experiment. Your earnings depend on the decisions you make and on luck.

The boxes

At the beginning of each round, the computer shows to each subject two closed boxes, Box A and Box B. Each box may contain coins or may be empty.

In particular, Box A has a 25% chance of containing 100 coins and Box B has a 50% chance of containing X coins, where X is an integer from 41 to 99, which is announced at the beginning of each round (every number in this range has the same probability of being selected)

At the beginning of each round, you will not know what each box contains, except for the total amount of coins each box may contain and the probability that it contains them. That is, at the beginning of each round you will see an image like this: (The numbers here are random and refer only to the example below)

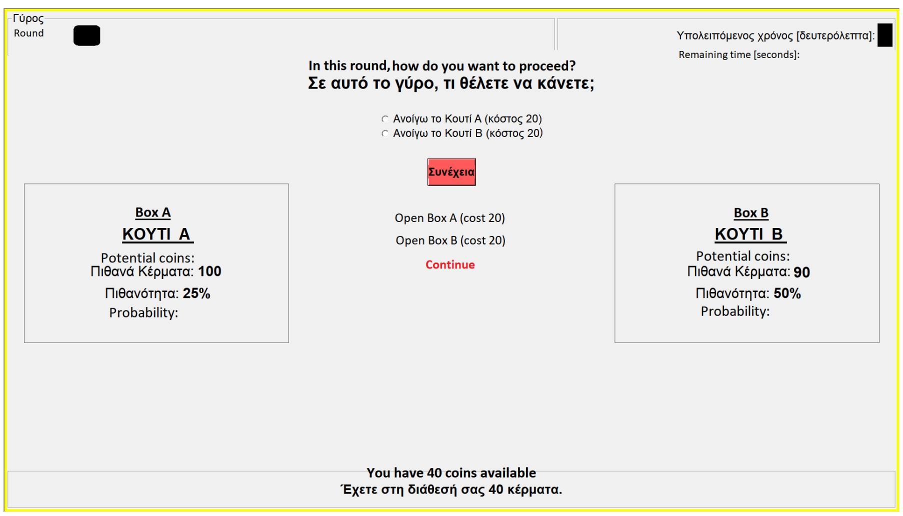

The procedure

At the beginning of each round, each subject is asked to open a box, Box A or Box B. Once a box is chosen its content is revealed.

Subsequently, each subject has the following options:
a) to keep the box that has been opened and receive its content.
b) to open the remaining box and choose to keep one of the two, receiving the content of the selected box.

Note: Each subject, at the end of each round can only keep one box.

Initial coins

At the beginning of each round, each subject will have 40 coins.

Opening cost

To open a box each subject has each time to pay a fixed cost. This cost is 20 coins per box she/he chooses to open, which are deducted from the initial coins of each round.

Payoffs

At the end of each round, each subject's payoff is calculated as:
Payoff = coins included in the selected box + initial coins - opening cost

Final earnings

At the end of the experiment, 5 rounds will be selected randomly and your final earnings will be based on your payoffs in these rounds plus the show-up fee (5 euros). The rate is 1 euro for every 60 coins. Each of the one hundred rounds has the same probability to be selected.

$$\textbf{Final Earnings} = \frac{1}{60}\ \textbf{x (sum of the points earned in 5 randomly selected rounds)} + \textbf{5}$$

Before the experiment begins, we will run three trial rounds to make sure that everyone understood the procedure. The coins that you will win during the trial rounds will not be included in your final profits.

B.2 Treatment: Flexible Search Conditions

Thank you for participating in this session. The experimental session will be run using a computer and all answers will be given through it. Please do not talk to each other and keep quiet during the session. Please note that the use of mobile phones and other electronic devices is not permitted. Please read the instructions carefully, and if you have any questions, raise your hand. The answer that will be given will be announced to everyone.

The experiment

The experiment consists of one hundred rounds and it is individual. That is, each of the participants will not be able to interact with other participants. The rules will be the same throughout the experiment. Your payoffs depend on the decisions you make and on luck.

The boxes

At the beginning of each round, the computer shows to each subject 2 closed boxes, Box A and Box B. Each box may contain coins or may be empty.

In particular, Box A has a 25% chance of containing 100 coins and Box B has a 50% chance of containing X coins, where X is an integer from 41 to 99, which is announced at the beginning of each round (every number in this range has the same probability of being selected)

At the beginning of each round, you will not know what each box contains, except for the total amount of coins each box may contain and the probability that it contains them. That is, at the beginning of each round you will see an image like this: (The numbers here are random and refer only to the example below)

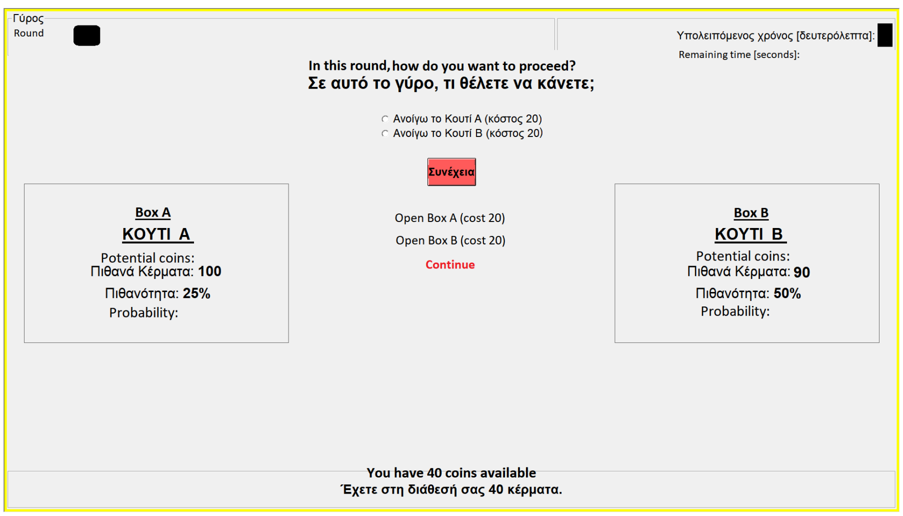

The procedure

At the beginning of each round, each subject is asked to open a box, Box A or Box B. Once the box is chosen, then its content is revealed.

Subsequently, each subject has the following options:
a) to keep the box that has been opened and receive its content.
b) to open the remaining box and keep one of the two, receiving the content of the selected box.
c) to keep the closed box without opening it and receive its content.

Note: Each subject, at the end of each round can only keep one box.

Initial coins

At the beginning of each round, each subject will have 40 coins.

Opening cost

To open a box each subject has to pay a fixed cost each time. This cost is 20 coins per box she/he chooses to open, which are deducted from the initial coins of each round.

Payoffs

At the end of each round, each subject's payoff is calculated as:
Payoff = coins included in the selected box + initial coins - opening cost

Final earnings

At the end of the experiment, 5 rounds will be selected randomly and your profits will be based on your payoffs in these rounds plus the show-up fee (5 euros). The rate is 1 euro for every 60 coins. Each of the one hundred rounds has the same probability to be selected.

$$\textbf{Final Earnings} = \frac{1}{60}\ \textbf{x (sum of the points earned in 5 randomly selected rounds)} + \textbf{5}$$

Before the experiment begins, we will run three trial rounds to make sure that everyone understood the procedure. The coins that you will win during the trial rounds will not be included in your final profits.

C Screenshots from the experiment

In this section, we present screenshots from all stages of the experiment, translated in English. Figures A5, A6, A7 and A8 correspond to the strict search conditions and Figures A9, A10, and A11 to the flexible search conditions.

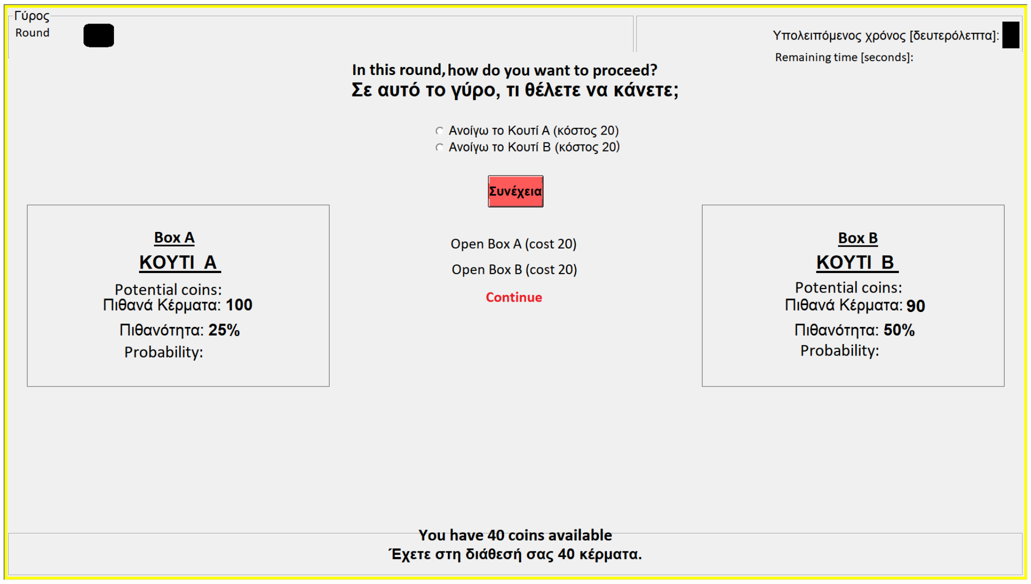

Figure A5: Strict search conditions: Stage 1. The subject is asked to open a box, Box A or Box B.

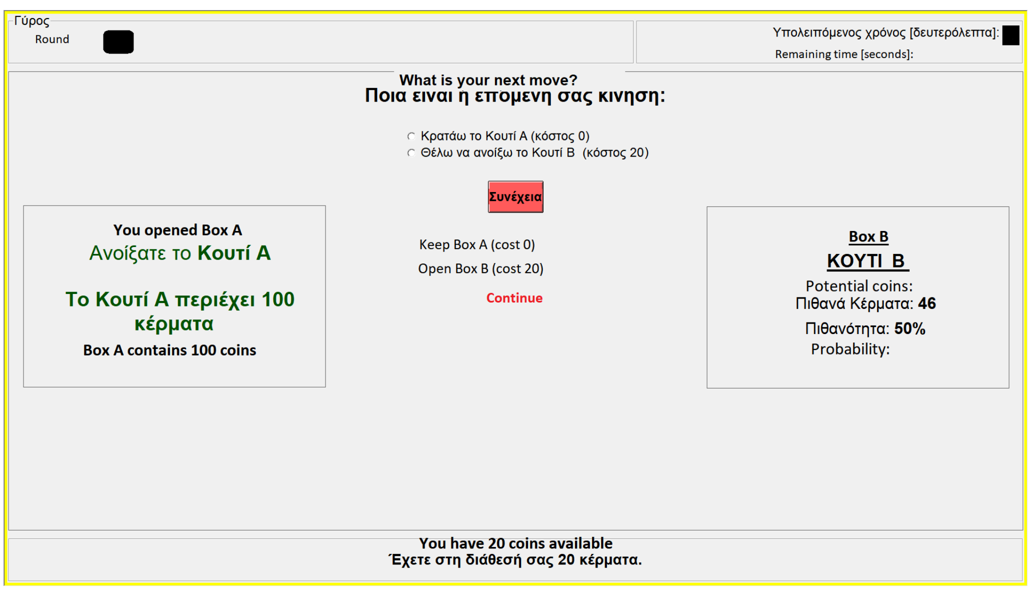

Figure A6: Strict search conditions: Stage 2. The subject opened Box A and its content was revealed. The subject can keep the inspected box and receive its content or proceed with inspecting Box B.

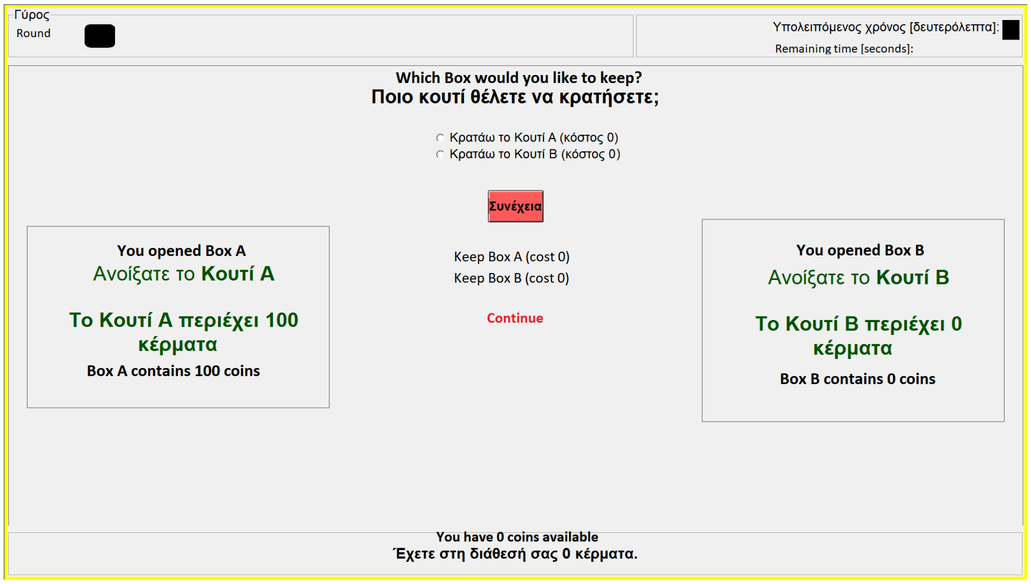

Figure A7: Strict search conditions: Stage 2. The subject also opened Box B and its content was revealed. The subject can keep either Box A or Box B.

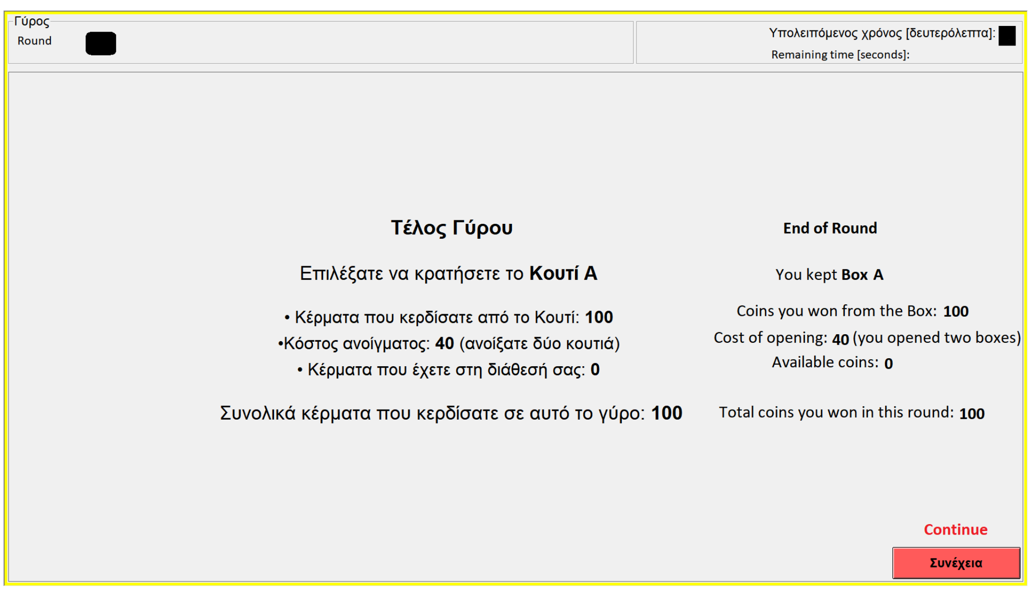

Figure A8: Strict search conditions: End of round. This is a summary of the round based on the subject's choices.

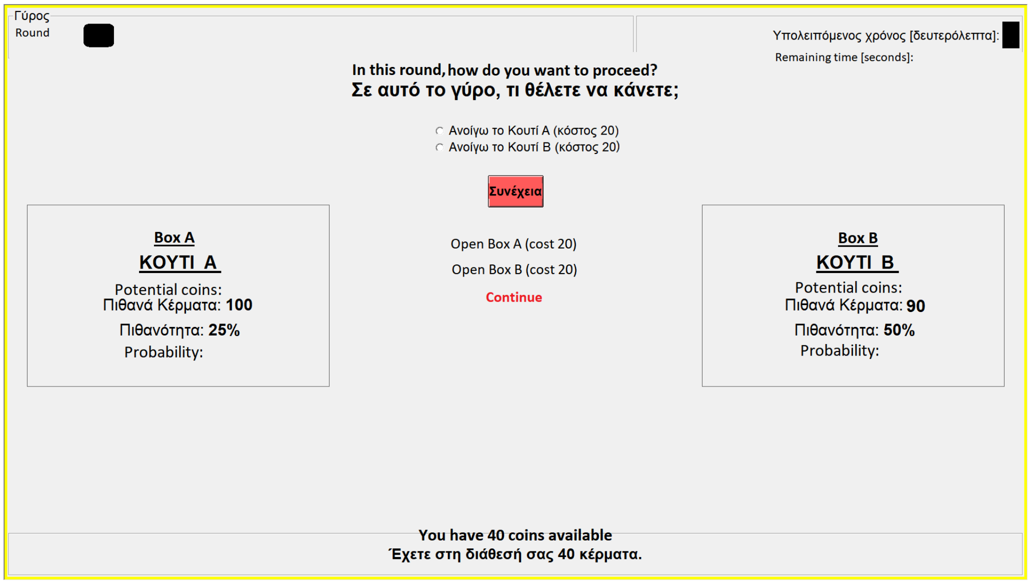

Figure A9: Flexible search conditions: Stage 1. The subject is asked to open a box, Box A or Box B.

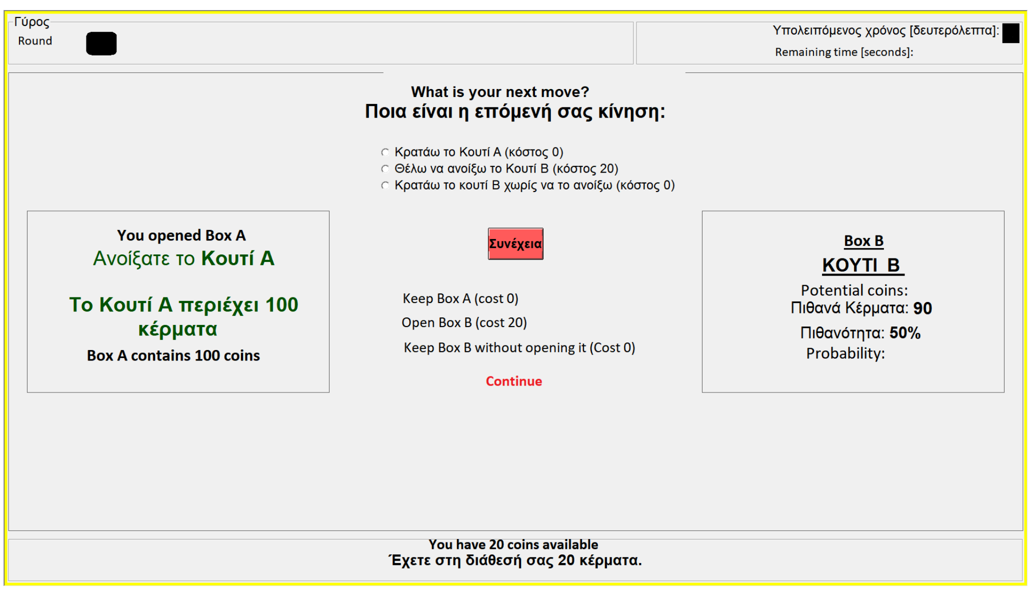

Figure A10: Flexible search conditions: Stage 2. The subject opened Box A and its content was revealed. The subject can keep the inspected box and receive its content, proceed with inspecting Box B, or take Box B without first opening it.

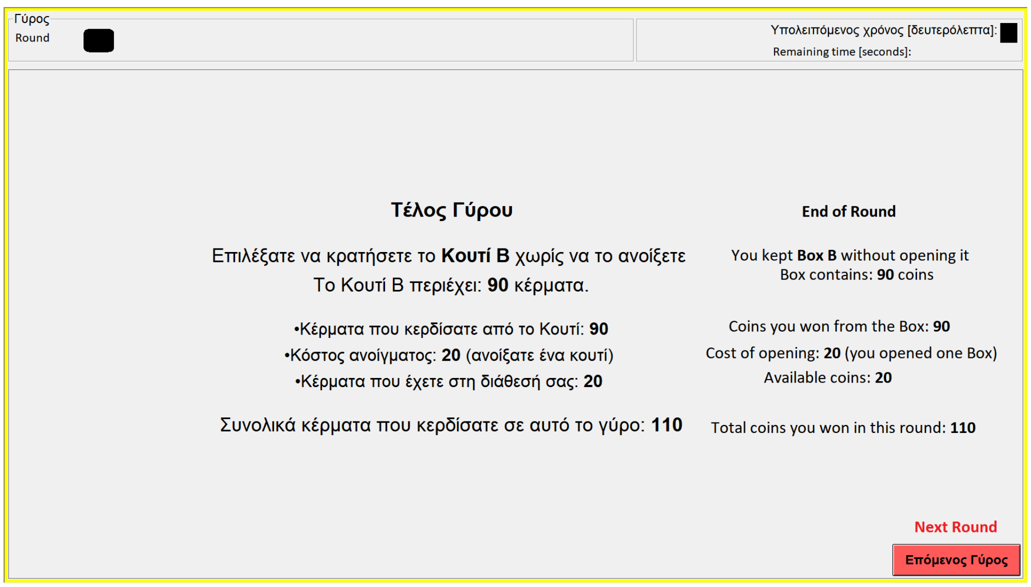

Figure A11: Flexible search conditions: End of round. This is a summary of the round based on subject's choices.

D Additional checks

Table A1: Opened boxes across treatments

Dependent Variable: Opened boxes	(1)
Treatment	-0.3810***
	(0.035)
Observations	6.000
R-squared	0.234

Notes: Standard errors clustered on a subject level are in parentheses. *** denote statistical significance at the 1% level. *Treatment*=0 for strict search conditions and *Treatment*=1 for flexible search conditions. Round dummies and a constant term are included.

Table A2: Total and average payoffs

Time	Rounds 1-100		Rounds 51-100		Rounds 1-100		Rounds 51-100	
Treatment	Strict	Flexible	Strict	Flexible	Strict	Flexible	Strict	Flexible
Session 1	5561.00	7049.73	2817.67	3571.27	370.73	469.98	187.84	238.08
Session 2	5486.73	6928.40	2731.20	3463.47	365.78	461.89	182.08	230.90
Average	5523.87	6989.07	2774.43	3517.37	368.26	465.94	184.96	234.49

Notes: Columns 2-5 correspond to total payoffs per treatment while columns 6-9 correspond to per-subject average payoffs.

Table A3: Threshold test

Rounds	Threshold Value	LM-test statistic	p-value
1-100	67	510.05	0.000
51-100	68	245.40	0.000

Notes: We test the existence of a threshold of the potential content of the safe box (X) on first inspection (*safeboxfirst*), controlling for round effects and including a constant term, against the alternative of no threshold using Hansen (2000). Number of bootstrap replications = 1000.

E Additional experimental sessions

Table A4: Summary of the additional experimental sessions

Treatment	Fixed X values	Subjects	Observations
Strict	50	10	1000
Strict	90	10	1000
Flexible	50	10	1000
Flexible	90	10	1000

Notes: X values refer to the potential content of the safe box.

Table A5: Average payoffs across treatments

Treatment	Strict_{50}	Strict_{90}	Flexible_{50}	Flexible_{90}
Rounds 1-100	4568	6337	6092	7577
Rounds 51-100	2329	3136	3046	3809

Table A6: Subject-level tests: Payoffs across treatments

	strict_{50} versus flexible_{50}	
Tests	t-test	Wilcoxon rank-sum test
Rounds 1-100	-12.07 (0.000)	-3.78 (0.000)
Rounds 51-100	-7.78 (0.000)	-3.71 (0.000)

	strict_{90} versus flexible_{90}	
Tests	t-test	Wilcoxon rank-sum test
Rounds 1-100	-5.45 (0.000)	-3.63 (0.000)
Rounds 51-100	-4.47 (0.000)	-3.25 (0.001)

Notes: For the one-sided t-test we report t-statistics and for the Wilcoxon rank-sum test we report z-statistics. P-values are in parentheses. All tests are based on unpaired data on a subject level.

Table A7: Payoffs across treatments

Dependent Variable: Payoff	(1)	(2)
Treatment	13.8200***	13.8200***
	(2.935)	(1.287)
SafeBox		0.4068***
		(0.032)
Observations	4.000	4.000
R-squared	0.050	0.083

Notes: Standard errors clustered on a subject level are in parentheses. *** denote statistical significance at the 1% level. *SafeBox* refers to the potential content of the safe box. Treatment takes the value 0 for strict_{50} and strict_{90} and equals 1 for flexible_{50} and flexible_{90}. Round dummies and a constant term are included in all specifications.

Table A8: Success rate across treatments

Dep. Var:	First move			
	Strict_{50}	Strict_{90}	Flexible_{50}	Flexible_{90}
Rounds	-0.0005	0.0010*	0.0025	0.0043**
	(0.001)	(0.000)	(0.001)	(0.002)
Observations	1,000	1,000	1,000	1,000
R-squared	0.001	0.005	0.048	0.064

Notes: Standard errors clustered on a subject level are in parentheses. ** and * denote statistical significance at the 5% and 10% level, respectively. A constant term is included in all specifications.

Table A9: Marginal Effects

Dependent Variable: safeboxfirst	(1)	(2)
Treatment	-0.4143***	-0.3487***
	(0.0548)	(0.0874)
Observations	2,000	2,000

Notes: Table A9 reports the average marginal effects of the variables of interest after estimating a probit model with *safeboxfirst* as the dependent variable and *Treatment* and round dummies as covariates. In Column (1), *Treatment*=0 for Strict_{50} and *Treatment*=1 for Flexible_{50}. In Column (2), *Treatment*=0 for Strict_{90} and *Treatment*=1 for Flexible_{90}. *safeboxfirst*=0 if the risky box is opened first and *safeboxfirst*=1 if the safe box is opened first. Delta-method robust standard errors are in parentheses. *** denote statistical significance at the 1% level.