Supplement to "Bandits in the Lab"

Johannes Hoelzemann^{*} Nicolas Klein[†]

February 28, 2021

A Analysis of Individual Games

In Section 5 and 6 of Hoelzemann and Klein (2021), 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.

Figures 1, 2 and 3 display the evolution of players' action choices over all six games. Players' actions are described by dots, the width of which corresponds to one second of time. For each of the six games, we conducted four treatments with ten groups each, the parameters of which (i.e., their duration, the quality of the risky arm and the timing of successes on the risky arm in case it was good) we had simulated ahead of time, as explained in Section 4 of the main text. As the figures show, the duration of the games ranged from 32 seconds for Game 5 to 230 seconds for Game 4. As is furthermore evident from the figures, players change their behaviors over time. While often playing risky at the beginning, players seem to grow less inclined to use the risky arm the longer it has unsuccessfully been used before. This shows that our subjects adapted to the evolving information about their environment.

^{*}University of Toronto. Mailing address: University of Toronto, Department of Economics, Max Gluskin House, 150 St. George Street; Toronto, ON M5S 2E9, Canada. e-mail: j.hoelzemann@utoronto.ca.

[†]Université de Montréal and CIREQ. Mailing address: Université de Montréal, Département de Sciences Économiques, C.P. 6128 succursale Centre-ville; Montréal, H3C 3J7, Canada. e-mail: kleinnic@yahoo.com.



Games 1 and 6 are shown. Players' actions are described by dots, the width of which corresponds to one second of time. Groups 1-10 correspond to the strategic treatment for two-player groups; groups 11-20 are the corresponding control treatments. Groups 21-30 played the strategic treatment for three-player groups, while groups 31-40 were the corresponding control treatments. In each group, we refer to the lowermost player as 'player 1', while 'player 2' will denote the player right above, and 'player 3' is the uppermost player. The x-axis represents calendar time. A *red* dot indicates that a player is playing *risky* in a given second, while a *blue* dot indicates that the player is playing *safe*. A black square indicates the arrival of a lump sum on the risky arm.

Figure 1: Action Choices by Players over Time, Games 1 & 6



Games 2 and 3 are shown. Players' actions are described by dots, the width of which corresponds to one second of time. Groups 1-10 correspond to the strategic treatment for two-player groups; groups 11-20 are the corresponding control treatments. Groups 21-30 played the strategic treatment for three-player groups, while groups 31-40 were the corresponding control treatments. In each group, we refer to the lowermost player as 'player 1', while 'player 2' will denote the player right above, and 'player 3' is the uppermost player. The x-axis represents calendar time. A *red* dot indicates that a player is playing *risky* in a given second, while a *blue* dot indicates that the player is playing *safe*. A black square indicates a the arrival of a lump sum on the risky arm.

Figure 2: Action Choices by Players over Time, Games 2 & 3



Games 4 and 5 are shown. Players' actions are described by dots, the width of which corresponds to one second of time. Groups 1-10 correspond to the strategic treatment for two-player groups; groups 11-20 are the corresponding control treatments. Groups 21-30 played the strategic treatment for three-player groups, while groups 31-40 were the corresponding control treatments. In each group, we refer to the lowermost player as 'player 1', while 'player 2' will denote the player right above, and 'player 3' is the uppermost player. The x-axis represents calendar time. A *red* dot indicates that a player is playing *risky* in a given second, while a *blue* dot indicates that the player is playing *safe*. There were no lump-sum arrivals in Games 4 and 5.

Figure 3: Action Choices by Players over Time, Games 4 & 5

A.1 Experimentation Intensity

In order to illustrate subjects' dynamically evolving incentives for public-good provision, Figure 4 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. 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 the end. Behavior in the control treatment, however, provides a sharp contrast, as most curves are monotonically decreasing, indicating cut-off behavior.



Experimentation Intensity for each subject.

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

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 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.

¹Corresponding figures for the other games look qualitatively similar.

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 *free-riding region*, as this game lasts only 32 seconds, implying that the *free-riding 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 *free-riding region* corresponds to three individual players in one group in the control treatment that have not reached the *free-riding region* either on account of an early success or because they did not use the risky arm enough. Table I summarizes our findings for each game separately by belief region.

		<i>n</i> = 2				<i>n</i> = 3				
		Strategic Treatment Control Treatment		ol Treatment	Strategic Treatment		Control Treatment			
Game	Belief	Obs.	Exp.	Obs.	Exp.	Obs.	Exp.	Obs.	Exp.	
	Region		Intensity		Intensity		Intensity		Intensity	
1	All	10	.508 [.065]	10	.730 [.238]	10	.455 [.107]	10	.797 [.217]	
2	—	10	.512 [.116]	10	.696 [.234]	10	.543 [.227]	10	.833 [.125]	
3	—	10	.565 [.086]	10	.878 [.176]	10	.457 [.169]	10	.866 [.199]	
4	—	10	.519 [.120]	10	.678 [.194]	10	.383 [.089]	10	.728 [.183]	
5	—	10	.653 [.204]	10	.984 [.051]	10	.596 [.264]	10	.953 [.110]	
6	_	10	.810 [.259]	10	.941 [.113]	10	.800 [.314]	10	.857 [.182]	
1	Risky 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	Free-Riding	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]	

Table I: Average Experimentation Intensity by Belief Regions and by Game

Average [st. dev.] experimentation intensity using group averages. For n = 3 in the control treatment, only players in nine groups entered the *free-riding region*.

Also, at the game level, the average experimentation intensity is substantially lower in the strategic treatment for both belief regions. In the *free-riding region* for groups of size n = 2, the corresponding *p*-values are 0.1679, 0.0176, 0.0089, and 0.0186 for Games 1-4, respectively. For groups of size n = 3, the same is true, 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. As for the *risky dominant region*, experimentation intensities in the two-player groups are lower in the strategic treatment, which is statistically significant at least at the 5%-level in Games 1, 3, and 5; there is no significant difference for Games 2 and 4. The *p*-values of the two-sided

Wilcoxon ranksum test amount to 0.0249, 0.1364, 0.0044, 0.1180, and 0.0013 for Games 1 to 5, respectively. The difference is statistically significant at least at the 10%-level for groups of size n = 3, however, with *p*-values amounting to 0.0823, 0.0138, 0.0097, 0.0215, and 0.0013 for Games 1-5, respectively.

We further conduct our "difference-in-differences"-analysis by comparing the difference in intensities across belief regions and across treatments for Games 1-4. For n = 2, this difference-in-differences is higher in the strategic treatment in all four games, though not statistically significant. The corresponding *p*-values are 0.4490, 0.1255, 0.5366, and 0.3239 for Games 1 to 4, respectively. For groups of size n = 3, there is no statistical evidence of a difference in treatments with *p*-values of 0.9393, 0.1304, 0.7090, and 0.6477 for Games 1 to 4, respectively.

While we cannot establish statistical significance on the individual-game level in the difference-in-differences analysis, we can do so by directly testing for different experimentation intensities between the two belief regions. In particular, in the strategic treatment for two-player groups, the difference between the *risky dominant* and the *free-riding 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. 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 strategic treatment with n = 3, even though average experimentation intensities decrease when moving from the *risky dominant region* to the *free-riding region*, no such 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 indication 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 outlier here, with behavior across regions exhibiting significant differences (the *p*-value is 0.0138).

A.2 Payoffs

Table II 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. For n = 2 (n = 3), the *p*-values are 0.0233 (0.0004), 0.0007 (0.0015), 0.0081 (0.0007), and 0.0012 (0.0013) for Games 2-5, 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. Game 1 is an outlier in that average final payoffs are statistically significantly higher *in the control treatment*, with *p*-values of 0.0342 (0.0820). 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).

		Strategic Tre	atment		Control Treatment				
Game	Obs.	Final Payoff	Min	Max	Obs.	Final Payoff	Min	Max	
Panel 1	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	

Table II: Average Final Payoffs by Game

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



Payoffs over time for each subject.

Figure 5: Evolution of Payoffs over Time by Player

Figure 5 displays the evolution of each player's cumulated payoff over time for Game 2. Positive slopes correspond to periods during which a subject played safe; flat parts indicate hapless risky play, while jumps denote lump sums arrivals from the risky arm.

In Table III, we provide the theoretically expected payoffs *conditional* on the realizations of the stochastic processes, which we had simulated ahead of time.² Of course, *conditionally* on a particular realization of the stochastic process, *ex ante* optimal behaviors may do very poorly, while *ex ante* very eccentric behaviors may well be optimal.³ Note, for instance, that, in Game 4, the equilibrium strategy, which gives up earlier, does better than the efficient solution. In fact, for groups of size n = 2, the best PBE does weakly better than the efficient solution for all six games. As differences in predicted payoffs are to a large extent driven by the timing of the big lump-sum payoffs from the risky arm, Table III provides a cautionary tale against ascribing excessive inferential value to observed payoff differences; except for Result 2, which compares payoffs across treatments for *given realizations* of the stochastic processes, we have relied on differences in observed behavior for our inference. Indeed, as subjects did not know the realizations of the stochastic processes when choosing their actions, observed behavioral differences will "filter out" the considerable additional noise that stems from the—"very stochastic"—mapping of behavior into realized payoffs.

Table III: Average Predicted Payoffs

	Efficient		Best	Best PBE		MPE		Single-Agent	
Game	<i>n</i> = 2	<i>n</i> = 3							
1	470.00	1666.67	950.00	1150.00	955.00	1150.00	360.00	1073.33	
2	780.00	2500.00	1260.00	1460.00	1265.00	1460.00	670.00	1280.00	
3	0	0	280.00	480.00	280.00	480.00	0	0	
4	1230.00	1340.00	1710.00	1900.00	1705.00	1900.00	1120.00	1100.00	
5	0	0	0	0	5.00	70.00	0	0	
6	3750.00	3333.33	3750.00	3333.33	3750.00	3333.33	3750.00	2500.00	

A.3 Cut-Off Behavior

We now turn to the frequency of cut-off behavior. As we have seen in Result 3, 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

²To get our MPE estimates, we assume each player hypothetically splitting each instant 50:50 between the two arms in the *free-riding region*, which, for the purpose of this table, we equate to the belief region $(p_1^*, \frac{p^*+\bar{p}}{2})$.

³Indeed, the equilibrium strategy in the matching-pennies game, for instance, while being an *ex ante* best response, will do rather poorly conditionally on a *particular realization* of the opponent's equilibrium strategy; by contrast, a pure strategy does strictly better than the equilibrium strategy given a particular realization of the opponent's equilibrium mixed strategy. Our game is no different in this respect. For example, in Game 4, the best course of action conditionally on the realizations of the random variables would have been to play safe throughout, even though *ex ante* "safe" is a dominated action at the start of the game. (Indeed, $p_0 > p^m$, so that even a myopic player should play risky.) By the same token, in Game 1, players should have switched to safe right after player 1 first obtained a success, had they known that the game ended before the second success would arrive.

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.

		<i>n</i> = 2		<i>n</i> = 3			
		Strategic Treatment	Control Treatment		Strategic Treatment	Control Treatment	
Game	Obs.	Tot. (Rel.)	Tot. (Rel.)	Obs.	Tot. (Rel.)	Tot. (Rel.)	
		Freq.	Freq.		Freq.	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)	

Table IV: Frequency of Cut-Off Behavior by Game

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 the frequency of cut-off behavior 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 increase 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.

When we use our "continuous" measure of cut-off behavior, differences across treatments are mostly highly statistically significant as well. Recall that this measure is defined as 1 minus the proportion of time in which a subject plays safe before ever playing risky, or plays risky after they had previously switched from risky to safe, before his risky arm is revealed to be good or the end of the game, whichever arrives first. For groups of size n = 2, the *p*-values are 0.0014, 0.0002, 0.0004, 0.0001, 0.0682, and 0.4247 for Games 1-6, respectively. For n = 3, the corresponding *p*-values are 0.0014, 0.0099, 0.0070, 0.0717, 0.0501, and 0.9628 for Games 1-6, respectively.

A.4 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 4.

		<i>n</i> = 2			<i>n</i> = 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]

Table V: Proportion of Time with a Single Pioneer by Game

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

As Table V highlights, also at the individual game level, we can confirm for all games that the addition of one (two) perfectly positively correlated arm(s) 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 for Games 1-5, but not for Game 6, in the 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 incidence of switches 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 success for two players: after 9 seconds of exploration by Player 1 and after 39 and 44 seconds of exploration by Player 1.

A.5 Switches of Action

In any Markov Perfect Equilibrium, we should expect players to switch roles at least once. As theory predicts and Result 5 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 VI displays the average number of switches per player across games for our four treatments. As in the main text of Hoelzemann and Klein (2021), we perform our statistical tests on the average *incidence* (rather than the *number*) of switches, and find that the average 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.

		<i>n</i> = 2			<i>n</i> = 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]

Table VI: Average Number of Switches per Player by Game

Average [st. dev.] switches per player using group averages.

The number of observations refers to both strategic and control treatment.

A.6 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 paying attention to the additional information their partner(s) provided them, a necessary condition for free-riding. 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.

		<i>n</i> = 2			<i>n</i> = 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]

Table VII: Average Fixation Intensity by Game

Average [st. dev.] fixation intensity using group averages.

The number of observations refers to both strategic and control treatment.

As Table VII 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 6 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 6 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.



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).

Figure 6: Heatmaps of Four Treatments

References

HOELZEMANN, J., AND N. KLEIN (2021): "Bandits in the Lab," Quantitative Economics, forthcoming.