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

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

$$E^{\sigma}_{rstop} = \frac{p_r\Big((c+X_r)^{1-\sigma}-1\Big)}{1-\sigma} + \frac{(1-p_r)\Big(c^{1-\sigma}-1\Big)}{1-\sigma} \tag{2}$$

$$E^{\sigma}_{sr} = \frac{p_s\Big((c+X_s)^{1-\sigma}-1\Big)}{1-\sigma} + \frac{(1-p_s)p_r\Big(X_r^{1-\sigma}-1\Big)}{1-\sigma} \tag{3}$$

$$E^{\sigma}_{sstop} = \frac{p_s\Big((c+X_s)^{1-\sigma}-1\Big)}{1-\sigma} + \frac{(1-p_s)\Big(c^{1-\sigma}-1\Big)}{1-\sigma} \tag{4}$$

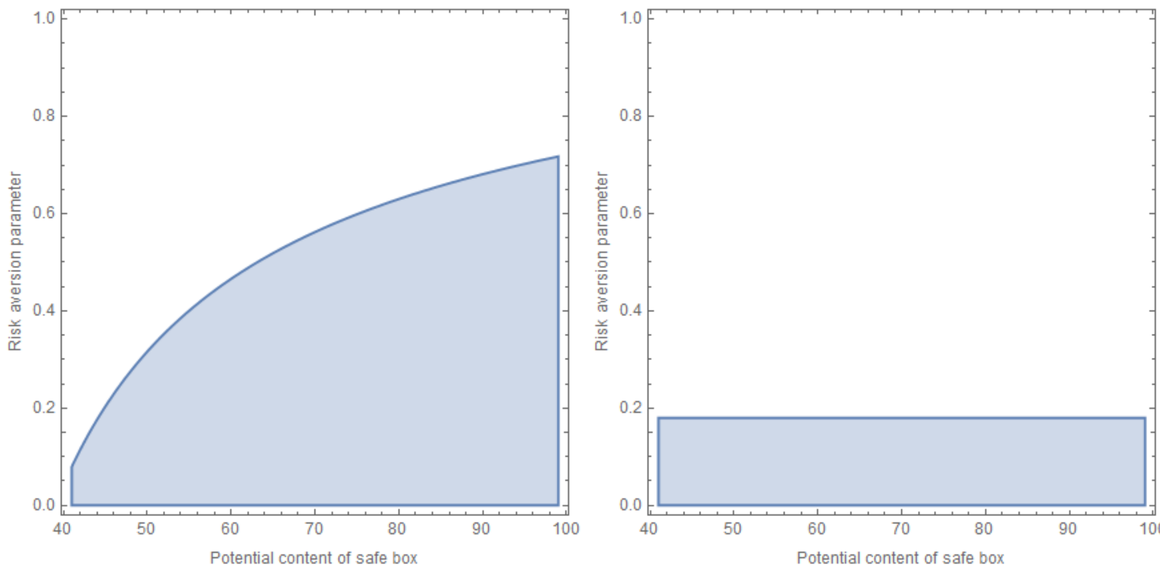

Figure A1: Left panel: Region where $E^{\sigma}_{rs} > E^{\sigma}_{rstop}$. Right panel: Region where $E^{\sigma}_{sr} > E^{\sigma}_{sstop}$. Both panels refer to strict search conditions

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.

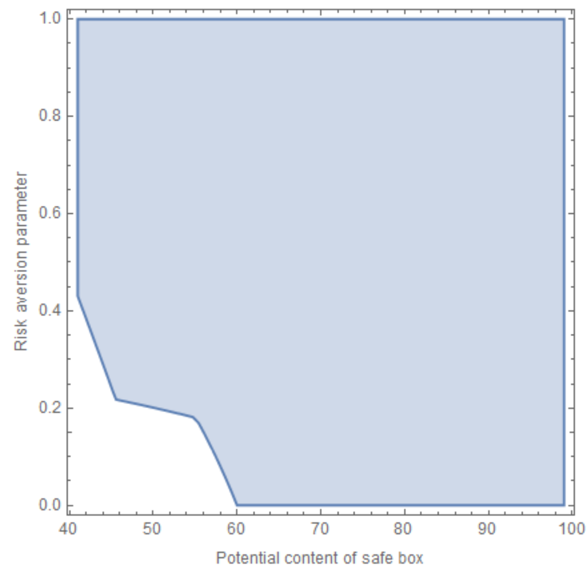

Figure A2: Region where inspecting the safe box first dominates inspecting the risky box first under strict search conditions

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

$$E_{rs}^{\sigma\,\prime} = \frac{p_r\Big((c+X_r)^{1-\sigma}-1\Big)}{1-\sigma} + \frac{(1-p_r)p_s\Big((c+X_s)^{1-\sigma}-1\Big)}{1-\sigma} \tag{5}$$

and

$$E_{sr}^{\sigma\,\prime} = \frac{p_s\Big((c+X_s)^{1-\sigma}-1\Big)}{1-\sigma} + \frac{(1-p_s)p_r\Big((c+X_r)^{1-\sigma}-1\Big)}{1-\sigma} \tag{6}$$

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.[1] The results from our numerical analysis can be seen as confidently demonstrating that risk aversion should amplify the expected outcome

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

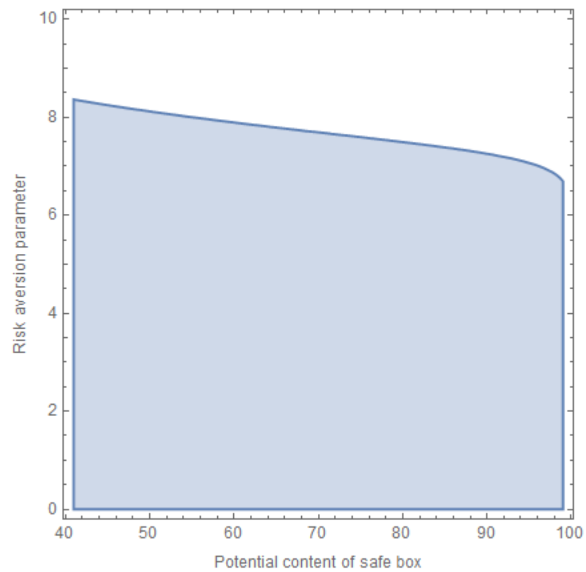

Figure A3: Region where inspecting the risky box first dominates inspecting the safe box first under flexible search conditions

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.

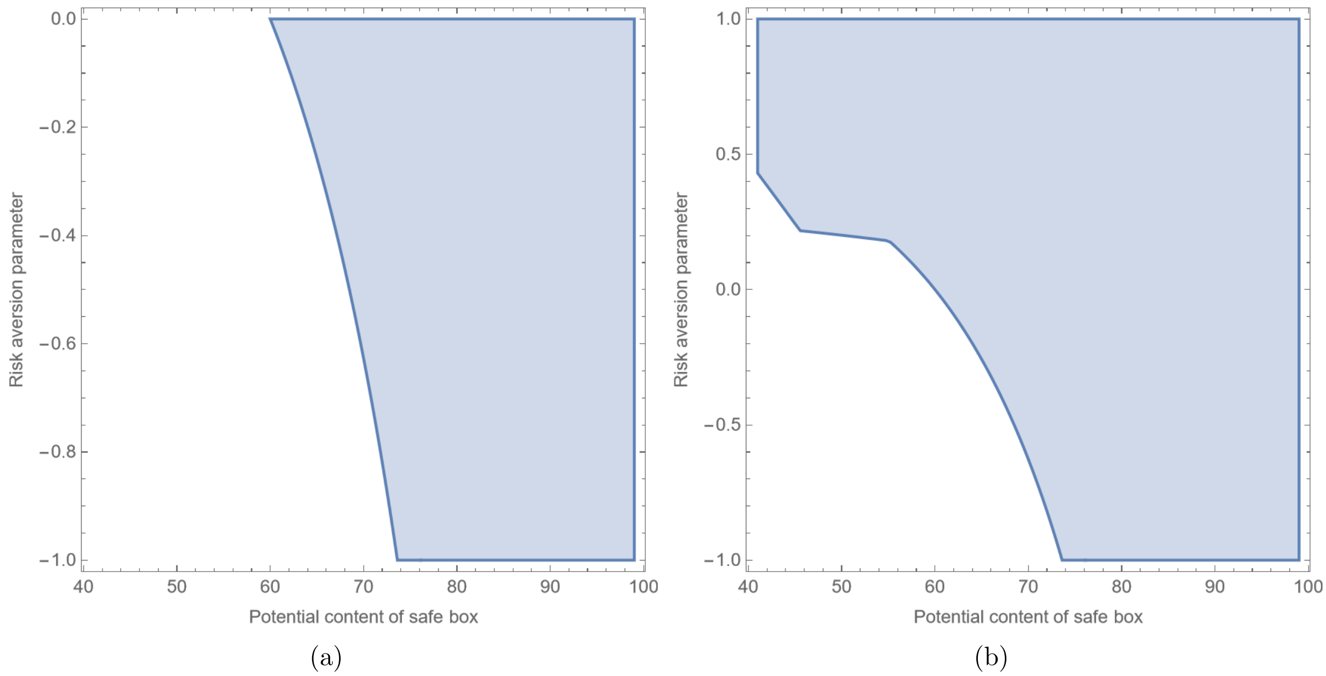

Figure A4: Region where inspecting the safe box first dominates inspecting the risky box first under strict search conditions

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)

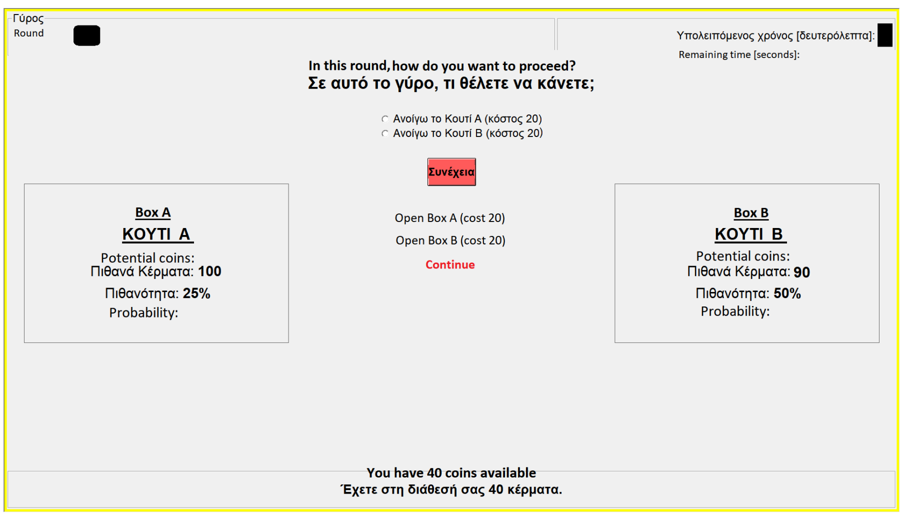

The procedure

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)

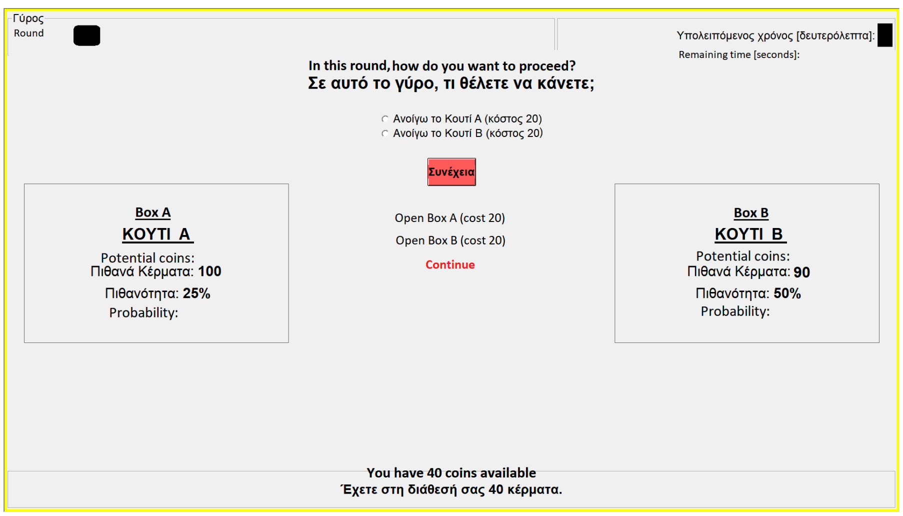

The procedure

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.

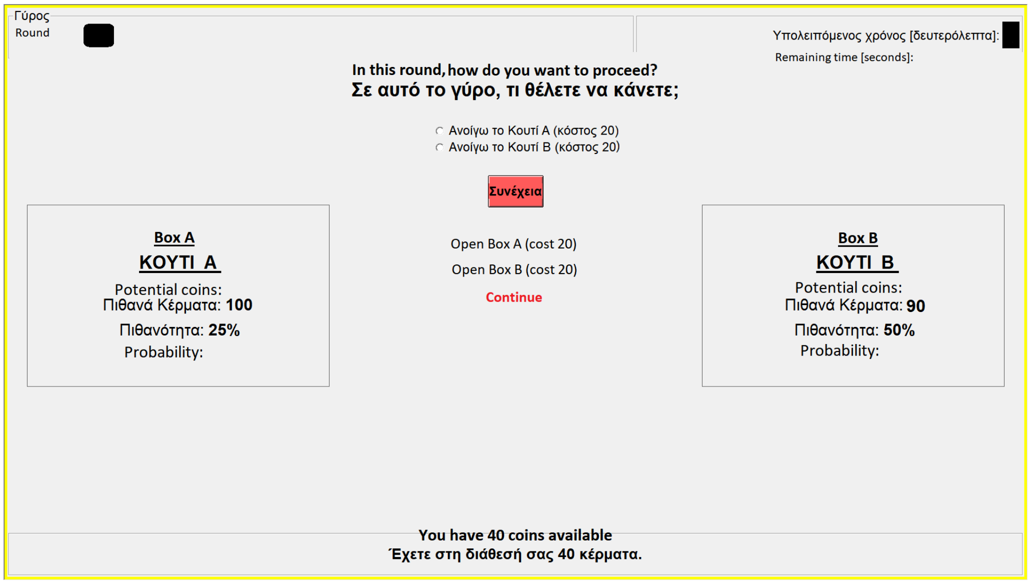

Figure A5: Strict search conditions: Stage 1. The subject is asked to open a box, Box A or Box B.

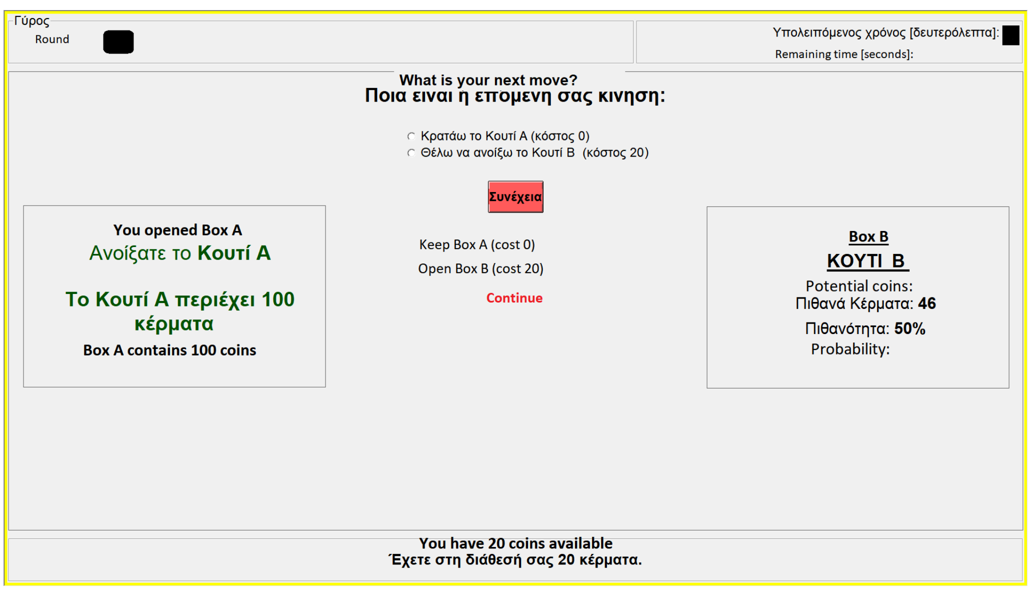

Figure A6: Strict search conditions: Stage 2. The subject opened Box A and its content was revealed. The subject can keep the inspected box and receive its content or proceed with inspecting Box B.

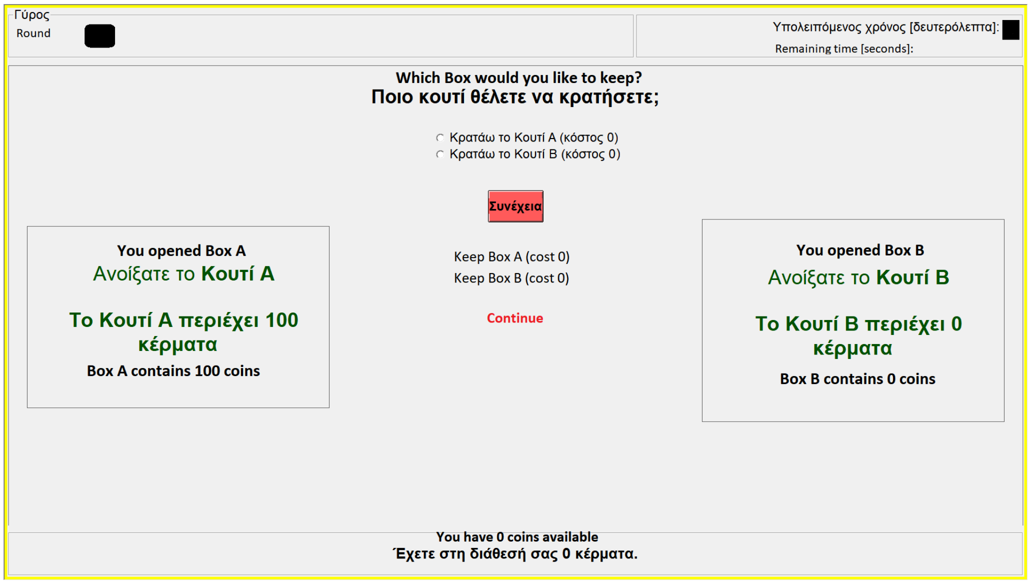

Figure A7: Strict search conditions: Stage 2. The subject also opened Box B and its content was revealed. The subject can keep either Box A or Box B.

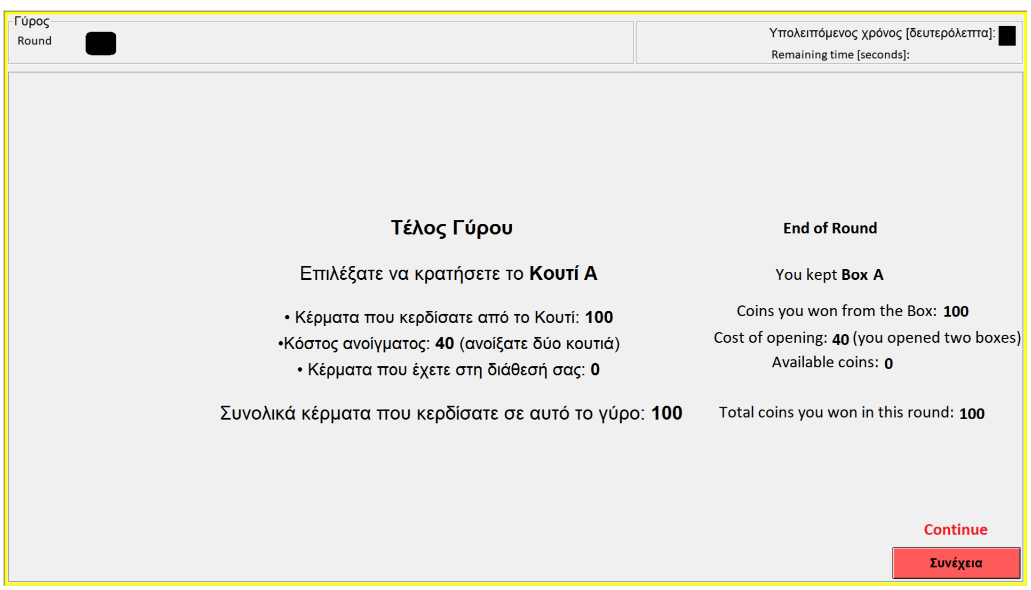

Figure A8: Strict search conditions: End of round. This is a summary of the round based on the subject's choices.

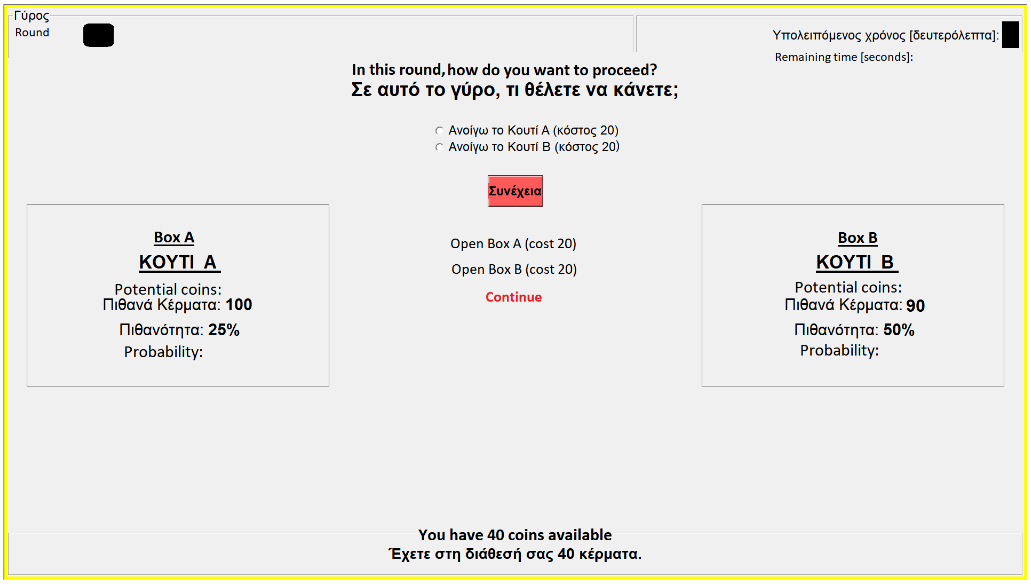

Figure A9: Flexible search conditions: Stage 1. The subject is asked to open a box, Box A or Box B.

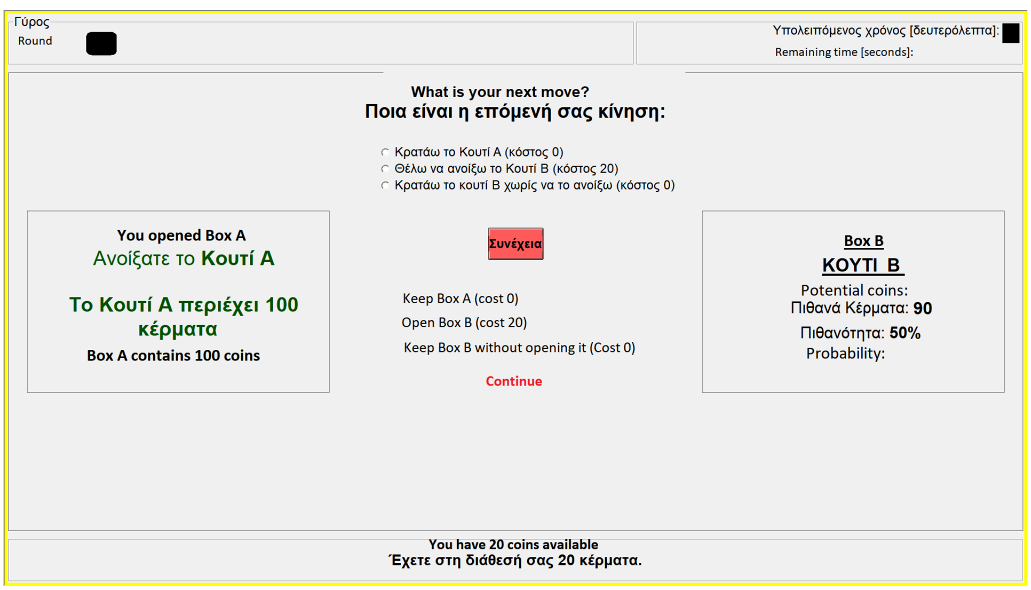

Figure A10: Flexible search conditions: Stage 2. The subject opened Box A and its content was revealed. The subject can keep the inspected box and receive its content, proceed with inspecting Box B, or take Box B without first opening it.

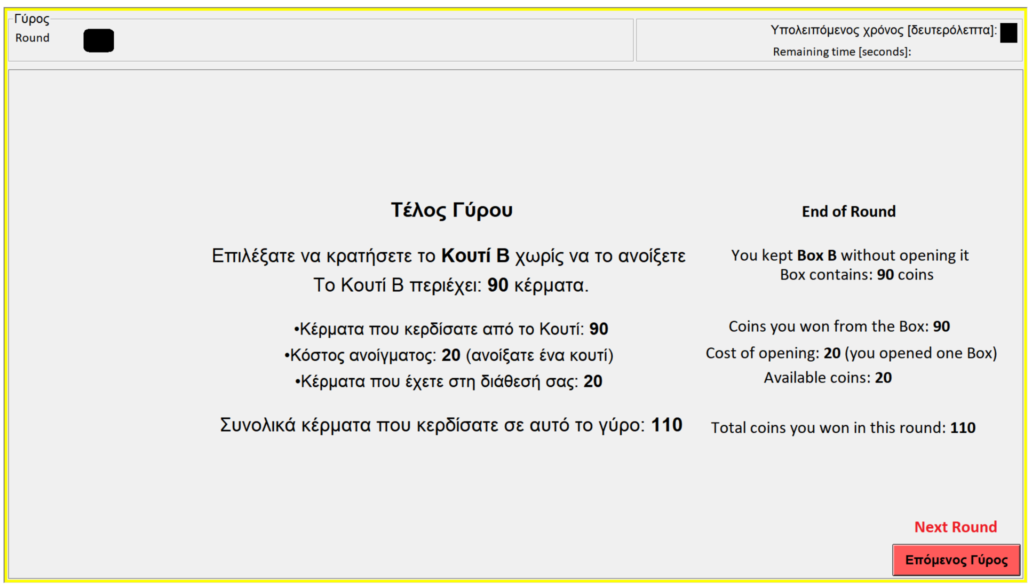

Figure A11: Flexible search conditions: End of round. This is a summary of the round based on subject's choices.

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.

References

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.