Supplement to “Bandits in the Lab”
by Johannes Hoelzemann & Nicolas Klein

A Analysis of Individual Games

In Section 5, we have presented our aggregate results for all six games. We now conduct a separate analysis of the several games, which differed in the realizations of the underlying random processes we simulated ahead of time, as Figures 1 – 3 show. Indeed, insights that hold in all, or most, of these six games might be considered more robust than results that held only on average over the games.

A.1 Experimentation Intensity

In order to illustrate subjects’ dynamically evolving incentives for public-good provision, Figure 1 displays the evolution of each player’s cumulated experimentation intensity over time in Game 1.[1] In the strategic treatment, increasing and flat parts at level 1, of a player’s curve correspond to periods in which the player actively provides information to the group by exploring the risky arm. By contrast, the player relies on his partner’s experimentation efforts when the curve is decreasing or flat at level 0.

Figure 1: Evolution of Cumulated Experimentation Intensity over Time, by Player

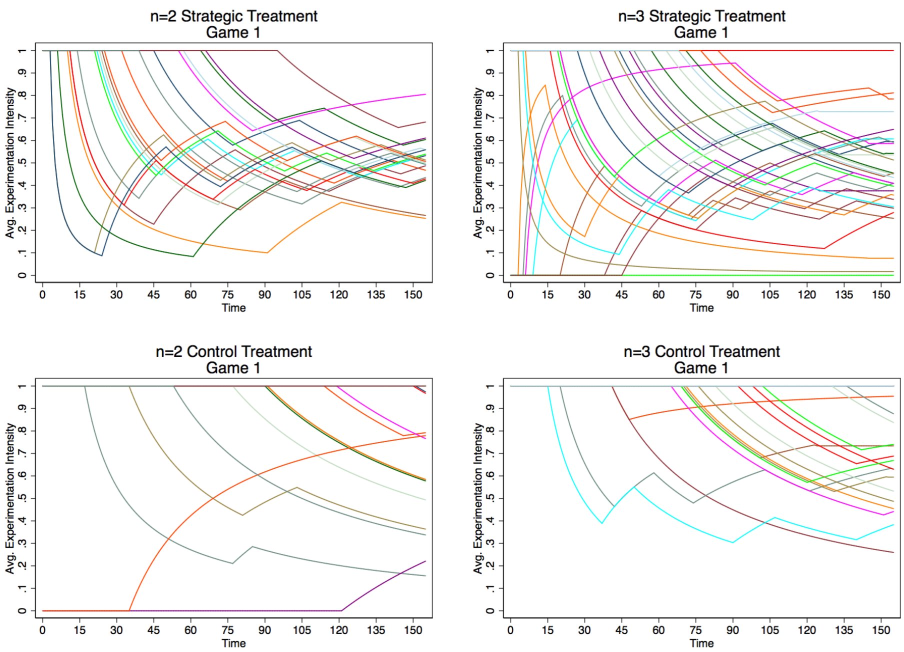

Experimentation Intensity for each subject.

The figure shows that, when players are still optimistic at the start of the game, they overwhelmingly tend to play risky. This is followed by a period in which subjects tended to alternate between safe and risky, with the safe action becoming more frequent toward

[1] For all games please see Figure 4 in the Appendix.

the end. Behavior in the control treatment, however, provides a sharp contrast, as most curves are monotonically decreasing, indicating cut-off behavior.

Also at the individual game level, the additional presence of one (two) perfectly positively correlated arms leads to lower experimentation intensities in all games. When considering all belief regions of a game, this is statistically significant for Games 1-5, but not for Game 6, in both settings with $n = 2$ and $n = 3$. The corresponding p-values in the case of $n = 2$ are 0.0155, 0.0493, 0.0009, 0.0102, 0.0013, and 0.3748 for Games 1-6, respectively. In the setting with $n = 3$, the average experimentation intensity is also lower in the strategic treatment (p-values of 0.0019, 0.0081, 0.0011, 0.0007, 0.0013, and 1.0000 for Games 1 to 6, respectively). As Figure 1 (in main text) highlights, Game 6 features an early success by Player 2 after 9 seconds of exploration, as well as successes by Player 1 after 39 and 44 seconds of exploration, respectively.

Table 1: Average Experimentation Intensity by Belief Regions, by Game

		$n = 2$				$n = 3$			
		Strategic Treatment		Control Treatment		Strategic Treatment		Control Treatment	
Game	Belief Region	Obs.	Exp. Intensity	Obs.	Exp. Intensity	Obs.	Exp. Intensity	Obs.	Exp. Intensity
1	All	10	.508 [.065]	10	.730 [.293]	10	.455 [.107]	10	.797 [.217]
2	—	10	.512 [.116]	10	.696 [.283]	10	.543 [.227]	10	.833 [.125]
3	—	10	.565 [.086]	10	.878 [.235]	10	.457 [.169]	10	.866 [.199]
4	—	10	.519 [.120]	10	.678 [.239]	10	.383 [.089]	10	.728 [.183]
5	—	10	.653 [.204]	10	.984 [.072]	10	.596 [.264]	10	.953 [.110]
6	—	10	.810 [.259]	10	.941 [.167]	10	.800 [.314]	10	.857 [.182]
1	R dominant	10	.648 [.217]	10	.835 [.184]	10	.709 [.310]	10	.935 [.133]
2	—	10	.723 [.254]	10	.888 [.192]	10	.649 [.291]	10	.976 [.040]
3	—	10	.617 [.189]	10	.906 [.155]	10	.593 [.303]	10	.906 [.204]
4	—	10	.732 [.261]	10	.880 [.177]	10	.613 [.275]	10	.889 [.218]
5	—	10	.653 [.204]	10	.984 [.051]	10	.596 [.264]	10	.953 [.110]
1	Mutually BR	10	.503 [.171]	10	.726 [.365]	10	.537 [.230]	10	.760 [.273]
2	—	10	.445 [.114]	10	.752 [.350]	10	.549 [.261]	10	.674 [.299]
3	—	10	.589 [.184]	10	.895 [.225]	10	.482 [.204]	9	.884 [.168]
4	—	10	.484 [.128]	10	.732 [.301]	10	.471 [.204]	10	.807 [.189]

Average [st. dev.] experimentation intensity using group averages. For $n = 3$ in the control treatment, only players in nine groups entered the "Mutually BR" region.

We proceed with our analysis by conducting our parameter tests separately by belief region. As player 2 has a success after 9 seconds of using the risky arm, we omit Game 6 from these tables. We furthermore omit Game 5 from the tables for the "mutually BR" region, as this game lasts only 32 seconds, implying that the "mutually BR" region cannot be attained in the control treatment and only lasts for a few seconds in the strategic treatment, if it is attained at all. For Game 3 in the three-player set-up, the missing observation for the "mutually BR" region corresponds to three individual players in one group in the control treatment that have not reached the "mutually BR" region either on account of an early success or because they did not use the risky arm enough. Table 1 summarizes our findings for each game separately by belief region.

Also, at the game level, the average experimentation intensity is substantially lower in the strategic treatment, for *both* belief regions. We first turn to the set-up for $n = 2$, and focus on the "R dominant" region, where the effect is statistically significant at least at the 5%-level in Games 1, 3, and 5 with the p-values of the two-sided Wilcoxon ranksum test amounting to 0.0249, 0.1364, 0.0044, 0.1180, and 0.0013 for Games 1 to 5, respectively. Now, let us consider the "mutually BR" region. Here, the contrast between the strategic and the control treatment is even more pronounced and statistically significant at least at the 5%-level for all games, with the exception of Game 1. The corresponding p-values are 0.1679, 0.0176, 0.0089, and 0.0186 for Games 1-4, respectively. Recall that the "mutually BR" region is not reached in Game 5.

We now turn to the set-up for $n = 3$. When considering the "R dominant" region, we find the difference in average experimentation intensities between the two treatments to be statistically significant as well. The p-values are 0.0823, 0.0138, 0.0097, 0.0215, and 0.0013 for Games 1-5, respectively. The same is true for the "Mutually BR" region, with the exception of Game 2. This is most likely due to an early success by player 3 after only 44 seconds of exploration. The p-values are 0.0603, 0.2395, 0.0023, and 0.0023 for Games 1 - 4, respectively.

It is of interest whether our findings from section 5.2 also hold at the individual game-level, in particular, whether players distinguish between the two belief regions. In the strategic treatment for two-player groups, the difference between the "R dominant" and the "Mutually BR" region is statistically significant with p-values of 0.0340, 0.0152, and 0.0154 for Games 1, 2 and 4, respectively but not for Game 3 where the p-value amounts to 0.7336. This suggests that subjects attempted to play MPE. In the control treatment, where no difference between these two belief regions is predicted to arise, we document p-values of 0.6791, 0.4247, 0.8425, and 0.2937 for Games 1-4.

In the setting with $n = 3$, even though average experimentation intensity decreases when moving from the "R dominant" region to the "Mutually BR" region, no statistical evidence can be established. The p-values are 0.3603, 0.4710, 0.8189, and 0.4473 for Games 1-4, respectively. Thus, the comparison between the two belief regions does not yield any evidence for MPE-type behavior for $n = 3$, whereas it does for $n = 2$. In the control treatment, where no difference between the regions is predicted to arise, we find no statistically significant differences for Games 1, 3, and 4 (p-values of 0.1592, 0.2576, and 0.2145). However, Game 2 is an exception, with behavior across regions exhibiting significant differences (the p-value is 0.0138).

A.2 Payoffs

Due to (positive) informational externalities, strategic interaction is predicted to arise among players in the strategic treatment. As information is a *public good*, the information produced by their partners allows players to make better decisions and hence to secure themselves higher payoffs. As a result, players' payoffs are predicted to be higher on average in the strategic treatment. Table 2 displays the average [st. dev.] final payoffs using group averages across all six games for our four treatments.

Figure 2 displays the evolution of each player's cumulated payoff over time. Positive slopes correspond to periods during which a subject played safe; flat parts indicate hapless risky play, while jumps denote lump sums arriving from the risky arm.[2]

Figure 2: Evolution of Payoffs over Time, by Player

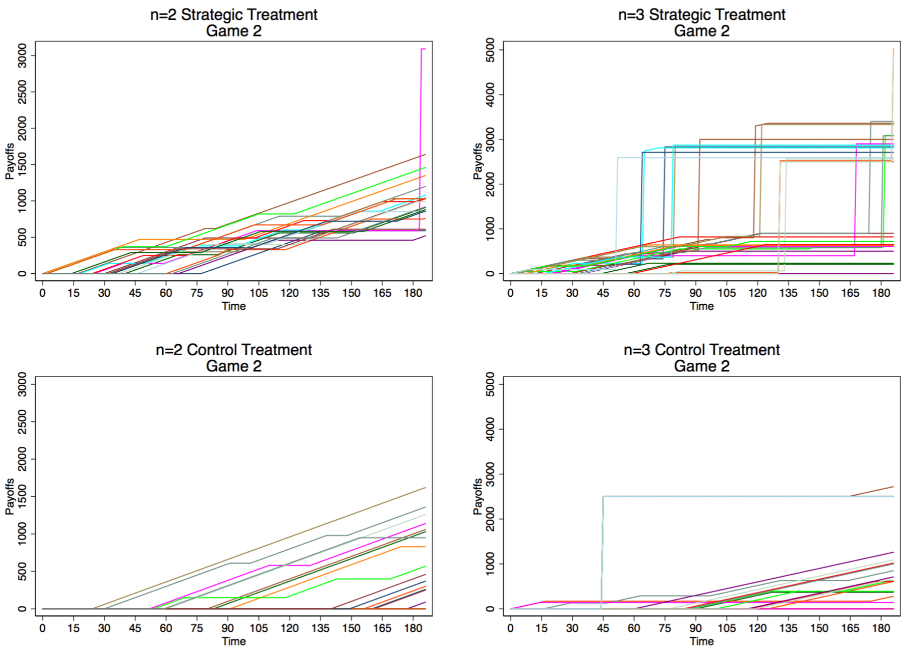

Payoffs over time for each subject.

Table 2 displays average final payoffs per group per game. With the exception of Game 1, average final payoffs are much higher in the strategic treatment than in the control treatment, for both group sizes. This is statistically significant for all Games, with the exception of the two-player groups in Game 6. For $n = 2$ ($n = 3$), the p-values are 0.0342 (0.0820), 0.0233 (0.0004), 0.0007 (0.0015), 0.0081 (0.0007), and 0.0012 (0.0013) for Games 1-5, respectively. The average-payoff difference is not statistically significant with p-value of 0.1145 for $n = 2$ in Games 6; however, such statistical evidence can be established in the setting with $n = 3$, with a p-value of 0.0028. Thus, also at the game-level, our subjects indeed take advantage of the positive informational externalities in the strategic treatment (with the exception of Game 1).

[2] For all games please see Figure 5 in the Appendix.

Table 2: Average Final Payoffs, by Game

	Strategic Treatment				Control Treatment			
Game	Obs.	Final Payoff	Min	Max	Obs.	Final Payoff	Min	Max
Panel A: n = 2								
1	10	817.50 [111.61]	670.00	1060.00	10	1176.50 [440.24]	500.00	1765.00
2	10	1092.00 [452.30]	755.00	2220.00	10	577.50 [446.27]	0.00	1210.00
3	10	407.00 [86.35]	230.00	535.00	10	109.50 [158.03]	0.00	405.00
4	10	1181.00 [279.65]	895.00	1910.00	10	761.00 [460.14]	0.00	1710.00
5	10	115.00 [68.39]	0.00	200.00	10	5.50 [17.39]	0.00	55.00
6	10	3800.50 [69.82]	3750.00	3945.00	10	3554.50 [677.30]	1630	3870.00
Panel B: n = 3								
1	10	1177.67 [365.25]	703.33	1743.33	10	1488.67 [411.23]	610.00	1910.00
2	10	2110.33 [433.68]	1370.00	2686.67	10	1161.00 [232.98]	833.33	1616.67
3	10	496.33 [153.39]	226.67	686.67	10	123.33 [180.88]	0.00	510.00
4	10	1465.33 [209.24]	1166.67	1790.00	10	641.67 [433.23]	0.00	1453.33
5	10	137.00 [88.71]	0.00	250.00	10	15.33 [35.28]	0.00	106.67
6	10	3135.00 [354.57]	2373.33	3363.33	10	2457.33 [433.21]	1276.67	2860.00

Average [st. dev.] final payoffs using group averages.

A.3 Eye-Tracking Data, by Game

Players in the strategic treatment focus much more intensively on their partners' actions and payoffs. Also at the individual game-level, our eye-tracking data further confirms that players were indeed taking advantage of the additional information their partner(s) provided them, a necessary condition for free-riding.

Table 3: Average Fixation Intensity, by Game

	$n = 2$			$n = 3$		
		Strategic Treatment	Control Treatment		Strategic Treatment	Control Treatment
Game	Obs.	Fixation Intensity	Fixation Intensity	Obs.	Fixation Intensity	Fixation Intensity
1	10	.620 [.066]	.870 [.046]	10	.384 [.080]	.710 [.091]
2	10	.620 [.099]	.882 [.085]	10	.365 [.069]	.709 [.119]
3	10	.600 [.050]	.874 [.105]	10	.392 [.079]	.762 [.065]
4	10	.615 [.047]	.875 [.116]	10	.389 [.094]	.700 [.091]
5	10	.633 [.116]	.876 [.105]	10	.383 [.089]	.745 [.129]
6	10	.594 [.125]	.814 [.073]	10	.382 [.070]	.646 [.111]

Average [st. dev.] fixation intensity using group averages. The number of observations refers to both strategic and control treatment.

By contrast, in the corresponding control treatments, where the information generated by their partners is of no value as the risky arms are uncorrelated, subjects seemed to focus almost exclusively on their own stream of payoffs, thus confirming our theoretical prediction according to which a rational player should completely ignore a partner's actions and payoffs in the control treatments.

As Table 3 highlights, the average fixation intensity using group averages is significantly lower in the strategic treatment, irrespective of the group size. This is highly statistically significant for all six games for both group sizes. For $n = 2$ ($n = 3$) the corresponding p-values are 0.0002 (0.0002), 0.0002 (0.0002), 0.0002 (0.0002), 0.0015 (0.0002), 0.0007 (0.0002), 0.0009 (0.0003) for Games 1-6, respectively.

Figure 3 displays (non-representative) heatmaps to illustrate the different information acquisition behavior in our four treatments. The measure of interest is the total number of fixations. For each heatmap, the accumulated number of fixations is calculated for an entire game and the image corresponds to the last point in calendar time before the game ends. A color gradient is employed to display the areas that attained more fixations (low=green to high=red). As Figure 3 illustrates, players not only switch actions more frequently in the strategic treatment but also focus much more intensively on their partners' actions and payoffs. This is in sharp contrast to the corresponding control treatment, where players seem to focus almost exclusively on their own streams of payoffs.

Figure 3: Heatmaps of Four Treatments

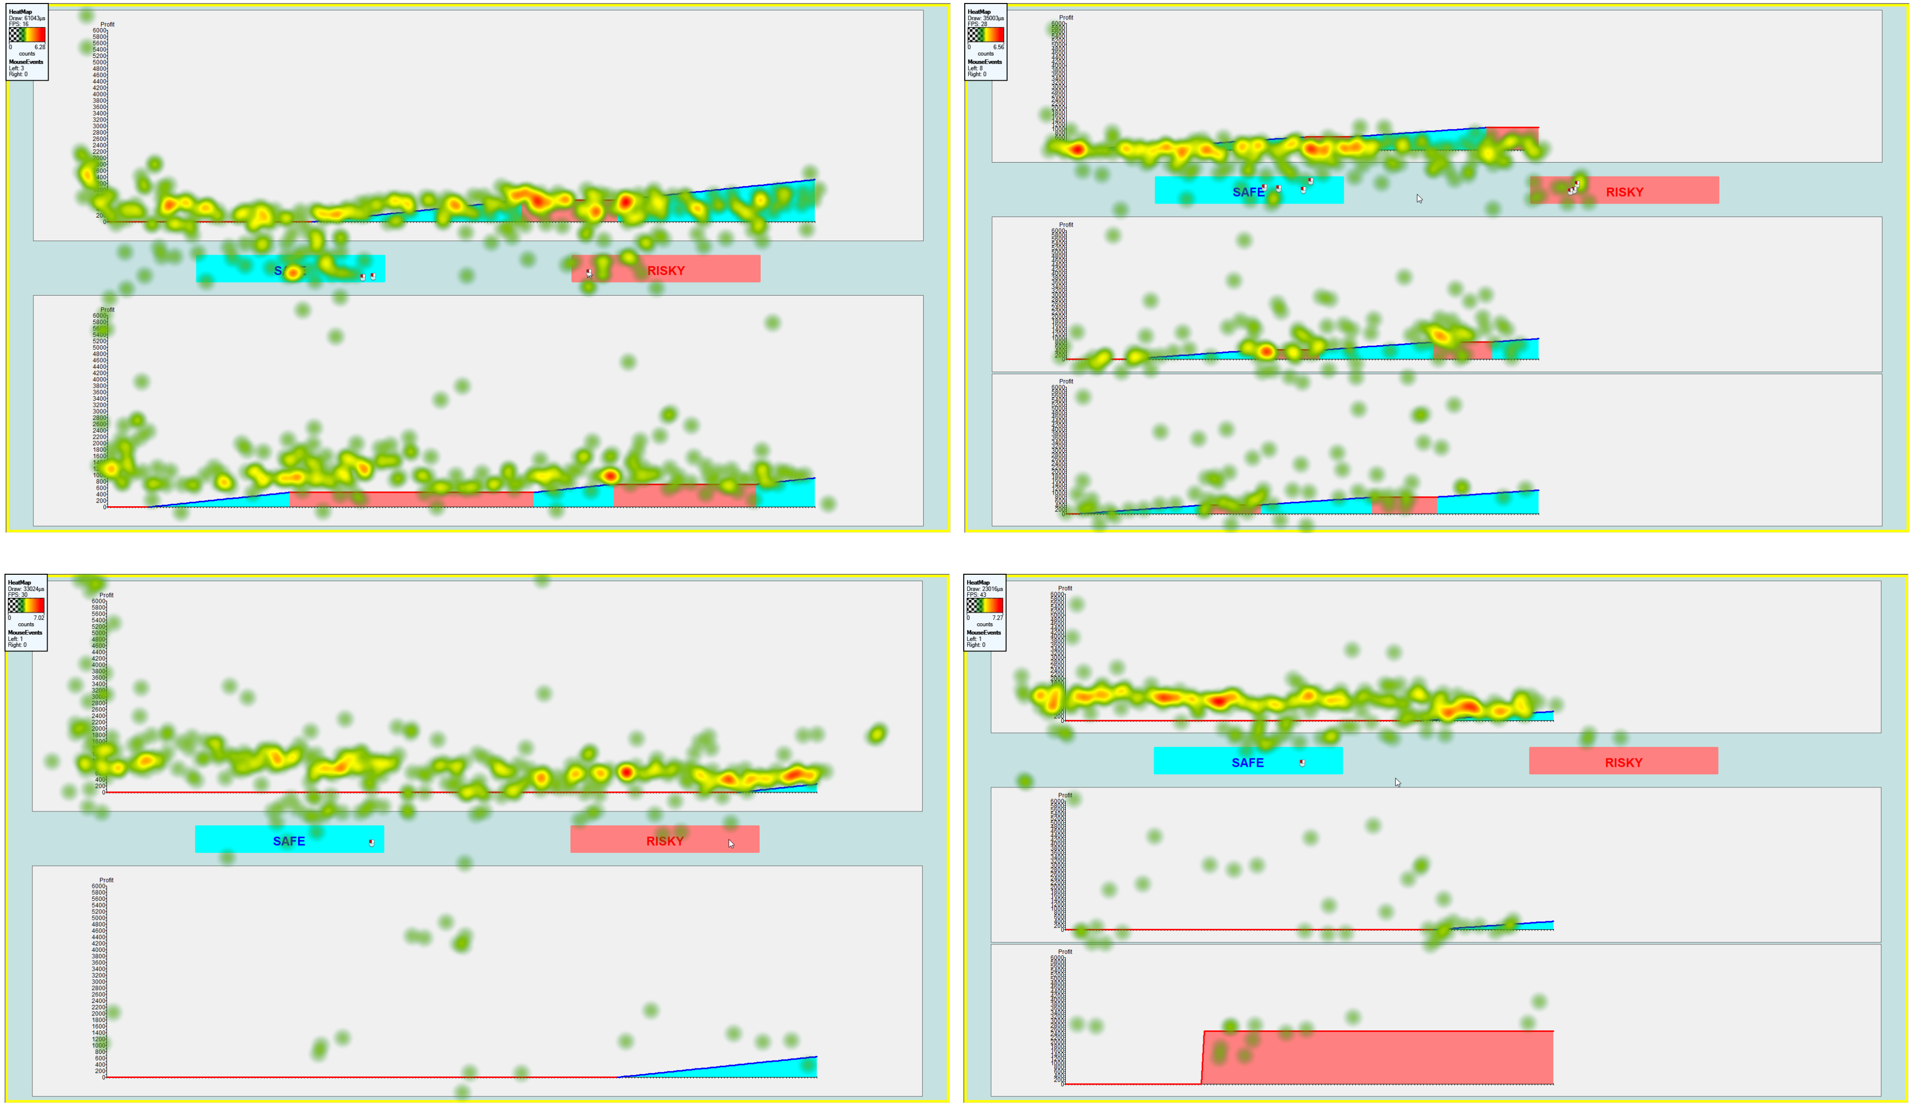

In the top-left corner, the strategic treatment with $n = 2$ is illustrated, with the corresponding control treatment represented just below. In the top-right corner, the strategic treatment with $n = 3$ is displayed, while the control treatment with $n = 3$ is shown at the bottom-right. All four heatmaps show the total number of fixations. The accumulated number of fixations is calculated for an entire game (Game 4 in the $n = 2$ set-up and Game 2 in the $n = 3$ set-up). Each fixation made has the same value and is indepentent of its duration. A color gradient is used to indicate the areas with more fixations (low=green to high=red).

A.4 Cut-Off Behavior

We now turn to the frequency of cut-off behavior. As we have seen in Result 5.5, cut-off behavior is much more frequent in the control treatment than in the strategic treatment for both group sizes. While it increases sharply in Games 5 and 6, as compared to Games 1-4, in the strategic treatments, it is still higher in the corresponding control treatments for either group size. In Game 5, this sharp increase is most likely due to the short duration of that game. In Game 6, it is most likely driven by the resolution of uncertainty very early in the game, with Player 2 achieving a success after exploring for 9 seconds.

Table 4: Frequency of Cut-Off Behavior, by Game

	$n = 2$			$n = 3$		
		Strategic Treatment	Control Treatment		Strategic Treatment	Control Treatment
Game	Obs.	Tot. (Rel.) Freq.	Tot. (Rel.) Freq.	Obs.	Tot. (Rel.) Freq.	Tot. (Rel.) Freq.
1	20	0 (0)	15 (.75)	30	3 (.10)	21 (.70)
2	20	0 (0)	15 (.75)	30	3 (.10)	22 (.73)
3	20	5 (.25)	19 (.95)	30	11 (.37)	26 (.87)
4	20	0 (0)	14 (.70)	30	6 (.20)	19 (.63)
5	20	17 (.85)	20 (1)	30	17 (.57)	29 (.97)
6	20	13 (.65)	17 (.85)	30	19 (.63)	25 (.83)

Total number of cut-offs (number of cut-offs divided by total observations). The number of observations refers to both strategic and control treatment.

We find the difference in cut-off behavior-frequency between the two treatments to be highly statistically significant for Games 1-4, for both group sizes. All p-values are 0.0001 for Games 1-4, respectively, with the exception of Game 4 for $n = 3$ where the p-value amounts to 0.0007. In the last two games where we observe a sharp uplift in cut-off behavior in the strategic treatment for the reasons outlined above, the corresponding p-values for $n = 2$ ($n = 3$) are 0.0754 (0.0003) and 0.1492 (0.0824) for Games 5 and 6, respectively.

A.5 Pioneers

There is a range of beliefs containing $(p_1^*, p^\ddagger)$ such that safe and risky are mutually best responses in any Markov Perfect Equilibrium, so that there exists a range of beliefs in which just *one pioneer* should play risky in MPE while the other player(s) free-ride(s). By contrast, in the control treatment as well as in the best PBE, players are predicted to play risky on $(p_1^*, \frac{1}{2}]$. In this belief region, conditionally on no success arriving, players should switch from risky to safe only once, and do so at the same time, at which their beliefs reach p_1^*. At the game-level too, we confirm Result 5.6.

As Table 5 highlights, also at the individual game level, we can confirm for all games that the addition of one (two) perfectly positively correlated arms leads to a much higher proportion of time where just one pioneer plays risky while the other remaining player(s) free-ride. This is highly statistically significant for all games in the three-player set-up and

Table 5: Proportion of Time with a Single Pioneer, by Game

	$n = 2$			$n = 3$		
		Strategic Treatment	Control Treatment		Strategic Treatment	Control Treatment
Game	Obs.	Single Pioneer	Single Pioneer	Obs.	Single Pioneer	Single Pioneer
1	10	.724 [.156]	.284 [.258]	10	.670 [.178]	.097 [.156]
2	10	.708 [.176]	.315 [.254]	10	.425 [.352]	0 [0]
3	10	.745 [.156]	.187 [.253]	10	.563 [.348]	.136 [.256]
4	10	.757 [.175]	.294 [.214]	10	.741 [.171]	.249 [.198]
5	10	.581 [.360]	.029 [.092]	10	.361 [.304]	0 [0]
6	10	.288 [.399]	.078 [.246]	10	.219 [.369]	0 [0]

Average [st. dev.] proportion of time with a single pioneer in a group. The number of observations refers to both strategic and control treatment.

for Games 1-5, but not for Game 6, in setting with $n = 2$. The corresponding p-values in the case of $n = 2$ are 0.0011, 0.0019, 0.0003, 0.0007, 0.0013, and 0.1494 for Games 1-6, respectively. In the setting with $n = 3$, the average experimentation intensity is also lower in the strategic treatment (p-values of 0.0002, 0.0006, 0.0026, 0.0003, 0.0019, and 0.0682 for Games 1 to 6, respectively). Recall that Game 6 is characterized by an early successes by two players: after 9 seconds of exploration by Player 1 and after 39 and 44 seconds of exploration by Player 1.

A.6 Switches of Action

In any Markov Perfect Equilibrium, we should expect players to switch roles at least once once. As theory predicts and Result 5.7 shows for the aggregate data, significantly more switches are observed in the strategic treatment than in the control treatment, for both group sizes. Recall that we have defined the incidence of switches as the number of a player's changes in action choice in a given game per unit of effective time.

Table 6 displays the average number of switches per player across games for our four treatments. The incidence of switches in the strategic treatment is much higher than in the control treatment in all games (for $n = 2$ with p-values of 0.0001, 0.0003, 0.0001, 0.0002, 0.0019, and 0.1352 for Games 1-6, respectively; in the $n = 3$ setting with p-values are of 0.0040, 0.0005, 0.0073, 0.0336, 0.0018, and 0.3526 for Games 1-6, respectively). Here again, the early success in Game 6 reveals the risky arm to be good and thus resolves all uncertainty at the very beginning of the game.

Table 6: Average Number of Switches per Player, by Game

	$n = 2$			$n = 3$		
		Strategic Treatment	Control Treatment		Strategic Treatment	Control Treatment
Game	Obs.	Switches Per Pl.	Switches Per Pl.	Obs.	Switches Per Pl.	Switches Per Pl.
1	10	4.45 [1.74]	.90 [.66]	10	3.40 [1.77]	1.13 [1.23]
2	10	4.50 [1.87]	1.35 [1.13]	10	2.77 [1.65]	.97 [.81]
3	10	2.20 [1.03]	.30 [.42]	10	1.73 [1.14]	.47 [.69]
4	10	6.05 [1.57]	1.85 [1.56]	10	4.00 [2.82]	1.7 [1.63]
5	10	.60 [.39]	.05 [.16]	10	.70 [.73]	.03 [.11]
6	10	.60 [.74]	.30 [.54]	10	.97 [1.29]	.37 [.55]

Average [st. dev.] switches per player using group averages. The number of observations refers to both strategic and control treatment.

B Interfaces

In this Appendix, we exhibit examples of the interfaces subjects saw during the game, showing the evolution of the screen over intervals of 30 seconds. In the top half (third) of his screen, a subject could see his own past actions and payoffs, while the bottom half (two thirds) of the screen showed his fellow group members' actions and payoffs. A blue (red) part of the payoff curve indicated that the player used the safe (risky) arm over the corresponding period. The x-axis represented calendar time, while the y-axis gave the player's cumulated total earnings up to each point in time. There was no prior indication of the point in time the game would end.

$n = 2$ Strategic Set-up: Example for Game 1

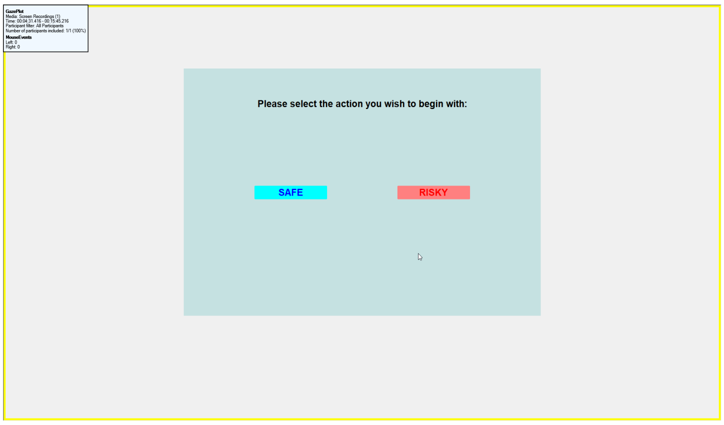

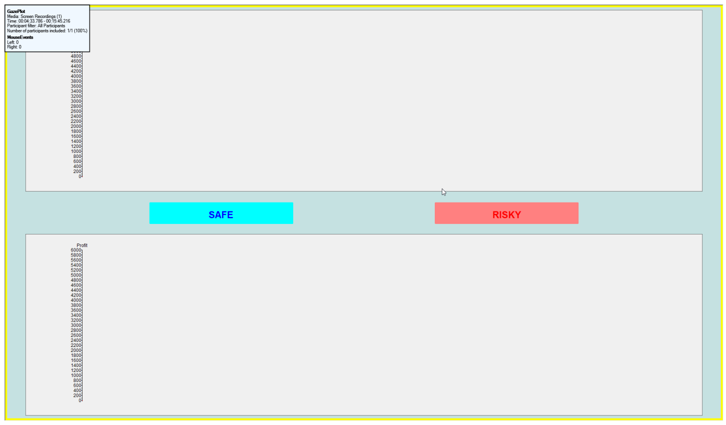

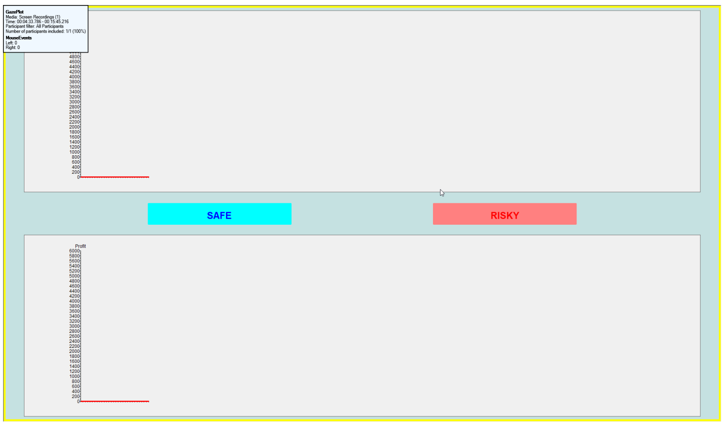

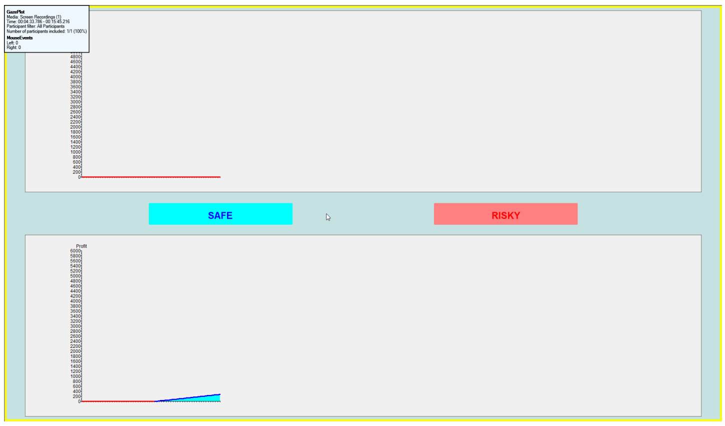

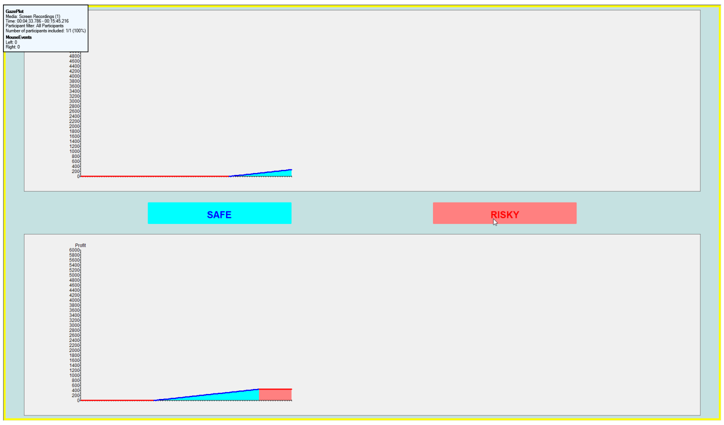
SAFE
RISKY

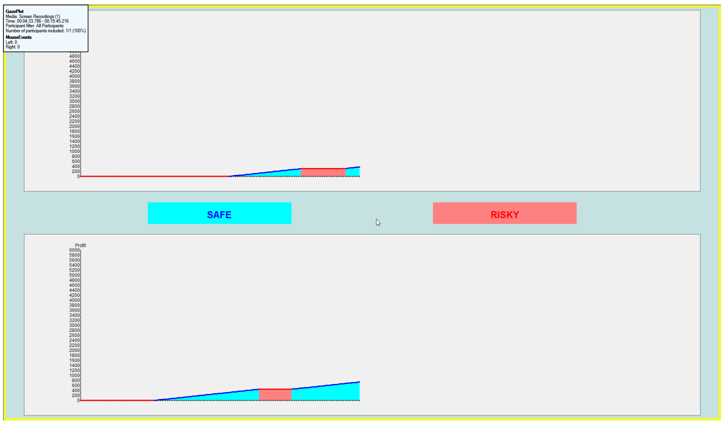
SAFE
RISKY

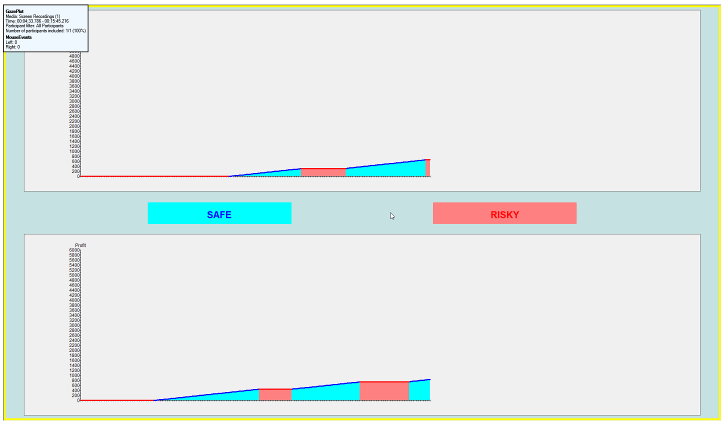
SAFE
RISKY

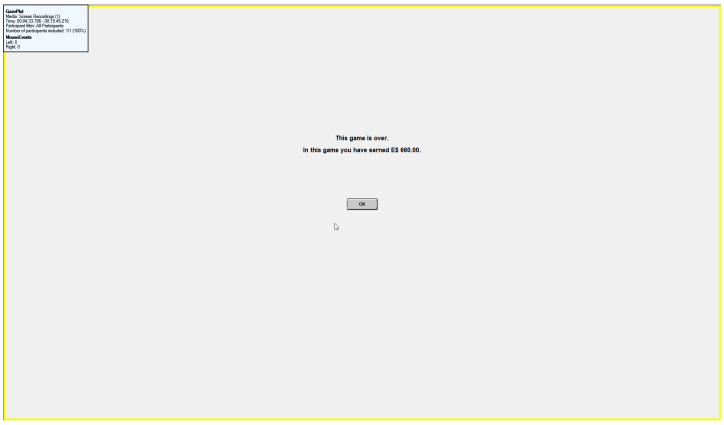
This game is over.
In this game you have earned E$ 660.00.
OK

n = 2 Control Set-up: Example for Game 1

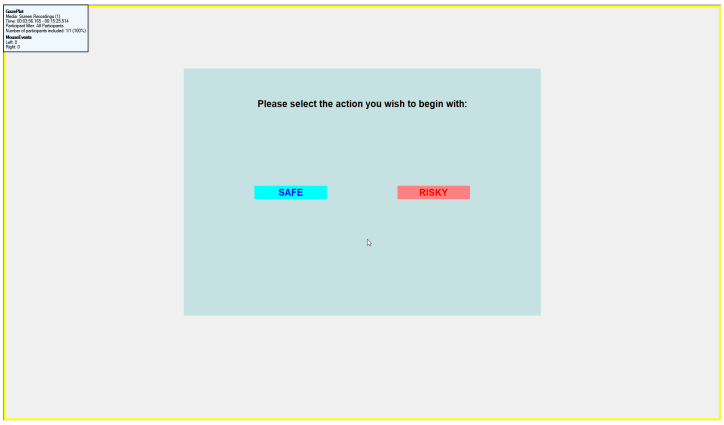

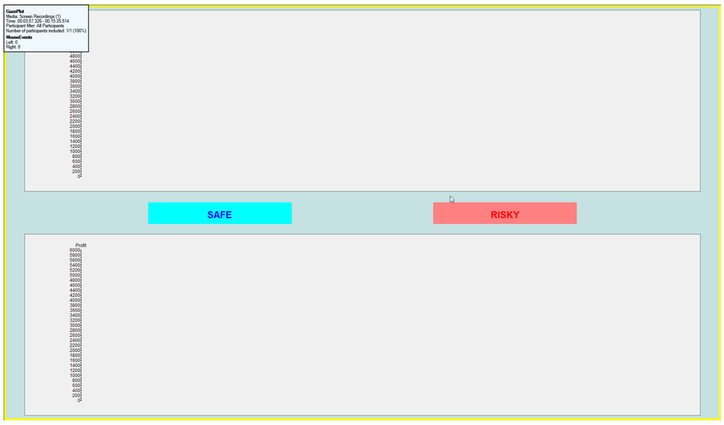

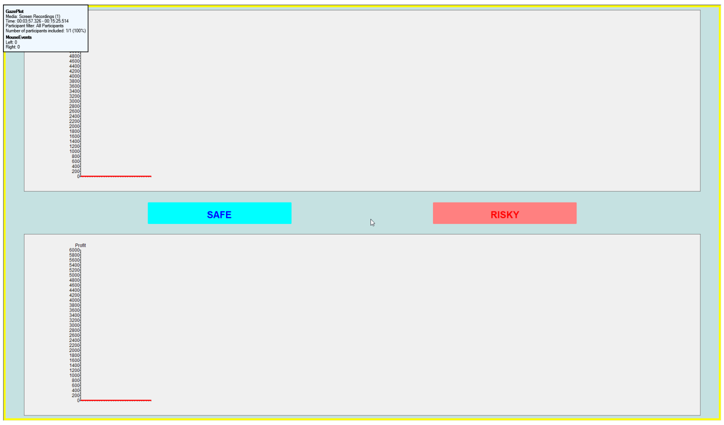

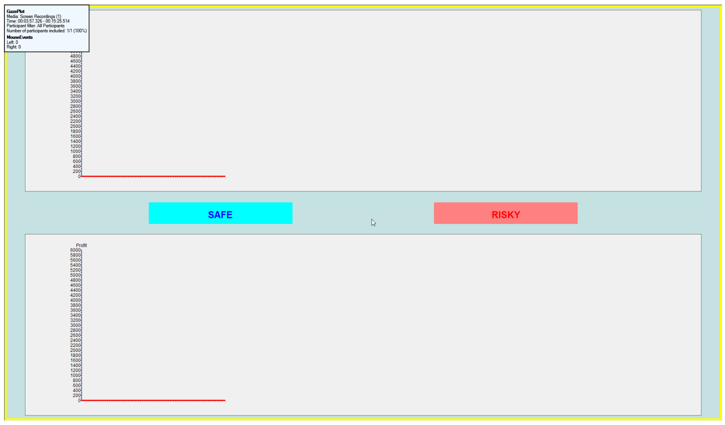

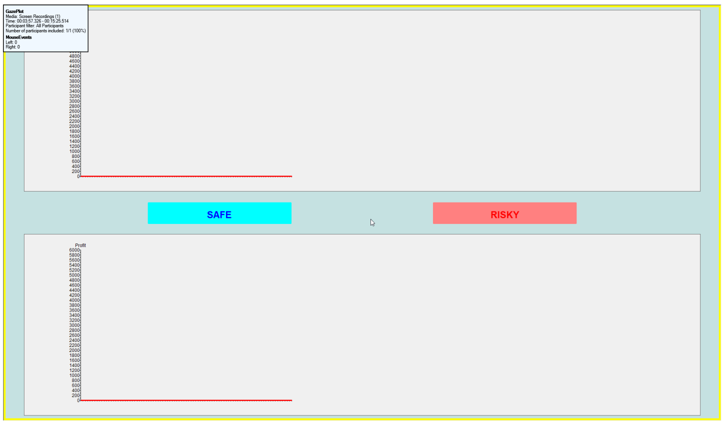
SAFE
RISKY
Profit

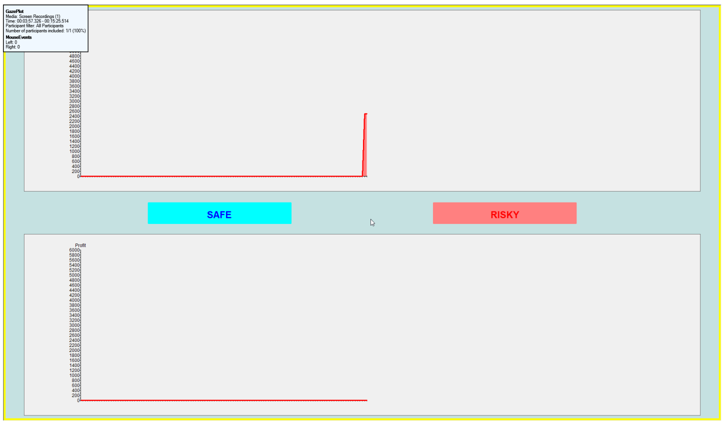
SAFE
RISKY
Profit

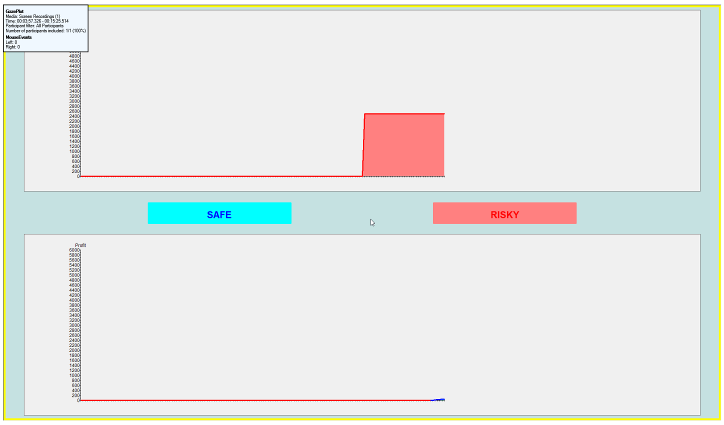
SAFE
RISKY
Profit

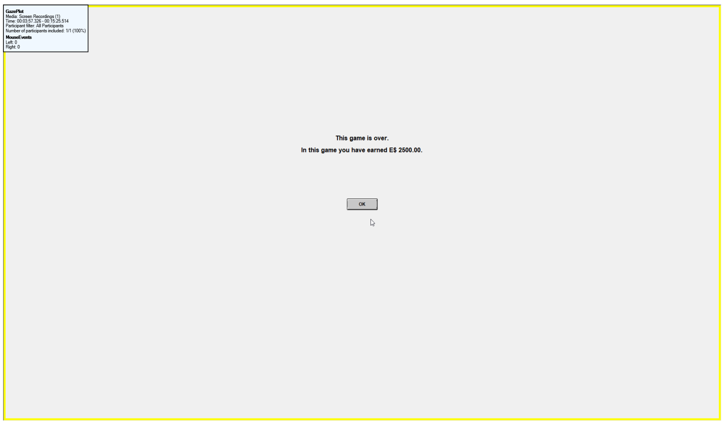
This game is over.
In this game you have earned E$ 2500.00.
OK

n = 3 Strategic Set-up: Example for Game 2

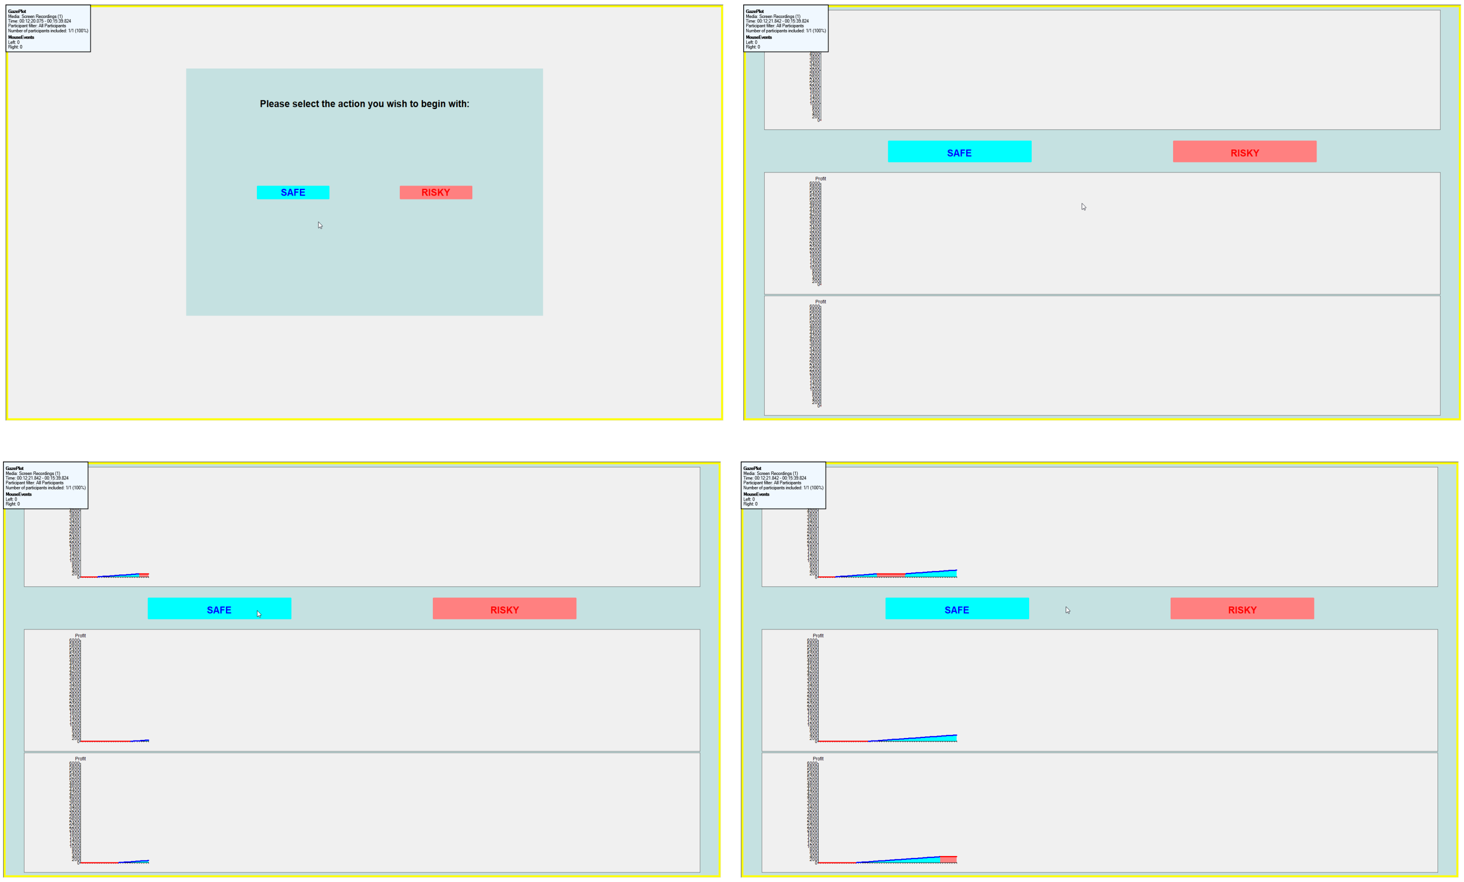


SAFE
RISKY
Profit
Profit


SAFE
RISKY
Profit
Profit

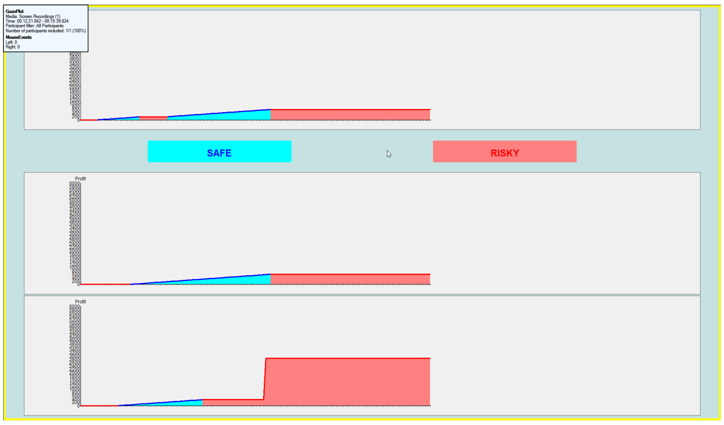
SAFE
RISKY
Profit
Profit

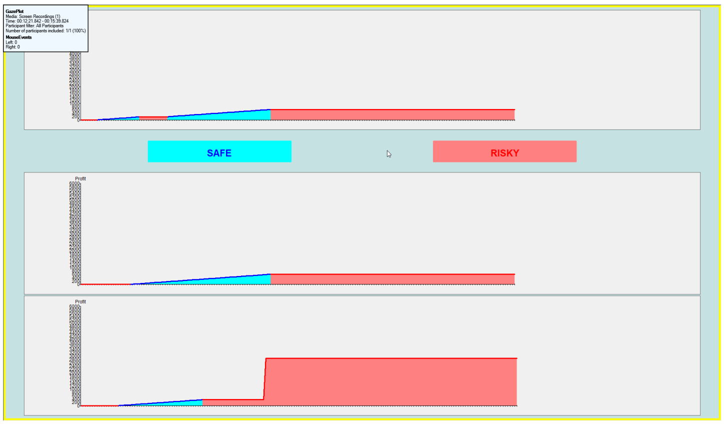
SAFE
RISKY
Profit
Profit

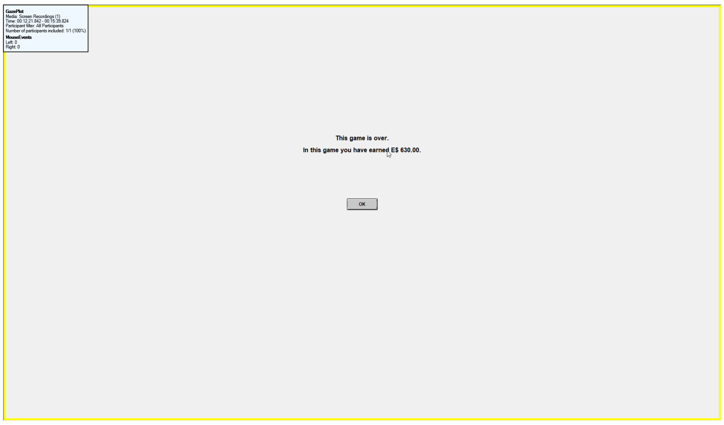
This game is over.
In this game you have earned E$ 630.00.
OK

n = 3 Control Set-up: Example for Game 2

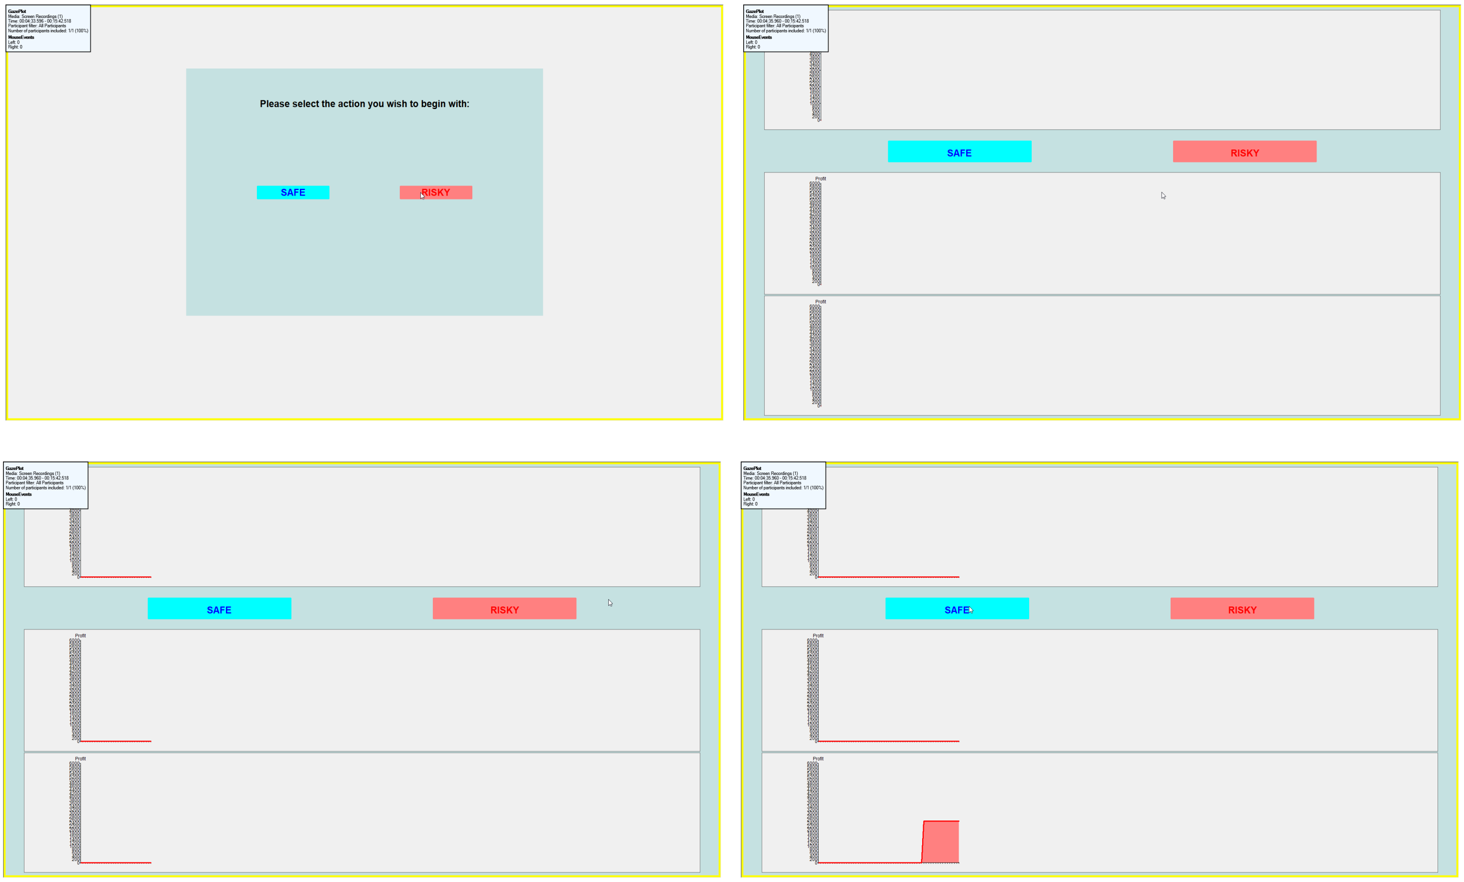

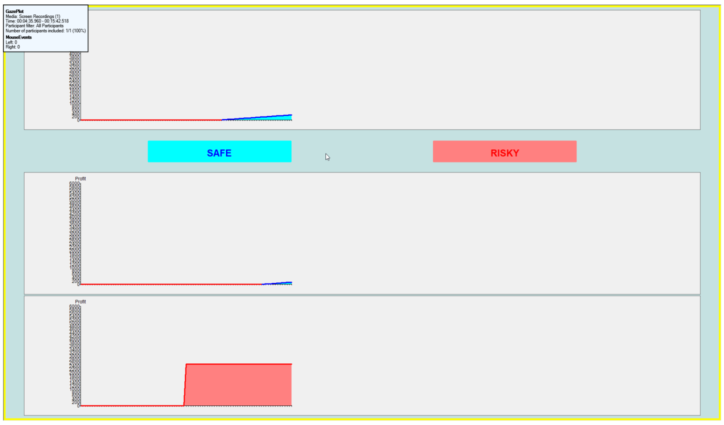
SAFE
RISKY
Profit
Profit

SAFE
RISKY
Profit
Profit
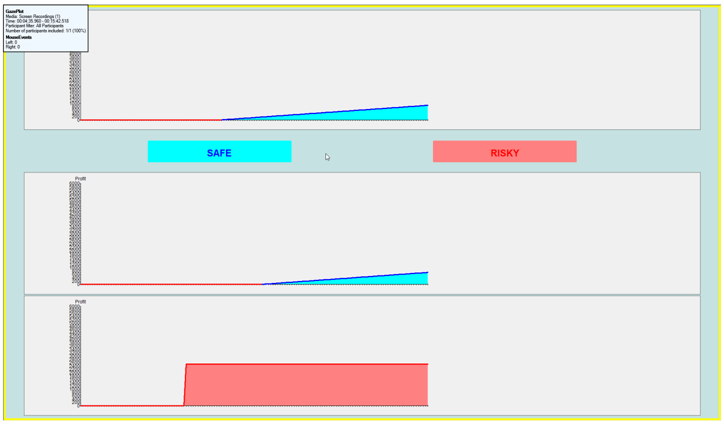
SAFE
RISKY
Profit
Profit
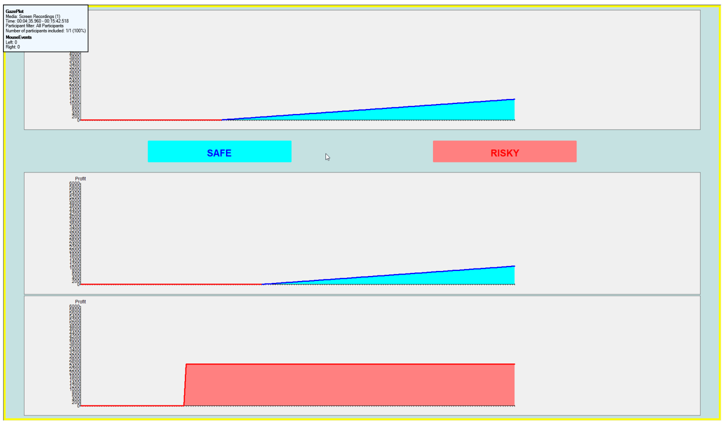
SAFE
RISKY
Profit
Profit
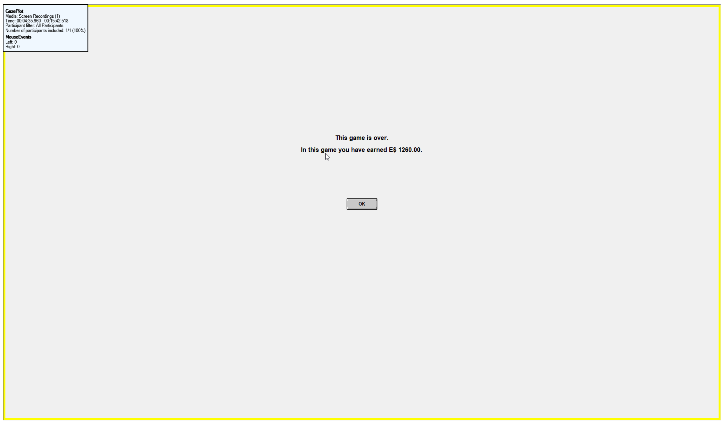
This game is over.
In this game you have earned E$ 1260.00.
OK

Figure 4: Evolution of Cumulated Experimentation Intensity over Time by Player, by Game

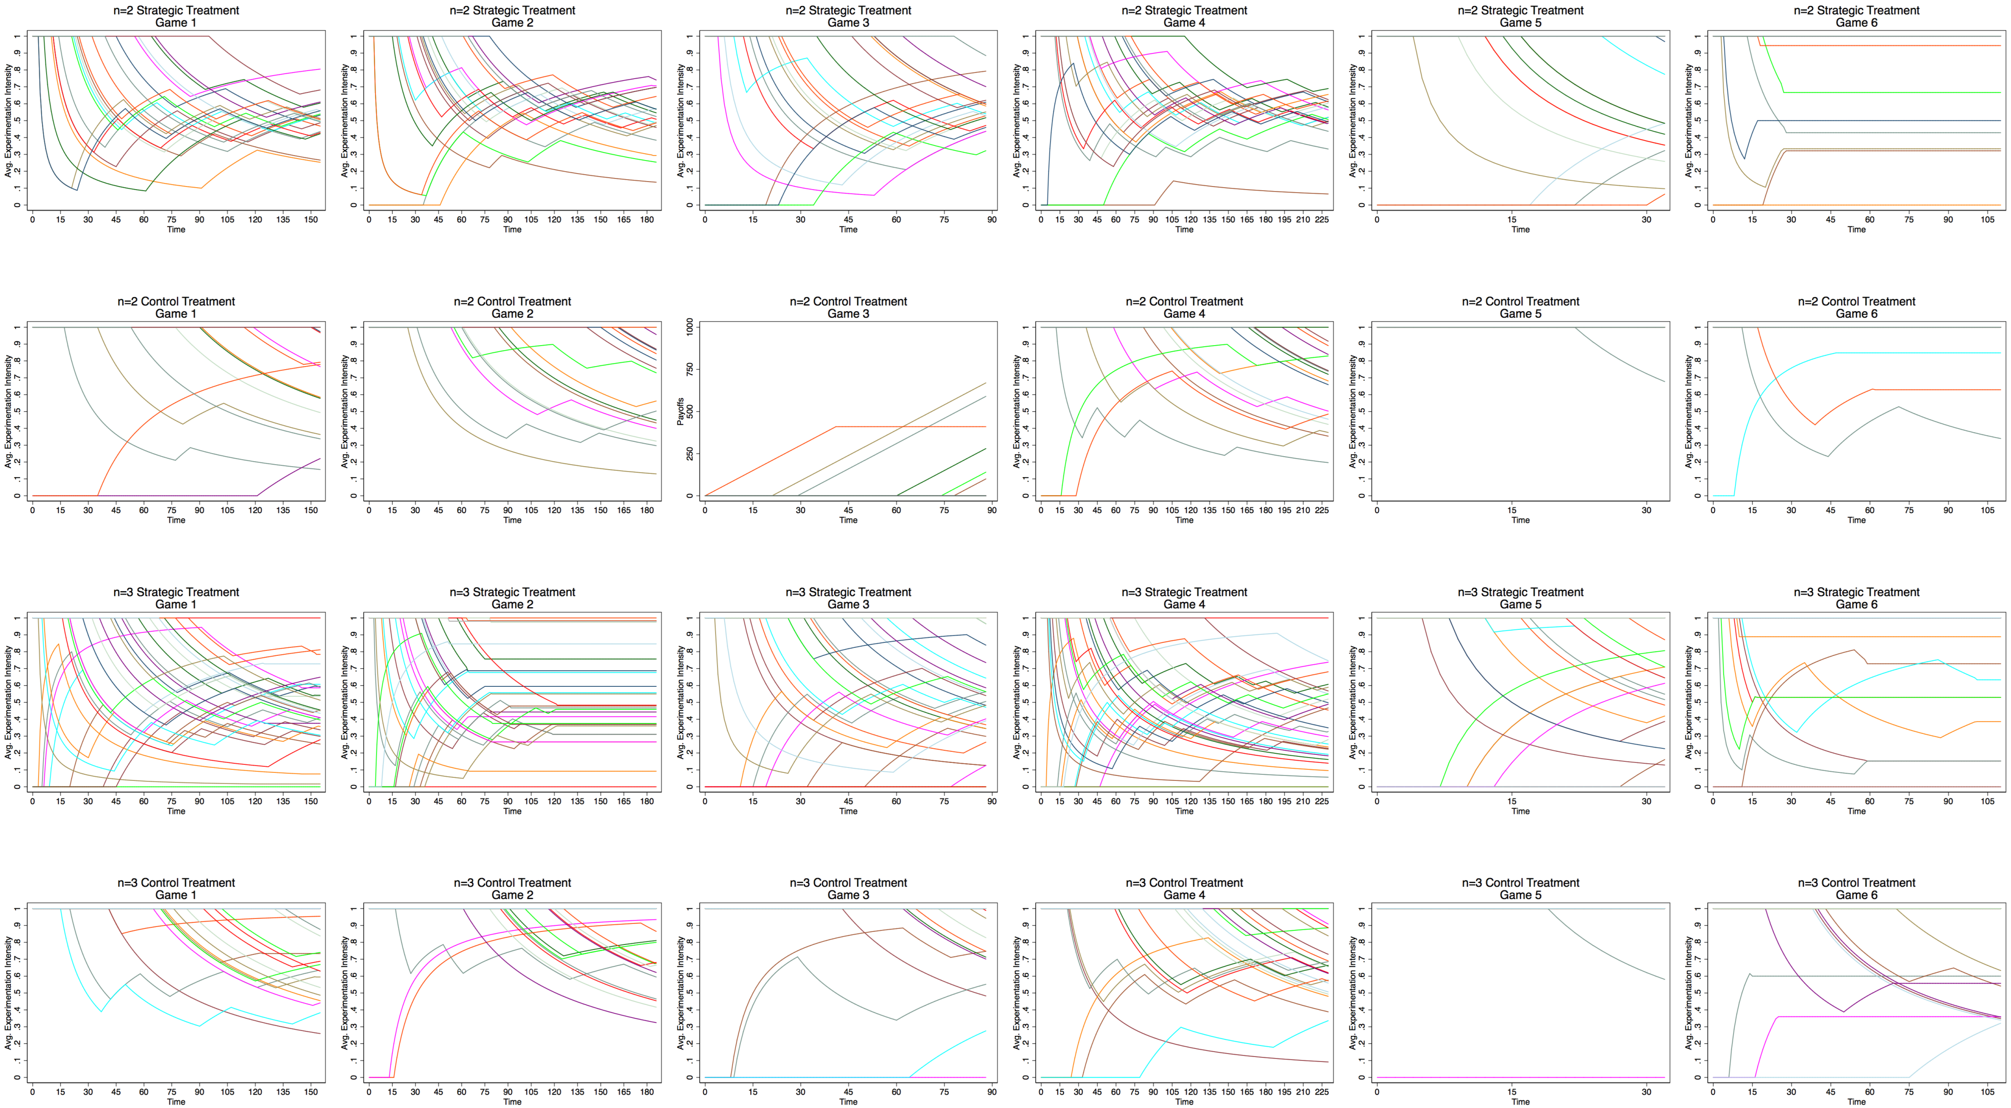

Figure 5: Evolution of Payoffs over Time by Player, by Game

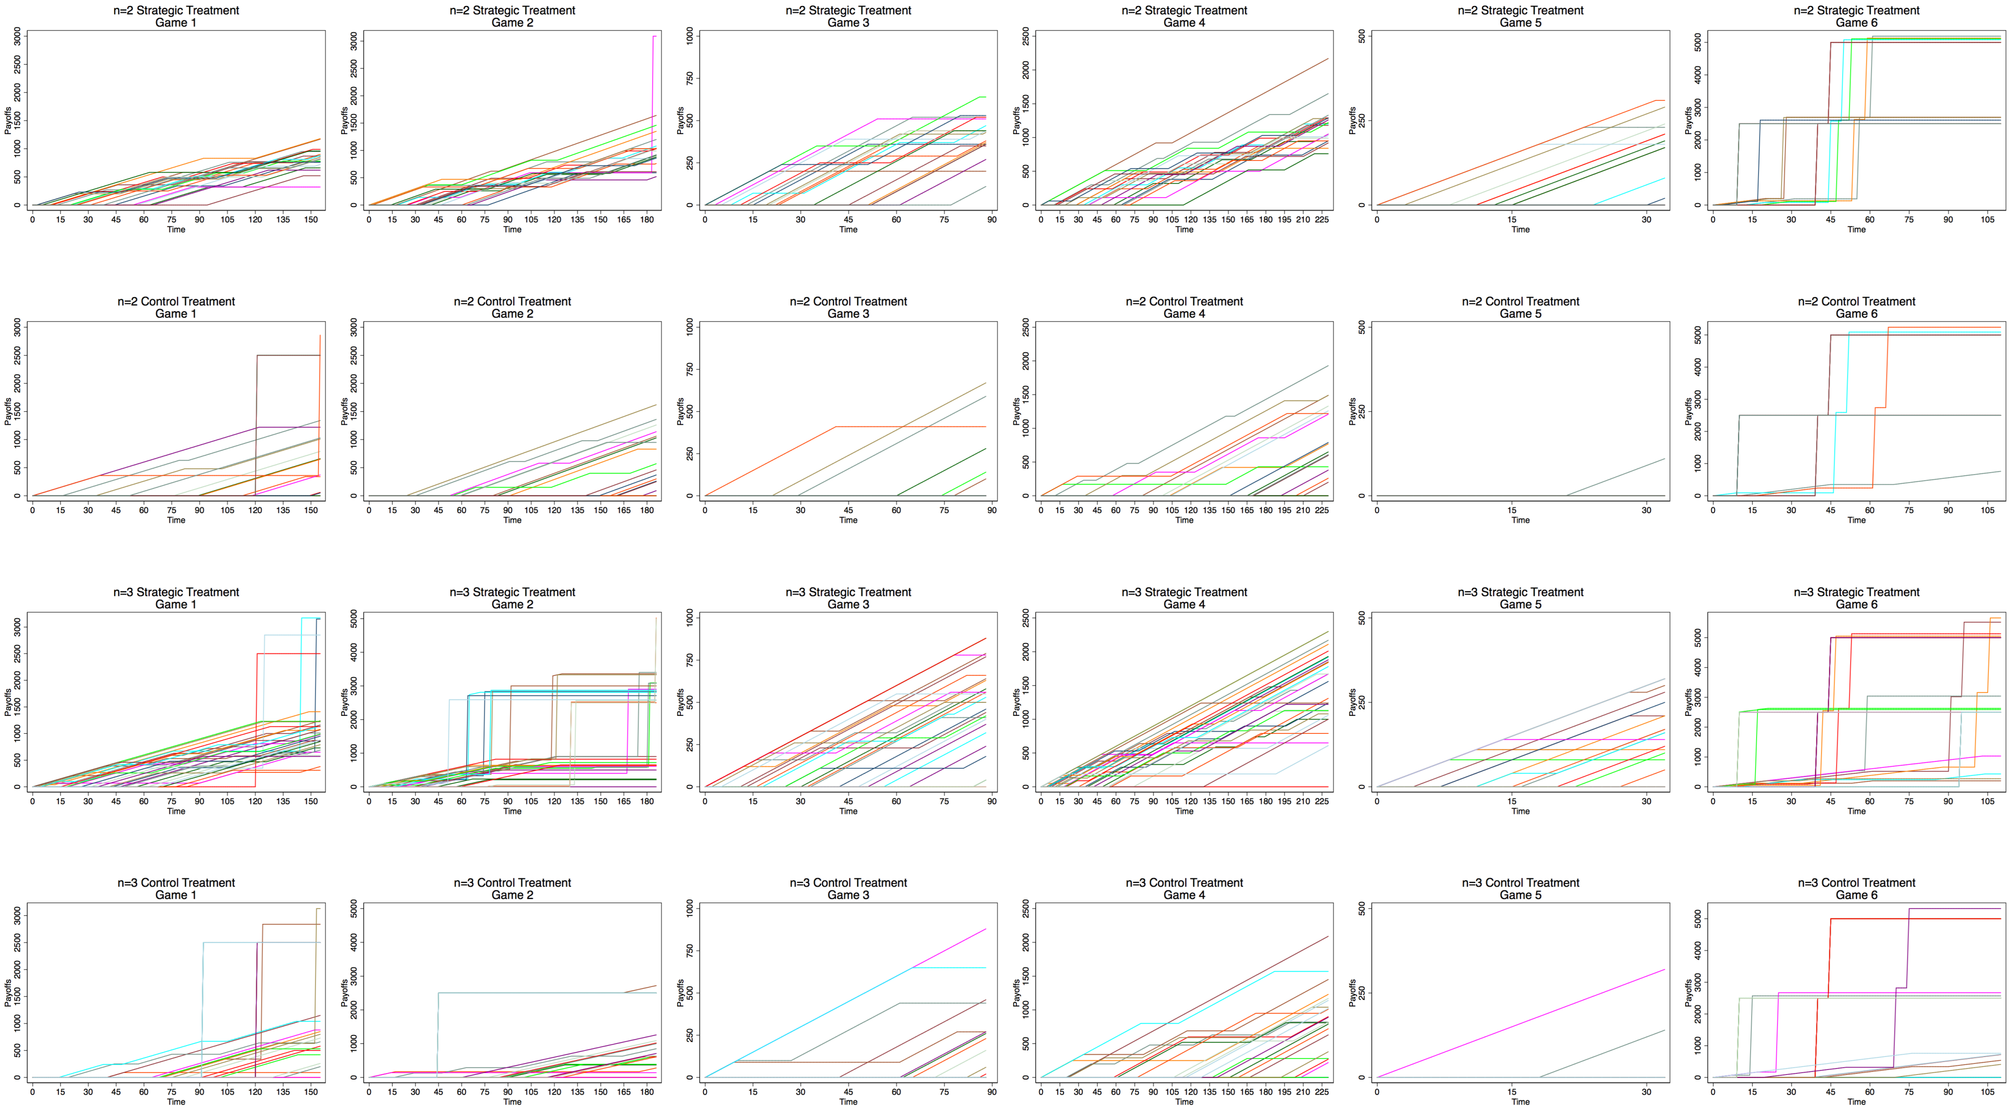

C Appendix: Instructions

The order of the instructions is as follows:

1. $n = 2$: Strategic Treatment
2. $n = 2$: Control Treatment
3. $n = 3$: Strategic Treatment
4. $n = 3$: Control Treatment

After reading the instructions, subjects answered a short comprehension check containing 5 questions.

1. Is your and your partner's risky option always of the same quality?
2. Can you learn from observing each other?
3. If the risky option is good, can you have more than one success?
4. Can these successes happen anytime?
5. The game lasts in expectation 120 seconds but can it end anytime?

After the subjects had successfully completed the test, all participants started the experiment at the same time.

Experiment Instructions

Ground Rules

Welcome to the experiment. Please read the instructions carefully. The earnings you make in this experiment will be paid to you, in cash, at the end of the session.

Your earnings will be determined by your choices and the choices of other participants.

Communication between participants is not allowed. Please use only the computer to input your decisions. Please do not start or end any programs, and do not change any settings.

How Groups are Organized

This experiment consists of six games in total. In the beginning of the first game, participants are randomly matched to pairs and the pairs stay the same in all six games. Therefore, in each game you will interact with the same participant.

How the Timing Works

Games will last on average ***120 seconds*** but may end at any time. The probability that the game ends is the same at each instant. Equivalently, the probability that the game ends during a given period of time depends only on the length of that period of time, and not on how long the game has already been going on. (Such processes are known as *exponential processes* in statistics.)

How the Game Works

In every game, you have to decide whether you want to play the "**safe**" or the "**risky**" option. You can switch between the two options at any time and as often as you like by clicking on the safe (Blue) or risky (Red) button on the screen.

Whenever you choose the **safe** option, your payoff will increase for sure at the rate ***E$ 10***. That means the **safe** option will give you a reward of ***E$ 10*** every second during which you use it.

When you choose the **risky** option, however, what you will be getting depends on the quality of that risky option. The quality of the **risky** option is determined by the computer once and for all at the start of each game; it never changes during the course of the game. We have programmed the computer so that the risky option will be **good** or **bad** with equal probability in each of the six games. The quality of the risky option in later games is independent of its quality in previous games. That is, in each of your six games, with probability ½ your risky option will be **good**; with probability ½ it will be **bad**. The same is true for your partner. Note that your risky option and that of your partner's might or might not be of the same quality.

If your risky option is **good**, it may give you a reward of ***E$ 2500***, but it will only ever do so if you use it. A good risky option yields such a reward after using it on average for 100 seconds. The probability that you get this reward from a good risky option during a given period of time during which you use it depends only on the length of that period of time; it does not depend on anything else, e.g. on how long the game has already been going on. Note that a good risky option may give you more than one reward of *E$ 2500* per game.

If your risky option is **bad**, it will never give you any reward.

You can switch back and forth between the risky option and the safe option at will and as many times as you like. All that matters for your chance of getting the reward is (1) the quality of the risky arm as determined by the computer before the game starts and (2) the overall amount of time you choose to spend on it.

The following graphic illustrates what you are going to see on your screen during the game. The graphs will be updated every second.

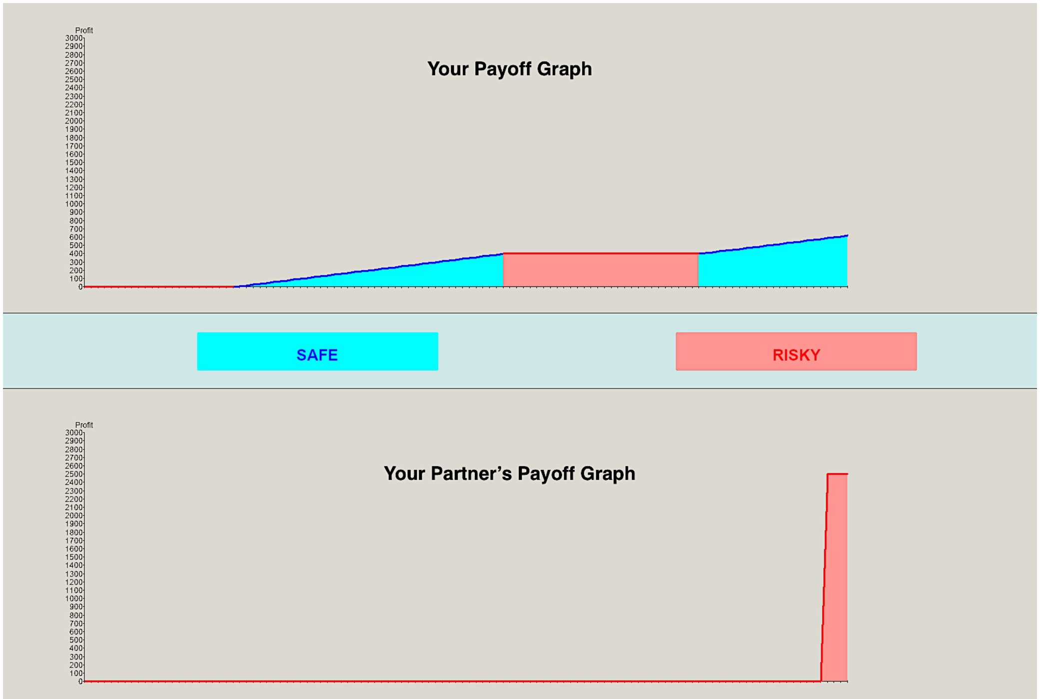

- The **upper** diagram always shows **your** actions and payoffs.
- In this example, you have started playing the risky option (highlighted in Red), then you have switched to the safe option (highlighted in Blue), then you have switched back again to the risky option, etc.
- The **lower** diagram always shows **your partner's** actions and payoffs.
- In this example, your partner has started playing the risky option and continues to do so.
- Note that, in this example, your partner's risky option was good and gave him once a reward of ***E$ 2500***.

The parameters are chosen in such a way that, *if you knew* the risky option to be good, you would be best off by **always** choosing it. Yet, *if you knew* the risky option to be bad, you would be best off by **always** choosing the safe option. In short:

Good risky option > Safe option > **Bad** risky option

Your partner is solving the exact same problem as you and has read the exact same instructions.

Payment

In the experiment you will be making decisions that will earn you E$ (Experimental Dollars). At the end of the experiment, the E$ you earned will be converted into Australian Dollars at an exchange rate of E$ 100 = AU$ 1, and paid out in cash. This amount will be added to your show-up fee of AU$ 5.

After completing the experiment, the computer will randomly select one out of the six games (this will be the same game for all participants), and this game will then be used to determine your payoffs.

Experiment Instructions

Ground Rules

Welcome to the experiment. Please read the instructions carefully. The earnings you make in this experiment will be paid to you, in cash, at the end of the session.

Your earnings will be determined by your choices and the choices of other participants.

Communication between participants is not allowed. Please use only the computer to input your decisions. Please do not start or end any programs, and do not change any settings.

How Groups are Organized

This experiment consists of six games in total. In the beginning of the first game, participants are randomly matched to pairs and the pairs stay the same in all six games. Therefore, in each game you will interact with the same participant.

How the Timing Works

Games will last on average ***120 seconds*** but may end at any time. The probability that the game ends is the same at each instant. Equivalently, the probability that the game ends during a given period of time depends only on the length of that period of time, and not on how long the game has already been going on. (Such processes are known as *exponential processes* in statistics.)

How the Game Works

In every game, you have to decide whether you want to play the "**safe**" or the "**risky**" option. You can switch between the two options at any time and as often as you like by clicking on the safe (Blue) or risky (Red) button on the screen.

Whenever you choose the **safe** option, your payoff will increase for sure at the rate ***E$ 10***. That means the **safe** option will give you a reward of ***E$ 10*** every second during which you use it.

When you choose the **risky** option, however, what you will be getting depends on the quality of that risky option. The quality of the **risky** option is determined by the computer once and for all at the start of each game; it never changes during the course of the game. We have programmed the computer so that the risky option will be **good** or **bad** with equal probability in each of the six games. The quality of the risky option in later games is independent of its quality in previous games. That is, in each of your six games, with probability ½ your (and your partner's!) risky option will be **good**; with probability ½ they will be **bad**. Note that your risky option and that of your partner's will always be of the **same** quality.

If your risky option is **good**, it may give you a reward of ***E$ 2500***, but it will only ever do so if you use it. A good risky option yields such a reward after using it on average for 100 seconds. The probability that you get this reward from a good risky option during a given period of time during which you use it depends only on the length of that period of time; it does not depend on anything else, e.g. on how long the game has already been going on. Note that a good risky option may give you more than one reward of *E$ 2500* per game.

If your risky option is **bad**, it will never give you any reward.

You can switch back and forth between the risky option and the safe option at will and as many times as you like. All that matters for your chance of getting the reward is (1) the quality of the risky arm as determined by the computer before the game starts and (2) the overall amount of time you choose to spend on it.

The following graphic illustrates what you are going to see on your screen during the game. The graphs will be updated every second.

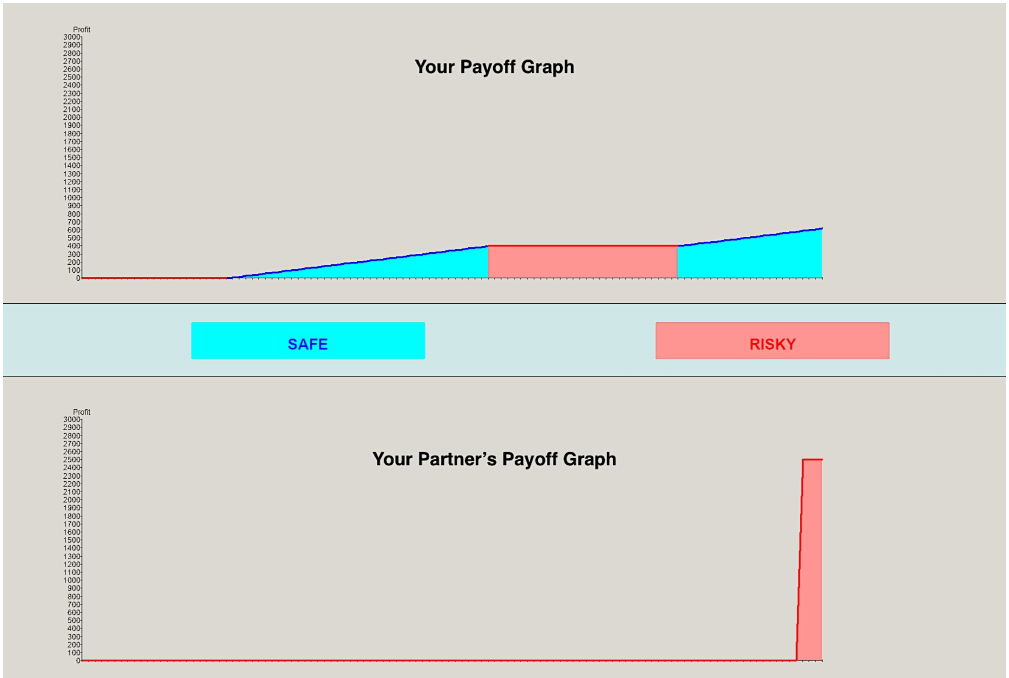

- The **upper** diagram always shows **your** actions and payoffs.
- In this example, you have started playing the risky option (highlighted in Red), then you have switched to the safe option (highlighted in Blue), then you have switched back again to the risky option, etc.
- The **lower** diagram always shows **your partner's** actions and payoffs.
- In this example, your partner has started playing the risky option and continues to do so.
- Note that, in this example, your partner's risky option was good and gave him once a reward of ***E$ 2500***. This means that your risky option was good too.

The parameters are chosen in such a way that, *if you knew* the risky option to be good, you would be best off by **always** choosing it. Yet, *if you knew* the risky option to be bad, you would be best off by **always** choosing the safe option. In short:

Good risky option > Safe option > **Bad** risky option

Your partner is solving the exact same problem as you and has read the exact same instructions. Note that by observing the behaviour of his risky option (provided he uses it) you can learn something about your risky option as well.

Payment

In the experiment you will be making decisions that will earn you E$ (Experimental Dollars). At the end of the experiment, the E$ you earned will be converted into Australian Dollars at an exchange rate of E$ 100 = AU$ 1, and paid out in cash. This amount will be added to your show-up fee of AU$ 5.

After completing the experiment, the computer will randomly select one out of the six games (this will be the same game for all participants), and this game will then be used to determine your payoffs.

Experiment Instructions

Ground Rules

Welcome to the experiment. Please read the instructions carefully. The earnings you make in this experiment will be paid to you, in cash, at the end of the session.

Your earnings will be determined by your choices and the choices of other participants.

Communication between participants is not allowed. Please use only the computer to input your decisions. Please do not start or end any programs, and do not change any settings.

How Groups are Organized

This experiment consists of six games in total. In the beginning of the first game, participants are randomly matched to groups of three players and the groups stay the same in all six games. Therefore, in each game you will interact with the same participants.

How the Timing Works

Games will last on average ***120 seconds*** but may end at any time. The probability that the game ends is the same at each instant. Equivalently, the probability that the game ends during a given period of time depends only on the length of that period of time, and not on how long the game has already been going on. (Such processes are known as *exponential processes* in statistics.)

How the Game Works

In every game, you have to decide whether you want to play the "**safe**" or the "**risky**" option. You can switch between the two options at any time and as often as you like by clicking on the safe (Blue) or risky (Red) button on the screen.

Whenever you choose the **safe** option, your payoff will increase for sure at the rate ***E$ 10***. That means the **safe** option will give you a reward of ***E$ 10*** every second during which you use it.

When you choose the **risky** option, however, what you will be getting depends on the quality of that risky option. The quality of the **risky** option is determined by the computer once and for all at the start of each game; it never changes during the course of the game. We have programmed the computer so that the risky option will be **good** or **bad** with equal probability in each of the six games. The quality of the risky option in later games is independent of its quality in previous games. That is, in each of your six games, with probability ½ your risky option will be **good**; with probability ½ it will be **bad**. The same is true for your partners. Note that your risky option and that of your partners' might or might not be of the same quality.

If your risky option is **good**, it may give you a reward of ***E$ 2500***, but it will only ever do so if you use it. A good risky option yields such a reward after using it on average for 100 seconds. The probability that you get this reward from a good risky option during a given period of time during which you use it depends only on the length of that period of time. It does not depend on anything else, e.g. on how long the game has already been going on. Note that a good risky option may give you more than one reward of *E$ 2500* per game.

If your risky option is **bad**, it will never give you any reward.

You can switch back and forth between the risky option and the safe option at will and as many times as you like. All that matters for your chance of getting the reward is (1) the quality of the risky arm as determined by the computer before the game starts and (2) the overall amount of time you choose to spend on it.

The following graphic illustrates what you are going to see on your screen during the game. The graphs will be updated every second.

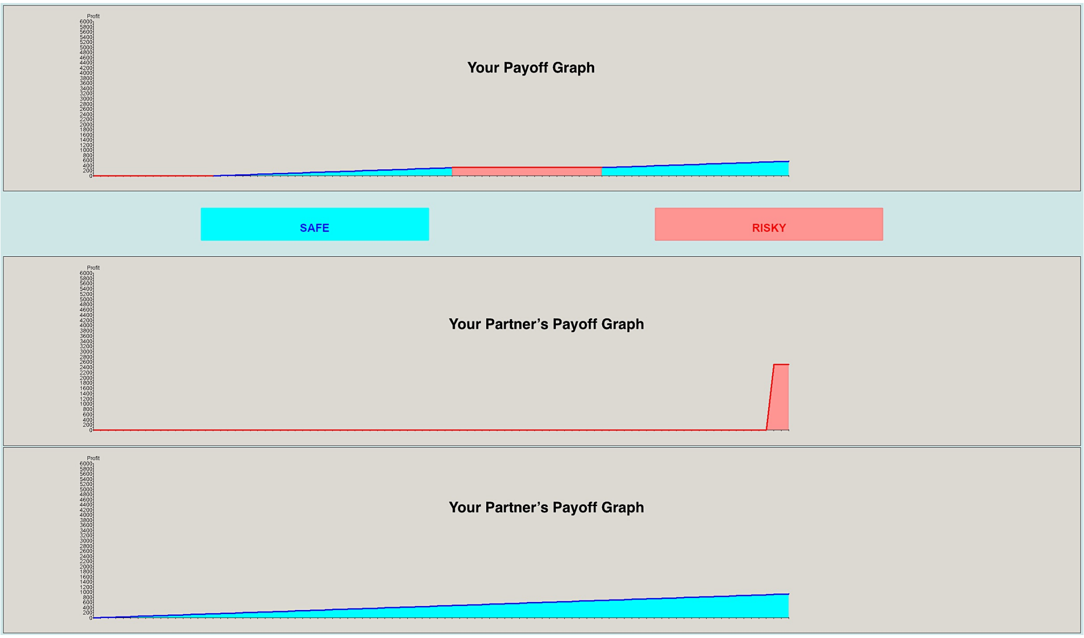

- The **upper** diagram always shows **your** actions and payoffs.
- In this example, you have started playing the risky option (highlighted in Red), then you have switched to the safe option (highlighted in Blue), then you have switched back again to the risky option, etc.
- The **lower** diagram always shows **your partner's** actions and payoffs.
- In this example, one of your partners has started playing the risky option and continues to do so. The other partner has started and continues playing the safe option.
- Note that, in this example, at least one of your partner's risky option was good and gave him once a reward of ***E$ 2500***.

The parameters are chosen in such a way that, *if you knew* the risky option to be good, you would be best off by **always** choosing it. Yet, *if you knew* the risky option to be bad, you would be best off by **always** choosing the safe option. In short:

Good risky option > Safe option > **Bad** risky option

Your partners are solving the exact same problem as you and have read the exact same instructions.

Payment

In the experiment you will be making decisions that will earn you E$ (Experimental Dollars). At the end of the experiment, the E$ you earned will be converted into Australian Dollars at an exchange rate of E$ 100 = AU$ 1, and paid out in cash. This amount will be added to your show-up fee of AU$ 5.

After completing the experiment, the computer will randomly select one out of the six games (this will be the same game for all participants), and this game will then be used to determine your payoffs.

Experiment Instructions

Ground Rules

Welcome to the experiment. Please read the instructions carefully. The earnings you make in this experiment will be paid to you, in cash, at the end of the session.

Your earnings will be determined by your choices and the choices of other participants.

Communication between participants is not allowed. Please use only the computer to input your decisions. Please do not start or end any programs, and do not change any settings.

How Groups are Organized

This experiment consists of six games in total. In the beginning of the first game, participants are randomly matched to groups of three players and the groups stay the same in all six games. Therefore, in each game you will interact with the same participants.

How the Timing Works

Games will last on average ***120 seconds*** but may end at any time. The probability that the game ends is the same at each instant. Equivalently, the probability that the game ends during a given period of time depends only on the length of that period of time, and not on how long the game has already been going on. (Such processes are known as *exponential processes* in statistics.)

How the Game Works

In every game, you have to decide whether you want to play the "**safe**" or the "**risky**" option. You can switch between the two options at any time and as often as you like by clicking on the safe (Blue) or risky (Red) button on the screen.

Whenever you choose the **safe** option, your payoff will increase for sure at the rate ***E$ 10***. That means the **safe** option will give you a reward of ***E$ 10*** every second during which you use it.

When you choose the **risky** option, however, what you will be getting depends on the quality of that risky option. The quality of the **risky** option is determined by the computer once and for all at the start of each game; it never changes during the course of the game. We have programmed the computer so that the risky option will be **good** or **bad** with equal probability in each of the six games. The quality of the risky option in later games is independent of its quality in previous games. That is, in each of your six games, with probability ½ your (and your partners'!) risky option will be **good**; with probability ½ they will be **bad**. Note that your and your partners' risky option will always be of the **same** quality.

If your risky option is **good**, it may give you a reward of ***E$ 2500***, but it will only ever do so if you use it. A good risky option yields such a reward after using it on average for 100 seconds. The probability that you get this reward from a good risky option during a given period of time during which you use it depends only on the length of that period of time. It does not depend on anything else, e.g. on how long the game has already been going on. Note that a good risky option may give you more than one reward of *E$ 2500* per game.

If your risky option is **bad**, it will never give you any reward.

You can switch back and forth between the risky option and the safe option at will and as many times as you like. All that matters for your chance of getting the reward is (1) the quality of the risky arm as determined by the computer before the game starts and (2) the overall amount of time you choose to spend on it.

The following graphic illustrates what you are going to see on your screen during the game. The graphs will be updated every second.

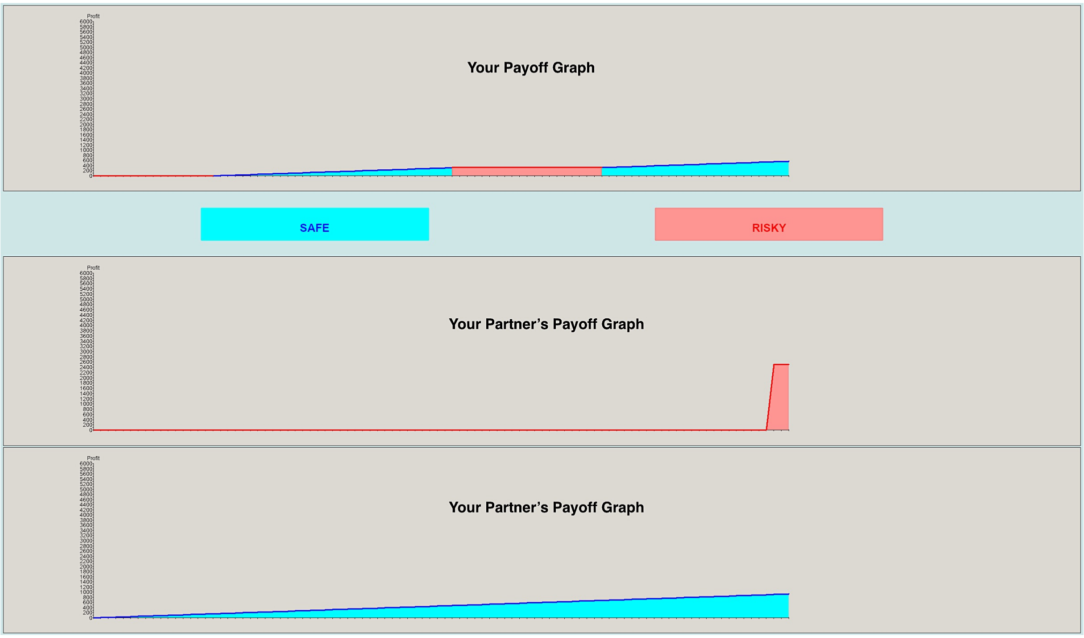

- The **upper** diagram always shows **your** actions and payoffs.
- In this example, you have started playing the risky option (highlighted in Red), then you have switched to the safe option (highlighted in Blue), then you have switched back again to the risky option, etc.
- The **lower** diagram always shows **your partners'** actions and payoffs.
- In this example, one of your partners has started playing the risky option and continues to do so. The other partner has started and continues playing the safe option.
- Note that, in this example, your partner's risky option was good and gave him once a reward of ***E$ 2500***. This means that your risky option was good too.

The parameters are chosen in such a way that, *if you knew* the risky option to be good, you would be best off by **always** choosing it. Yet, *if you knew* the risky option to be bad, you would be best off by **always** choosing the safe option. In short:

Good risky option > Safe option > **Bad** risky option

Your partners are solving the exact same problem as you and have read the exact same instructions. Note that by observing the behaviour of their risky option (provided they use it) you can learn something about your risky option as well.

Payment

In the experiment you will be making decisions that will earn you E$ (Experimental Dollars). At the end of the experiment, the E$ you earned will be converted into Australian Dollars at an exchange rate of E$ 100 = AU$ 1, and paid out in cash. This amount will be added to your show-up fee of AU$ 5.

After completing the experiment, the computer will randomly select one out of the six games (this will be the same game for all participants), and this game will then be used to determine your payoffs.