Breakdowns in the Lab
Supplemental Appendix: Additional Results

Johannes C. Hoelzemann Nicolas Klein

August 24, 2026

A Additional Robustness and Behavioral Diagnostics

This appendix provides additional analyses addressing inference with repeated observations, best responses to the empirical distribution of play, two auxiliary regressions discussed in the main text, and alternative descriptions of behavioral heterogeneity. Unless otherwise noted, standard errors are clustered at the matching-group level. There are 16 matching groups and 104 participants in the bad-news sample.

A.1 Alternative Clustering at the Participant Level

The main analysis clusters standard errors at the matching-group level, which is the level at which participants are repeatedly rematched. As an additional robustness check, Table A.1 reports participant-clustered standard errors for the principal individual-level treatment effects. The point estimates are unchanged. In particular, the unconditional first-period treatment effect remains 0.174 and is statistically significant under either clustering scheme.

Table A.2 repeats the history-dependence specifications from the main text using standard errors clustered at the participant level. The qualitative pattern is unchanged: own prior experimentation is the strongest predictor of current risky play, partner experimentation remains positively associated with current play, and the interaction between observability and partner experimentation remains statistically insignificant.

A.2 Best Responses to the Empirical Distribution of Play

We next ask whether the observed level of experimentation can be rationalized as a best response to the empirical frequency with which opponents experiment. This is a

Table A.1: Treatment Effects under Alternative Clustering Schemes

Outcome	Coefficient	MG-clustered s.e.	Participant-clustered s.e.	MG p	Participant p
First-period risky choice	0.174	0.067	0.054	0.020	0.002
Individual experimentation intensity	0.157	0.052	0.044	0.008	< 0.001
Individual payoff (points)	1,398	175	287	< 0.001	< 0.001

Notes: The first row uses all 2,600 first-period decisions. The second and third rows use 2,600 participant–game observations and control for the learning variable used in Table IV of the main text. "MG" denotes matching group. The coefficient is the effect of observable actions. Matching-group p-values use the finite-cluster t reference distribution; participant-clustered inference uses 104 participant clusters.

Table A.2: History Dependence with Participant-Clustered Standard Errors

	(1)	(2)	(3)	(4)
Observable actions	0.066	0.037	0.076	
	(0.036)	(0.051)	(0.037)	
Own prior risky choices	0.256	0.253	0.265	0.135
	(0.026)	(0.026)	(0.027)	(0.027)
Partner prior risky choices	0.063	0.049	0.069	0.069
	(0.017)	(0.020)	(0.018)	(0.014)
Observable × partner prior risky		0.032		
		(0.030)		
Period 3	-0.245	-0.246	-0.251	-0.163
	(0.024)	(0.024)	(0.025)	(0.023)
Observations	4,106	4,106	3,280	4,106
Individual fixed effects	No	No	No	Yes

Notes: Participant-clustered standard errors are in parentheses. Columns (1), (2), and (4) use period-2 and period-3 decisions following histories in which no breakdown had previously occurred. Column (3) uses the restricted sample employed in the corresponding specification in the main text. The observable-actions indicator is absorbed by the individual fixed effects in Column (4).

descriptive empirical-best-response exercise rather than an equilibrium test. We treat empirical opponent experimentation frequencies as plug-in beliefs and let the focal player optimize given those frequencies. In period 3 the calculation is exact because the decision is myopic. In earlier periods we solve backwards, taking the opponent's history-contingent empirical experimentation frequency as given.

For reference, under the efficient observable-actions equilibrium the expected payoff gain from following the prescribed period-1 risky action rather than deviating to safe and inducing reversion to safe play is 55.1 points.

Table A.3 reports the first-period empirical-best-response calculation. At the empirical opponent experimentation rates, safe yields the higher expected payoff in period 1 in both treatments. The risky-minus-safe payoff difference is approximately -150 points under observable actions and -196 points under unobservable actions. A matching-group bootstrap yields 95% intervals of approximately $[-181, -121]$ points and $[-249, -135]$ points, respectively.

Table A.3: Period-1 Best Response to Empirical Opponent Play

Treatment	Opponent risky rate	$U(R) - U(S)$	Best response
Observable actions	0.726	-150	Safe
Unobservable actions	0.552	-196	Safe

Notes: The calculation uses the empirical period-1 experimentation rate of the opponent and history-specific empirical period-2 experimentation rates. The focal player optimizes in periods 2 and 3. The exercise describes best responses to the empirical distribution of play and does not impose equilibrium consistency on the opponent's strategy.

The later-period calculations illustrate why substantial experimentation can nevertheless be individually optimal after informative histories. Under observable actions in period 2, risky is optimal after any first-period history containing at least one risky choice at the empirical continuation frequencies, but not after a history in which both players chose safe. The risky-minus-safe payoff differences for first-period action histories (own, other) $= (0,0), (0,1), (1,0), (1,1)$ are approximately -463, 12, 22, and 384 points, respectively. Under unobservable actions, the corresponding period-2 difference is approximately -217 points after the participant chose safe in period 1 and 243 points after the participant chose risky.

In period 3, with observable actions, risky is myopically optimal once at least two risky-arm uses have been observed without a breakdown, exactly as implied by the theoretical analysis. The risky-minus-safe payoff differences after $n = 0, 1, 2, 3, 4$ successful observed risky uses are approximately -523, -194, 157, 516, and 869 points.

With unobservable actions, using empirical opponent experimentation frequencies to integrate over the unobserved action history gives approximate period-3 differences of −158, 194, and 553 points after the participant has personally made zero, one, or two previous risky choices without a breakdown.

These calculations reinforce the interpretation in the main text. The treatment comparison is strong, but observed play should not be interpreted as literal implementation of the equilibrium strategies: at empirical opponent frequencies, risky is not a first-period best response in either treatment, while it becomes a best response after sufficiently informative no-breakdown histories.

A.3 Additional Regressions Reported in the Main Text

Table A.4 reports the regression relating period-2 risky play to the participant's own and the opponent's period-1 actions. The sample is restricted to histories with no period-1 breakdown. The positive coefficient on the opponent's period-1 risky choice is robust to either matching-group or participant-level clustering.

Table A.4: Opponent's Period-1 Action and Period-2 Risky Play

	Coefficient	MG-clustered s.e.	Participant-clustered s.e.
Observable actions	0.073	0.043	0.043
Own period-1 risky choice	0.265	0.069	0.043
Opponent period-1 risky choice	0.114	0.036	0.028
Observations	2,194		
R^2	0.120		

Notes: OLS linear probability model. The dependent variable is the participant's period-2 risky choice. The sample contains period-2 decisions following a first period without a breakdown. Matching-group clustered p-values for own and opponent period-1 risky choices are 0.002 and 0.006, respectively; participant-clustered p-values are both below 0.001.

Table A.5 reports the cross-game persistence specification using the actual chronological order of play. For each participant, the first-period risky choice in game order o is regressed on the participant's own first-period choice and the first-period choice of the participant's actual partner in game order $o - 1$, together with the observable-actions indicator and game order. Own behavior is strongly persistent, whereas there is no evidence of imitation of the previous partner.

Table A.5: Cross-Game Persistence and Imitation

	Coefficient	MG-clustered s.e.	Participant-clustered s.e.
Own previous first-period risky choice	0.313	0.030	0.040
Previous partner's first-period risky choice	-0.011	0.027	0.022
Observable actions	0.122	0.049	0.038
Game order	0.0008	0.0011	0.0009
Observations	2,496		
R^2	0.127		

Notes: OLS linear probability model. The dependent variable is the participant's first-period risky choice in game order o. The chronological game-order variable is used to construct the lag. With matching-group clustering, the p-values are $p < 0.001$ for own previous risky choice, $p = 0.694$ for the previous partner's risky choice, $p = 0.025$ for observable actions, and $p = 0.464$ for game order. Participant-clustered inference gives the same substantive conclusion.

A.4 Behavioral-Pattern Decomposition

The main text classifies participants using their overall pre-breakdown experimentation frequency. As a complementary analysis that does not rely on arbitrary participant-level cutoffs, we classify each participant–game according to the sequence of actions taken before the first breakdown. An *all-risky* path consists exclusively of risky choices up to the first breakdown (or the end of the game); an *all-safe* path consists exclusively of safe choices; and the remaining categories distinguish the direction of within-game switching. We then compute each participant's share of games in each category and compare these participant-level shares across treatments.

Table A.6: Behavioral Patterns within Games

Pattern	Observable	Unobservable	Difference	MG-clustered p
All risky until breakdown/end	0.550	0.418	0.132	0.129
All safe until breakdown/end	0.069	0.225	-0.155	0.001
Safe → risky	0.165	0.173	-0.008	0.861
Risky → safe	0.118	0.092	0.025	0.352
Multiple switches	0.098	0.092	0.007	0.741
Post-breakdown safe-choice rate	0.655	0.776	-0.121	0.088

Notes: The first five rows partition participant–games according to the sequence of choices made before the first breakdown. Treatment means are the averages of participant-level shares across the 25 games. The final row reports the participant-level fraction of post-breakdown decision opportunities in which the safe arm is chosen. Inference clusters by matching group.

The main treatment difference in this decomposition is a sharp reduction in all-safe paths: such paths account for 22.5% of games under unobservable actions but only 6.9% under observable actions. The shares of games involving within-game switching are very similar across treatments. Thus, the aggregate treatment effect is not simply an artifact of the particular one-third/two-thirds thresholds used for the participant typology in the main text.

A.5 Why Persistent Experimenters Earn Less than Intermediate Experimenters

Table A.7 decomposes the realized payoffs of intermediate and persistent experimenters in the observable-actions treatment using the classification in the main text. Persistent experimenters generate more successful risky-arm payoffs, but they are also exposed to substantially more breakdowns. Because a breakdown is much more costly than the gain from a successful risky choice, the additional realized breakdown losses more than offset the additional successful risky payoffs in this sample.

Table A.7: Payoff Decomposition by Behavioral Type, Observable Actions

	Intermediate	Persistent experimenters	Difference
Average payoff per game (points)	2,356	1,946	-410
Successful risky choices per game	1.526	2.037	0.511
Breakdowns per game	0.117	0.226	0.109
Breakdowns before any prior breakdown	0.108	0.218	0.111
Breakdowns after a prior breakdown	0.009	0.008	-0.002
Post-breakdown risky choices per game	0.126	0.163	0.037
Post-breakdown risky-choice rate	0.334	0.352	0.018

Notes: A successful risky choice pays 2,857 points, while a breakdown yields $2,857 - 20,000 = -17,143$ points. Persistent experimenters obtain approximately 0.511 more successful risky choices per game, worth about 1,459 additional points, but experience approximately 0.109 more breakdowns per game, contributing about 1,869 additional points of losses. The net difference is therefore approximately -410 points per game. The extra breakdown exposure occurs almost entirely before any breakdown has previously been observed, rather than through post-breakdown experimentation.

This decomposition clarifies that the lower realized payoff of persistent experimenters in the observable-actions treatment is not driven by a greater tendency to continue experimenting after a known breakdown. Instead, persistent experimenters simply expose themselves more often to the risky arm before a breakdown, producing both more successful risky outcomes and more breakdown losses; in the realized sample, the latter dominate.

A.6 Unit of Observation in the Empirical CDFs

The empirical CDFs in Figure 3 of the main text use the two-player group–game as the unit of observation. There are 650 group–game observations in each treatment (1,300 in total). Experimentation intensity and group payoff are computed for the same group–game observation. Because participants are rematched within matching groups, inferential procedures account for matching-group dependence; the CDF figure itself is descriptive.