Appendix

(For Online Publication)

Appendix A: *Piece-rate* vs. *Tournament* Treatments

In this Appendix, we describe the results from our incentive scheme treatments and explain why we have not included tests of the hypotheses related to the incentive schemes in the main text. In short, our incentive structure in the *tournament* treatment seems not to have been salient enough to induce changes in effort compared to the *piece-rate* treatment. This can be seen clearly in the summary statistics reported in Table A1. Mean learning typically varies only slightly when we compare *tournament* and *piece-rate* within any given treatment, and the sign of the difference is not consistent across treatments. The positive effect of *teaching* is still clearly visible, as is its negative interaction with *tracking* for the bottom half. We report formal statistical analyses below which also suggest that incentive variation had no significant effect on behavior.

Hypothesis A1 (Main Effect of *Tournament*): ambiguous

The effect of competition in the *tournament* treatment relative to the *piece-rate* treatment is theoretically ambiguous. In order to derive a comparative static prediction, we would need to know the optimal effort in the *piece-rate* and *tournament* treatment, but the optimal effort levels depend on the parameters (payoffs) as well as the unobserved costs of effort. As such, we cannot predict the effect of *tournament* treatment on learning relative to the *piece-rate* treatment.

Columns (1) and (2) in Panel A of Table A2 present the estimated treatment effects of competition on learning with the *piece-rate* treatment serving as the control treatment. There is no difference in learning between the two treatments even after controlling the individual covariates in Column (2).[1] Columns (3) and (4) present the estimated treatment effects by initial performance (rank). For both subjects in the top half and bottom half, these estimates are far from statistically significant at conventional levels and small in magnitude. Figure A1 confirms that the cumulative distributions of learning for the full sample, subjects in the bottom half, and subjects in the top half are almost identical in both treatments. The p-value for the two-sample Kolmogorov-Smirnov test for the equality of the distributions for the full sample (Panel A) is 0.618, while the p-value for subjects in bottom half (Panel B) is 0.982.

Hypothesis A2 (Interaction of *Tournament* and *Teaching*): negative

We also hypothesized that there would be interactions between the *tournament* and *teaching* treatments. Specifically, we expected that subjects would be teach each other less in the Practice Block when they anticipated competing against the other group members subsequently in the Evaluation Block—an effect which would be strongest when grouped with subjects of similar ability (i.e., in the *tracked* treatment). If someone is too effective as a teacher, their "student" may surpass their performance in the next stage, lowering their payoff. We found, however, no evidence of any such interaction effects. These and all of our estimates concerning the (non-)effects of the *tournament* treatment are available on request.

Overall, the *tournament* treatment does not seem to have affected learning. One possibility is that the payoff structure by rank in the *tournament* treatment was very nearly linear, and the payoffs were chosen so as to lead to similar earnings in the *piece-rate* treatment and the *tournament* treatment. As

[1] The raw means of learning in the *piece-rate* and *tournament* treatments are very similar (24.4 vs. 25.5 seconds).

such, the tournament incentives may have been insufficient to induce significant changes in behavior. We find similar results in the Nonogram experiments presented in Panel B of Table A2.

Figure A1: Cumulative Distributions of Learning by *Piece-rate* vs. *Tournament* (Sudoku)

A. Full Sample

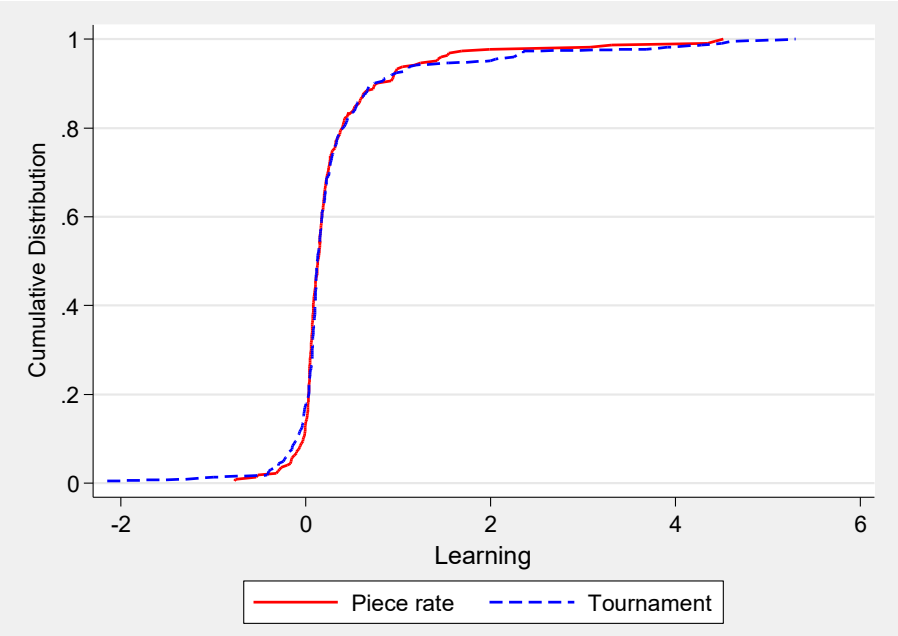

B. Among Subjects in Bottom Half

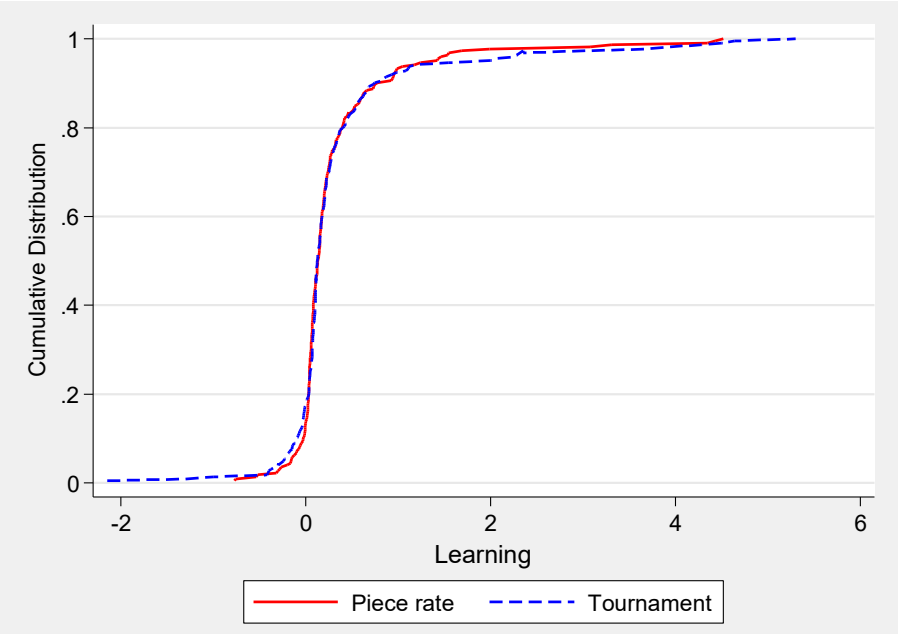

Notes: The cumulative distributions of learning for the *piece-rate* and *tournament* treatments are displayed. Panel A uses the full sample, while Panel B is limited to subjects in bottom half who ranked 5–8 in the Ability Block (T=0). Learning is calculated by subtracting the average solving time (AST) in Evaluation Block (T=1) from that in the Ability Block (T=0) so that higher values indicate *improvement* in solving time. Note that AST in the Evaluation Block (T=1) and in the Ability Block (T=0) is standardized by the mean and standard deviation of raw AST at T=0 before taking the difference so that standardized AST at T=0 has a mean of zero and standard deviation of 1. For Panel A, the p-value for the two-sample Kolmogorov-Smirnov test for the equality of the distributions is 0.618, while that for Panel B is 0.982. There are 448 subjects.

Table A1: Mean Learning by Treatment, Including Incentives

A. Sudoku

		No teaching		*Teaching*	
		Untracked	*Tracked*	*Untracked*	*Tracked*
Bottom half	*Piece rate*	0.340	0.447	0.886	0.371
		(0.109)	(0.164)	(0.159)	(0.159)
	Tournament	0.473	0.594	0.758	0.412
		(0.248)	(0.259)	(0.168)	(0.196)
Top half	*Piece rate*	0.098	0.103	0.098	0.048
		(0.017)	(0.015)	(0.015)	(0.038)
	Tournament	0.092	0.117	0.097	0.096
		(0.019)	(0.027)	(0.015)	(0.026)

B. Nonograms

		No teaching		*Teaching*	
		Untracked	*Tracked*	*Untracked*	*Tracked*
Bottom half	*Piece rate*	0.625	0.533	1.362	0.837
		(0.283)	(0.178)	(0.273)	(0.268)
	Tournament	0.865	0.815	0.982	0.648
		(0.171)	(0.222)	(0.346)	(0.322)
Top half	*Piece rate*	0.103	0.080	0.186	0.104
		(0.027)	(0.011)	(0.047)	(0.014)
	Tournament	0.129	0.166	0.142	0.128
		(0.026)	(0.020)	(0.049)	(0.053)

Notes: For each cell, we report the session level mean and its standard errors. The outcome is learning, which is calculated by subtracting the average solving time (AST) in the Evaluation Block (T=1) from that in the Ability Block (T=0), so that higher values indicate improvement in solving time. Note that AST in the Evaluation Block (T=1) and in the Ability Block (T=0) is standardized by the mean and standard deviation of raw AST at T=0 before taking the difference so that standardized AST at T=0 has a mean of zero and standard deviation of 1.

Table A2: Effect of *Tournament* on Learning

A. Sudoku

	(1)	(2)	(3)	(4)
A. Overall				
Tournament	-0.006	-0.005		
	(0.056)	(0.054)		
B. Heterogeneity				
Tournament for Top half			-0.006	-0.005
			(0.056)	(0.054)
Tournament for Bottom half			-0.006	-0.005
			(0.056)	(0.054)
Controls	No	Yes	No	Yes

B. Nonograms

	(1)	(2)	(3)	(4)
A. Overall				
Tournament	0.006	0.003		
	(0.070)	(0.068)		
B. Heterogeneity				
Tournament for Top half			0.023	0.010
			(0.021)	(0.024)
Tournament for Bottom half			-0.012	-0.007
			(0.134)	(0.131)
Controls	No	Yes	No	Yes

Notes: Each column reports the estimated treatment effects from a different OLS regression. The estimates in Columns (1) and (2) are obtained by estimating equation [2] with the *teaching* dummy being replaced by a *tournament* dummy. The estimates in Columns (3) and (4) are obtained by estimating equation [3] with the *teaching* dummy being replaced by a *tournament* dummy (the original estimates available upon request). The control group is the *piece rate* treatment. The estimated treatment effects and their standard errors were computed using the *lincom* command in STATA. Standard errors clustered at the group level are reported in parentheses. The outcome is learning, which is calculated by subtracting the average solving time (AST) in Evaluation Block (T=1) from that in the Ability Block (T=0) so that higher values indicate *improvement* in solving time. Note that AST in the Evaluation Block (T=1) and in the Ability Block (T=0) is standardized by the mean and standard deviation of raw AST at T=0 before taking the difference so that standardized AST at T=0 has a mean of zero and standard deviation of 1. The subjects in bottom half are those subjects ranked 5–8 in the Ability Block (T=0) and the top half those ranked 1–4. The controls include a dummy for being male, a dummy for being experienced with Sudoku, risk attitudes (0–9), prosociality (0–5), and a dummy for the eight subjects who could not solve any Sudoku puzzles in the Ability Block (T=0). See Table 2 for definitions of each control variable. For Panel A (Sudoku), there were 28 sessions with 224 subjects (8 subjects per session) each for the *piece rate* and *tournament* treatments, respectively. For Panel B (Nonograms), there were 16 sessions with 128 subjects each for the *piece rate* and *tournament* treatments, respectively. Each session consisted of two groups (4 subjects per group). Significance levels: *** p<0.01, ** p<0.05, * p<0.10

Appendix B: Additional Figures and Tables

Figure B1: Different measures of learning (Sudoku)

A. Δ in # of Sudoku solved

B. Implied percent change in average solve time (for a 2-puzzle improvement)

C. Δ in average solve time (our learning measure)

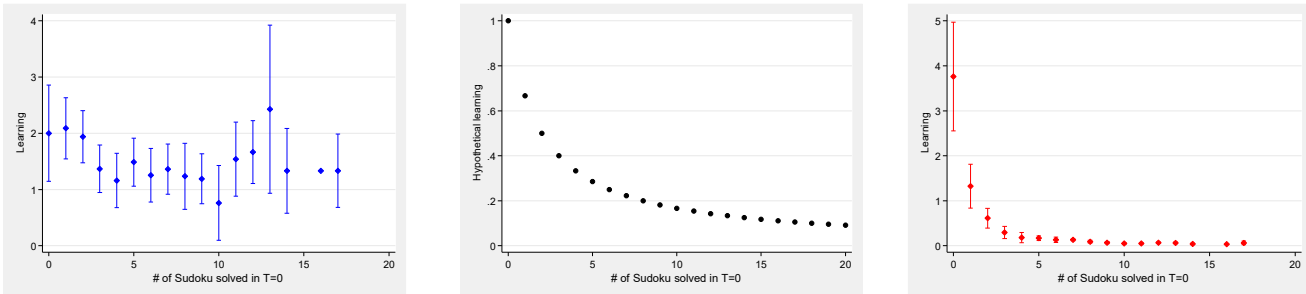

Notes: Panel A plots the average change in the number of Sudoku solved per 10 min between T=0 and T=1 together with 95% confidence intervals for subjects with a given number solved at T=0. Panel B shows the required relative improvement in average solve time necessary to produce a 2-puzzle improvement for individuals solving different numbers of puzzles at T=0, assuming that all the subjects spent the full time solving puzzles in T=0 and T=1. Suppose x is the number solved at T=0, the implied percent improvement can be written as y= [(10/x)-(10/(x+2))]/(10/x) (where it takes one as x approaches 0). Panel C plots our preferred learning measure: the change in standardized average solve time.

Here we provide more details on the validity of our chosen measure of learning. There are three advantages over using a naïve measure like change in puzzles solved (per 10 minutes).

First, large improvements in solving time do not always translate into significant changes in the number of puzzles solved given that subjects solve complex, time-consuming 6×6 Sudoku facing binding time-constraints.[2] The lack of variation makes it difficult to use changes in the number of Sudoku solved to identify meaningful effects of treatments.[3]

Second, the change in the number of puzzles solved masks variation in relative improvement in performance across the distribution of performance at T=0. Given that we expected ex ante that differences in learning and treatment effects might exist across the ability distribution, this is a significant reason to study changes in average solve time. To illustrate the point, consider two hypothetical subjects with identical improvement in the number of puzzles solved: Person A who solved 1 puzzle at T=0 and improved to 3 puzzles (per 10 minutes) at T=1, and Person B who solved 10 puzzles at T=0 and improved to 12 puzzles (per 10 minutes) at T=1. Assuming these subjects spent the whole time period working on the puzzles solved implies that Person A improved from 600 seconds per puzzle

[2] Indeed, the observed variation in the change in the number of Sudoku puzzles solved is much smaller than that of changes in average solving time: the coefficients of variation (COV), which divide a variable's standard deviation by its mean, are 2.36 (=73.5/31.1) and 0.55 (=2.8/5.1) for changes in average solving time and the number of Sudoku puzzles solved, respectively.

[3] Moreover, using the number of Sudoku puzzles correctly solved as an outcome measure makes no distinction between subjects who barely finish N puzzles and those who run out of time just before correctly completing the N+1th puzzle, while also implying a difference in performance between subjects who barely complete N puzzles and those who were about to finish the Nth puzzle when time expired. Using average solving time for correctly solved problems allows us to accurately compare such subjects.

to 200 seconds per puzzle, while Person B, who had the same improvement in number solved, improved from 60 seconds per puzzle to 50 seconds per puzzle. The same improvement in number solved requires vastly different improvements in average solve time for people of differing initial abilities.

Figure B1 summarizes both of the above points. Panel (a) shows, for each bin of number solved at T=0 in our data, the average change in the number of Sudoku solved per 10 min between T=0 and T=1, and the associated 95% confidence intervals. There is no obvious relationship between these variables, and we see limited variation in the change in the number of Sudoku solved across the entire initial ability distribution. On average, subjects across the initial ability distribution improved by only 1-2 puzzles per 10 minutes from T=0 to T=1. In panel (b), we plot the implied percent change in average solve time necessary to produce a 2-puzzle-per-10-minutes improvement, across the same distribution of initial performance: $y = [10/x - 10/(x+2)]/(10/x)$, where x is the number solved at T=0. Panel (c) plots our learning measure, which has a very similar shape to the hypothetical relative improvement in panel (b). Panel (c) highlights that our learning measure based on the change in average solve time exhibits considerable variation across the initial ability distribution. In particular, it captures the difference in *relative* improvement in performance for those who were lower performers in the Ability block (T=0).[4] Using the change in the number solved (Panel (a)) obscures large improvements in performance among low-ability (slow) subjects at T=0—precisely those subjects who learn the most. By contrast, our measure based on change in average solve time (Panel (c)) reveals this heterogeneity in learning clearly.

Third, measuring performance using solve time is consistent with studies of performance on complex puzzles when subjects face time limits; see, for example, Charness et al. [2015] and Koellinger and Block [2016]. In particular, Charness et al. note that, as a measure of performance, speed is less sensitive to arbitrary time limits when puzzles are time-consuming.

[4] The leftmost dot captures those who solved 0 puzzles at T=0 and, though it stands out visually, in our regression we include a dummy for those subjects in the analysis so that our main results do not depend on these extreme observations.

Figure B2: Cumulative Distributions of Learning by *Tracking* in the *Teaching* Treatment Separately for Top and Bottom Half Subjects (Sudoku)

A. Among Subjects in Top Half

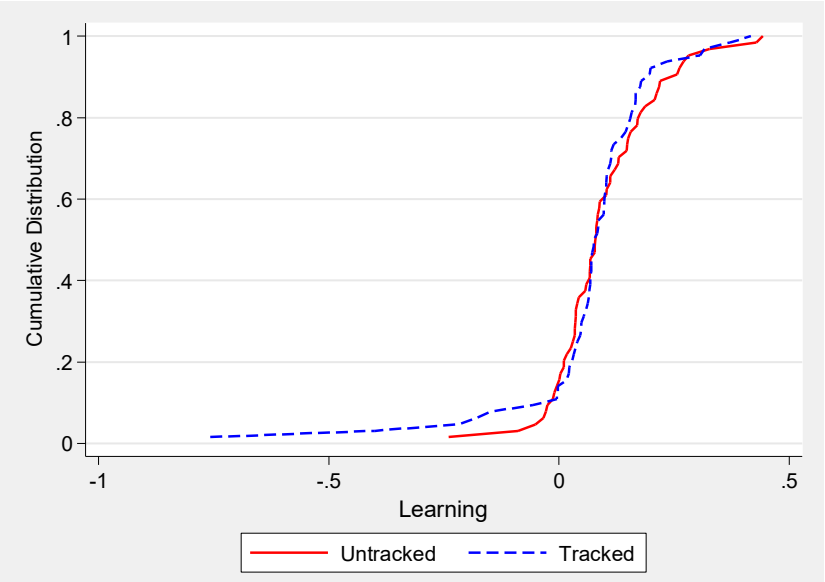

B. Among Subjects in Bottom Half

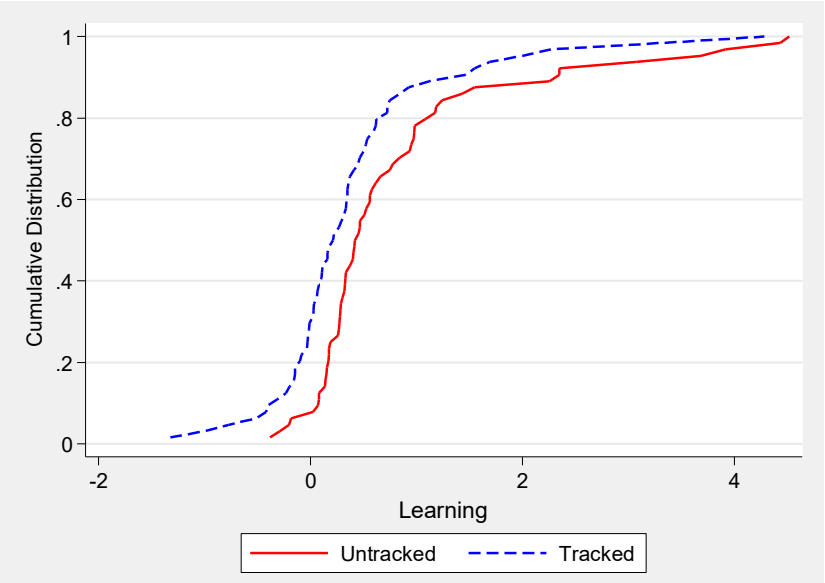

Notes: The cumulative distributions of learning reported in Figure 5-B are separately displayed for subjects in the top half (Panel A) and subjects in the bottom half (Panel B). Learning is calculated by subtracting the average solving time (AST) in the Evaluation Block (T=1) from that in the Ability Block (T=0) so that higher values indicate *improvement* in solving time. Note that AST in the Evaluation Block (T=1) and in the Ability Block (T=0) is standardized by the mean and standard deviation of raw AST at T=0 before taking the difference so that standardized AST at T=0 has a mean of zero and standard deviation of 1. Also note that the scales on the *x*-axes are different for the two graphs. The sample is limited to 32 *teaching* sessions with 256 subjects. For each session, there is one group for the top half and the bottom half; thus there are 32 groups each for the top half and bottom half. For Panel A, the p-value for the two-sample Kolmogorov-Smirnov test for the equality of the distributions in the *untracked* and *tracked* treatments is 0.843, while that for Panel B is 0.004.

Figure B3: Screenshot from the *Teaching* Treatment during the Practice Block (Nonograms)

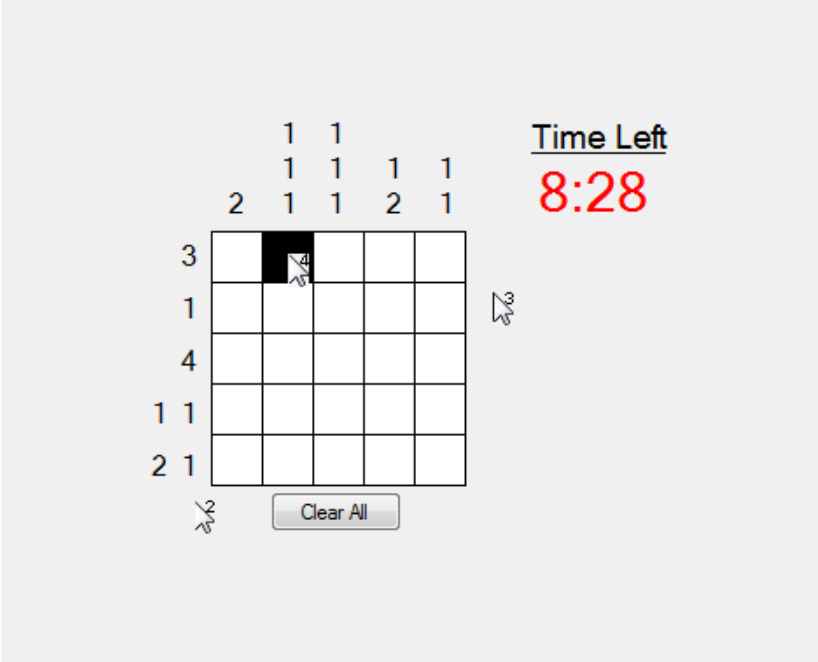

Notes: Subjects are able to simultaneously edit a common 5×5 Nonogram puzzle during the Practice Block. Each mouse arrow is labeled with the within-group performance rank of the person in the Ability Block (T=0). Performance is measured by the number of Nonogram puzzles solved with the average solving time serving as a tie-breaker. In the *no-teaching* treatment, the three arrows showing the within-group rank of the other subjects would not have been visible as each subject worked independently.

Figure B4: Cumulative Distributions of Learning by *Tracking* in the *Teaching* Treatment Separately for Top and Bottom Half Subjects (Nonograms)

A. Among Subjects in the Top Half

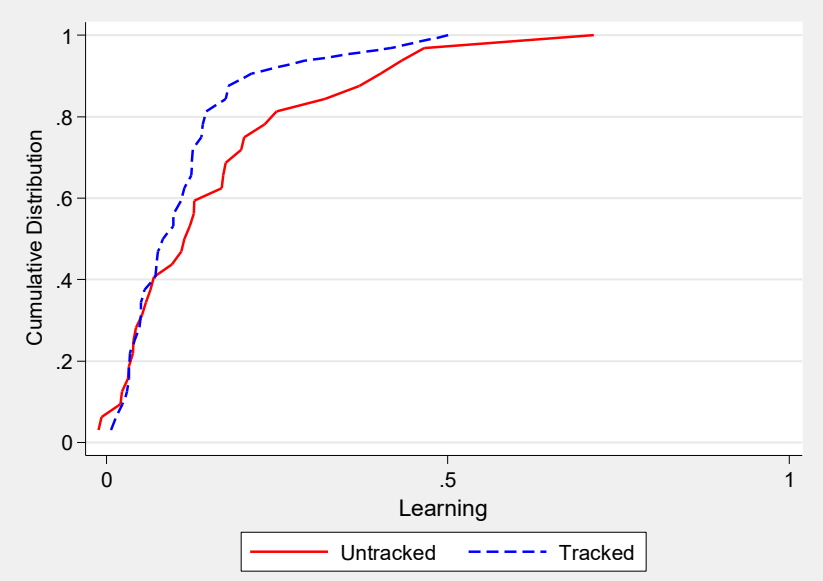

B. Among Subjects in the Bottom Half

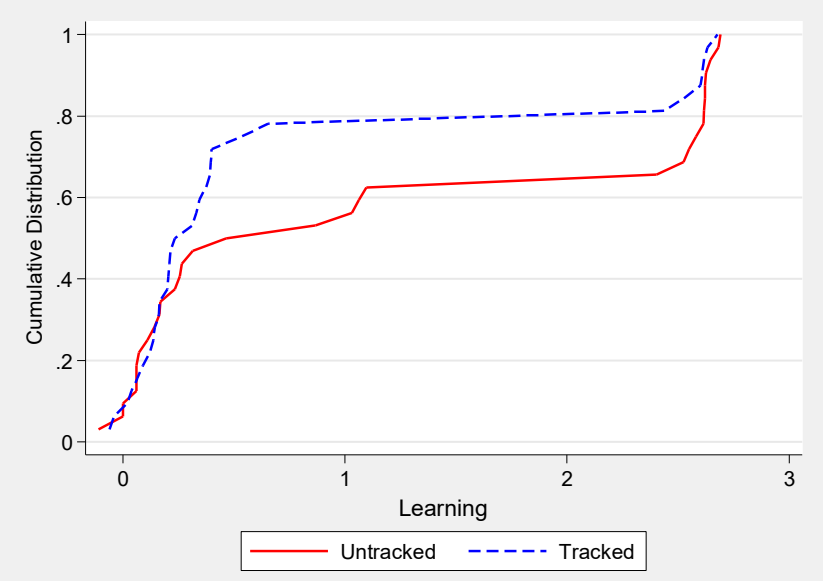

Notes: The cumulative distributions of learning reported in Figure 8-B are separately displayed for subjects in the top half (Panel A) and subjects in the bottom half (Panel B). Learning is calculated by subtracting the average solving time (AST) in Evaluation Block (T=1) from that in the Ability Block (T=0) so that higher values indicate *improvement* in solving time. Note that AST in the Evaluation Block (T=1) and in the Ability Block (T=0) is standardized by the mean and standard deviation of raw AST at T=0 before taking the difference so that standardized AST at T=0 has a mean of zero and standard deviation of 1. Also, note that the scales on the *x*-axes are different between the two graphs. The sample is limited to 16 *teaching* sessions with 144 subjects. For each session, there is one group of subjects in the top half and one group of subjects in the bottom half; thus there are 16 groups each for the top half and bottom half. For Panel A, the p-value for the two-sample Kolmogorov-Smirnov test for the equality of the distributions in the *untracked* and *tracked* treatments is 0.434, while that for Panel B is 0.160.

Table B1: Balance tests (Sudoku)

	Bivariate regression		Equality test	
	Teaching	Tracking	2×2×2	2×2
			(*p-value*)	
Variable	(1)	(2)	(3)	(4)
Male	-0.02	0.01	*0.77*	*0.96*
	(0.05)	(0.05)		
Experienced	0.04	0.03	*0.65*	*0.70*
	(0.05)	(0.04)		
Risk Attitudes (0-9)	0.26*	-0.01	*0.35*	*0.17*
	(0.16)	(0.16)		
Prosociality (0-5)	0.10	-0.13	*0.39*	*0.27*
	(0.10)	(0.10)		
Solved none at T=0	-0.01	0.00	*0.23*	*0.11*
	(0.01)	(0.01)		
[Raw] Average solve time at T=0 (sec)	6.32	-9.29	*0.31*	*0.15*
	(9.24)	(9.20)		
[Standardized] Average solve time at T=0	0.07	-0.10	*0.31*	*0.15*
	(0.10)	(0.09)		
# of Sessions	56	56	56	56
# of Group	112	112	112	112
# of Subjects	448	448	448	448

Notes: Columns (1) and (2) report a set of bivariate regressions that test how each variable in the far-left column is related to the teaching treatment (Column 1) and to the tracking treatment (Column 2). Standard errors are reported in parenthesis. Columns (3) and (4) report the p-values for each variable in the far-left column of the null hypotheses that the means are equal across 8 treatment combinations (Column 3) and 4 treatment combinations pooling across the incentive treatments (Column 4). Experienced takes a value of one if a subject reports having prior experience with Sudoku. Risk attitudes take on the values from 0 to 9 with higher numbers indicating more risk-loving subjects. Prosociality takes on the values from 0 to 5 with higher numbers indicating higher prosociality. See Appendix D for details on the elicitation of risk attitudes and prosociality and Appendix E for screenshots. There were 56 sessions with 448 subjects (8 subjects per session). Each session consisted of two groups (4 subjects per group). Significance levels: *** p<0.01, ** p<0.05, * p<0.10

Table B2: Coefficient Estimates for Equations [2] and [3] (Sudoku)

Outcome: Learning

	A. Overall		B. Heterogeneity	
	(1)	(2)	(3)	(4)
Teaching	0.113**	0.119**	-0.018	-0.016
	(0.052)	(0.052)	(0.017)	(0.020)
Bottom Half			0.185***	0.167***
			(0.053)	(0.052)
Teaching × Bottom Half			0.258***	0.256***
			(0.096)	(0.095)
Male		0.041		0.007
		(0.059)		(0.057)
Experienced		-0.203***		-0.066
		(0.076)		(0.069)
Risk Attitudes (0–9)		0.010		0.015
		(0.020)		(0.019)
Prosociality (0–5)		-0.007		-0.024
		(0.030)		(0.031)
None Correct at T=0	3.528***	3.416***	3.385***	3.356***
	(0.386)	(0.395)	(0.386)	(0.387)
Constant	0.191***	0.286**	0.102***	0.142
	(0.028)	(0.128)	(0.010)	(0.112)
Controls	No	Yes	No	Yes
R-squared	0.38	0.40	0.44	0.44
# of Sessions	56	56	56	56
# of Groups	112	112	112	112
# of Subjects	448	448	448	448

Notes: Each column reports the results from a different OLS regression. Standard errors clustered at the group level are reported in parentheses. The outcome is learning, which is calculated by subtracting the average solving time (AST) in Evaluation Block (T=1) from that of Ability Block (T=0) so that higher values indicate *improvement* in solving time. Note that AST in the Evaluation Block (T=1) and in the Ability Block (T=0) is standardized by the mean and standard deviation of raw AST at T=0 before taking the difference so that standardized AST at T=0 has a mean of zero and standard deviation of 1. Experienced takes a value of one if a subject indicates experience with Sudoku puzzles before the experiment. Risk attitudes take on the values from 0 to 9 with higher number indicating higher risk-aversion. Prosociality takes on the values from 0 to 5 with higher number indicating higher prosociality. See Appendix C for details on the elicitation of risk attitude and prosociality. We also include a dummy for the eight subjects who could not solve any Sudoku puzzles in the Ability Block (T=0). There were 56 sessions with 448 subjects (8 subjects per session) in the experiment. The estimates were used to produce Table 4 in the main text. Significance levels: *** p<0.01, ** p<0.05, * p<0.10

Table B3: Coefficient Estimates for Equations [4] and [5] (Sudoku)

Outcome: Learning

	A. Overall	B. Heterogeneity
	(1)	(2)
Tracked	-0.039	0.024
	(0.061)	(0.026)
Teaching	0.171**	0.012
	(0.073)	(0.024)
Tracked × *Teaching*	-0.105	-0.056
	(0.101)	(0.044)
Bottom half		0.231***
		(0.085)
Tracked × Bottom half		-0.119
		(0.107)
Teaching × Bottom half		0.322**
		(0.127)
Tracked × *Teaching* × Bottom half		-0.132
		(0.183)
Controls	Yes	Yes
R-squared	0.41	0.45
# of Sessions	56	56
# of Groups	112	112
# of Subjects	448	448

Notes: Each column reports the results from a different OLS regression. The outcome is learning, which is calculated by subtracting the average solving time (AST) in Evaluation Block (T=1) from that in the Ability Block (T=0) so that higher values indicate *improvement* in solving time. Note that AST in the Evaluation Block (T=1) and in the Ability Block (T=0) is standardized by the mean and standard deviation of raw AST at T=0 before taking the difference so that standardized AST at T=0 has a mean of zero and standard deviation of 1. The controls include a dummy for being male, a dummy for being experienced with Sudoku, risk attitudes (0–9), prosociality (0–5), and a dummy for the subjects who could not solve any Sudoku puzzles at T=0. The bottom half consists of those subjects ranked 5–8. There were 24 *no-teaching* sessions with 192 subjects, and 32 *teaching* sessions with 256 subjects (8 subjects per session). Each session consisted of two groups (4 subjects per group). The estimates were used to produce Table 5 in the main text. Significance levels: *** p<0.01, ** p<0.05, * p<0.10

Table B4: Effect of *Teaching* on Logged Learning, Robustness Checks (Sudoku)

Outcome: Learning (logged)

	(1)	(2)	(3)	(4)
A. Overall				
Teaching	0.030	0.031		
	(0.026)	(0.027)		
B. Heterogeneity				
Teaching for Top half			-0.021	-0.024
			(0.021)	(0.022)
Teaching for Bottom half			0.080**	0.079**
			(0.046)	(0.046)
Controls	No	Yes	No	Yes

Notes: Each column reports the results from a different OLS regression. Columns (1) and (2) come from equation [2] with and without controls using the full sample. Columns (3) and (4) come from the equation [3] with and without controls using the full sample. Here, the control group is the *no-teaching* treatment. The estimated treatment effects and their standard errors reported in the table were computed using the *lincom* command in STATA. The coefficient estimates from equations [2] and [3] are not reported to save space (results available upon request). Standard errors clustered at the group level are reported in parentheses. The outcome is logged learning, which is defined as the difference between logged average solving time (AST) in the Ability Block (T=0) and the Evaluation Block (T=1). The bottom half consists of those subjects ranked 5–8 and the top half those subjects ranked 1–4 in T=0. All regressions—even those labeled as including "no" controls—include a dummy for the eight subjects who could not solve any Sudoku puzzles at T=0. The controls further include a dummy for being male, a dummy for being experienced with Sudoku, risk attitudes (0–9), and prosociality (0–5). See Table 2 for definitions of each control variable. There were 56 sessions with 448 subjects (8 subjects per session). Each session consisted of two groups (4 subjects per group), and thus there were 112 groups. Significance levels for one-sided test: *** p<0.01, ** p<0.05, * p<0.10

Table B5: Effect of *Tracking* on Logged Learning, Robustness Checks (Sudoku)

Outcome: Learning (logged)

	In the *No-teaching* Treatment		In the *Teaching* Treatment	
	(1)	(2)	(3)	(4)
A. Overall				
Tracked	-0.025		-0.081**	
	(0.035)		(0.040)	
B. Heterogeneity				
Tracked for Top half		0.022		-0.021
		(0.026)		(0.034)
Tracked for Bottom half		-0.068		-0.152**
		(0.060)		(0.068)
Controls	Yes	Yes	Yes	Yes

Notes: The estimated treatment effects in Columns (1) and (3) come from equation [4], while the estimated treatment effects in Columns (2) and (4) come from equation [5]. The control treatment is the *untracked* treatment. The estimated treatment effects and their standard errors were computed using the *lincom* command in STATA. Standard errors clustered at the group level are reported in parentheses. The outcome is logged learning, which is defined as the difference between logged average solving time (AST) in the Ability Block (T=0) and logged AST in the Evaluation Block (T=1). Note that learning is calculated by subtracting logged AST at T=1 from that at T=0, so that higher values indicate *improvement* in solving time. The controls include a dummy for being male, a dummy for being experienced with Sudoku, risk attitudes (0–9), prosociality (0–5), and a dummy for the eight subjects who could not solve any Sudoku puzzles in the Ability Block (T=0). The bottom half consists of those subjects ranked 5–8 in the Ability Block (T=0). There were 24 *no-teaching* sessions with 192 subjects, and 32 *teaching* sessions with 256 subjects. Each session consisted of two groups (4 subjects per group). Significance levels: *** p<0.01, ** p<0.05, * p<0.10

Table B6: Group Mean, Group Standard Deviation, and Teaching Frequencies

	A. Sudoku				B. Nonograms			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Group mean	0.36***		0.02	0.28	0.20		0.04	1.54
	(0.10)		(0.21)	(0.32)	(0.16)		(0.35)	(0.98)
Group SD		0.36***	0.34**	0.40**		0.22	0.18	-0.42
		(0.09)	(0.17)	(0.18)		(0.16)	(0.35)	(0.50)
Group mean × Group SD				-0.18				-1.34
				(0.17)				(0.84)
Vuong test of Zero-Inflated model vs. Standard Poisson	z=3.61	3.57	3.53	3.31	2.91	2.83	2.74	2.83
	p=0.0002	0.0002	0.0002	0.0005	0.0018	0.0023	0.0031	0.0023
# of Groups	64	64	64	64	30	30	30	30

Notes: Each column reports the coefficient estimates from a different zero-inflated Poisson regression in which the number of teaching statements recorded for a given group is the dependent variable. Group mean is the group average of *raw* average solving time (AST) in the Ability Block (T=0) so that higher values indicate that the group consisted of lower-ability subjects while lower values indicate that the group consisted of higher-ability subjects. Because this is a group-level analysis, the sample size for Sudoku in Panel A is 64, while that for Nonograms in Panel B is 30 as there was one Nonogram session (two groups) in which the audio recording did not work. The Vuong tests of the zero-inflated model against standard Poisson models are reported in each column with the z-scores and corresponding p-values. All of these tests support the use of the zero-inflated model.

Table B7: Summary Statistics and Balance Tests (Nonograms)

	A. Overall			B. Heterogeneity		
	Mean	*p-value* of equality test		Bottom half (rank5-8)	Top half (rank1-4)	Dif (5)-(6)
		2×2×2	2×2			
Variable	(1)	(2)	(3)	(4)	(5)	(6)
Male	0.45	*0.60*	*0.43*	0.41	0.49	-0.09
	[0.52]			[0.49]	[0.55]	(0.07)
Experienced	0.02	*0.33*	*0.57*	0.00	0.03	-0.03**
	[0.12]			[0.00]	[0.17]	(0.02)
Risk Attitude (0–9)	3.63	*0.12*	*0.14*	3.63	3.63	0.00
	[1.60]			[1.60]	[1.60]	(0.18)
Prosociality (0–5)	1.79	*0.28*	*0.13*	1.85	1.72	0.13
	[0.97]			[0.99]	[0.94]	(0.11)
Solved None at T=0	0.20	*0.47*	*0.18*	0.39	0.00	0.39***
	[0.40]			[0.49]	[0.00]	(0.04)
Solved None at T=1	0.07	-	-	0.13	0.00	0.13***
	[0.25]			[0.34]	[0.00]	(0.03)
Raw Average solve time at T=0 (*sec*)	186.47	*0.26*	*0.14*	307.31	65.64	241.66***
	[209.92]			[240.90]	[31.91]	(21.47)
Raw Average solve time at T=1 (*sec*)	85.38	-	-	132.37	38.40	93.97***
	[140.16]			[186.77]	[10.46]	(16.23)
Raw Learning (=AST0-AST1) (*sec*)	101.09	-	-	174.94	27.25	147.69***
	[170.62]			[216.42]	[25.00]	(19.02)
Standardized average solve time at T=0	0.00	*0.26*	*0.14*	0.58	-0.58	1.15***
	[1.00]			[1.15]	[0.15]	(0.10)
Standardized average solve time at T=1	-0.48	-	-	-0.26	-0.71	0.45***
	[0.67]			[0.89]	[0.05]	(0.08)
Learning	0.48	-	-	0.83	0.13	0.70***
	[0.81]			[1.03]	[0.12]	(0.09)
# of Sessions	32			32	32	
# of Groups	64			32	32	
# of Subjects	256			128	128	

Notes: Column (1) reports means for the full sample with standard deviations in brackets. Columns (2) and (3) report the p-values for each variable in the far-left column of the null hypotheses that the means are equal across 8 treatment combinations (Column (2)) and 4 treatment combinations pooling across the incentive treatments (Column (3)). Columns (4) and (5) report the means by ranks in the Ability Block (T=0). The bottom half consists of those subjects ranked 5–8, and the top half consists of those subjects ranked 1–4. Column (6) reports the difference in means between subjects in the top half and subjects in the bottom half with standard errors clustered at the group level in parentheses. See the notes for Table 2 for descriptions of the variables. The experienced variable equals one if a subject reported prior experience with Nonograms and zero otherwise. There were total of 32 sessions with 256 subjects (8 subjects per session). Each session consisted of two groups (4 subjects per group). Significance levels: *** p<0.01, ** p<0.05, * p<0.10

Table B8: Mean Learning by Treatment (Nonograms)

A. Teaching

	No teaching	*Teaching*
Bottom half	0.709	0.957
	(0.104)	(0.152)
Top half	0.119	0.140
	(0.013)	(0.021)

B. Teaching x Tracking

	No teaching		*Teaching*	
	Untracked	*Tracked*	*Untracked*	*Tracked*
Bottom half	0.745	0.674	1.172	0.742
	(0.160)	(0.142)	(0.216)	(0.197)
Top half	0.116	0.123	0.164	0.116
	(0.018)	(0.020)	(0.032)	(0.026)

Notes: For each cell, we report the session level mean and its standard errors. The outcome is learning, which is calculated by subtracting the average solving time (AST) in the Evaluation Block (T=1) from that in the Ability Block (T=0), so that higher values indicate improvement in solving time. Note that AST in the Evaluation Block (T=1) and in the Ability Block (T=0) is standardized by the mean and standard deviation of raw AST at T=0 before taking the difference so that standardized AST at T=0 has a mean of zero and standard deviation of 1.

Table B9: Coefficient Estimates for Equations [2] and [3] for the Effect of *Teaching* on Learning (Nonograms)

Outcome: Learning

	A. Overall		B. Heterogeneity	
	(1)	(2)	(3)	(4)
Teaching	0.203***	0.195***	0.021	0.022
	(0.064)	(0.064)	(0.022)	(0.021)
Bottom Half			-0.018	-0.006
			(0.077)	(0.081)
Teaching × Bottom Half			0.358***	0.337***
			(0.121)	(0.121)
Male		0.040		0.039
		(0.071)		(0.072)
Experienced		-0.279***		-0.166***
		(0.062)		(0.039)
Risk Attitudes (0–9)		0.003		0.002
		(0.025)		(0.026)
Prosociality (0–5)		-0.042		-0.038
		(0.041)		(0.041)
None Correct at T=0	0.091**	0.145	0.119***	0.163
	(0.036)	(0.152)	(0.011)	(0.148)
Constant	1.477***	1.481***	1.390***	1.394***
	(0.172)	(0.171)	(0.180)	(0.179)
Controls	No	Yes	No	Yes
R-squared	0.53	0.53	0.55	0.55
# of Sessions	32	32	32	32
# of Groups	64	64	64	64
# of Subjects	256	256	256	256

Notes: Each column reports the results from a different OLS regression. The outcome is learning, which is calculated by subtracting the average solving time (AST) in the Evaluation Block (T=1) from that in the Ability Block (T=0) so that higher values indicate *improvement* in solving time. Note that AST in the Evaluation Block (T=1) and in the Ability Block (T=0) is standardized by the mean and standard deviation of raw AST at T=0 before taking the difference so that standardized AST at T=0 has a mean of zero and standard deviation of 1. The controls include a dummy for being male, a dummy for being experienced with Nonograms, risk attitudes (0–9), prosociality (0–5), and a dummy for the subjects who could not solve any Nonogram puzzles at T=0. The bottom half consists of those subjects ranked 5–8 in the Ability Block (T=0). There were 16 *no-teaching* and *teaching* sessions with 128 subjects, respectively (8 subjects per session). Each session consisted of two groups (4 subjects per group). The estimates were used to produce Table 8 in the main text. Significance levels: *** p<0.01, ** p<0.05, * p<0.10

Table B10: Coefficient Estimates for Equations [4] and [5] (Nonograms)

Outcome: Learning

	A. Overall	B. Heterogeneity
	(1)	(2)
Tracked	0.067	0.024
	(0.095)	(0.044)
Teaching	0.260***	0.066**
	(0.093)	(0.029)
Tracked × *Teaching*	-0.131	-0.087*
	(0.128)	(0.050)
Bottom half		-0.040
		(0.129)
Tracked × Bottom half		0.071
		(0.186)
Teaching × Bottom half		0.381**
		(0.187)
Tracked × *Teaching* × Bottom half		-0.089
		(0.243)
Controls	Yes	Yes
R-squared	0.53	0.55
# of Sessions	32	32
# of Groups	64	64
# of Subjects	256	256

Notes: Each column reports the results from a different OLS regression. The outcome is learning, which is calculated by subtracting the average solving time (AST) in the Evaluation Block (T=1) from that in the Ability Block (T=0) so that higher values indicate *improvement* in solving time. Note that AST in the Evaluation Block (T=1) and in the Ability Block (T=0) is standardized by the mean and standard deviation of raw AST at T=0 before taking the difference so that standardized AST at T=0 has a mean of zero and standard deviation of 1. The controls include a dummy for being male, a dummy for being experienced with Nonograms, risk attitudes (0–9), prosociality (0–5), and a dummy for the subjects who could not solve any Nonogram puzzles at T=0. The bottom half consists of those subjects ranked 5–8 in the Ability Block (T=0). There were 24 *no-teaching* sessions with 192 subjects, and 32 *teaching* sessions with 256 subjects (8 subjects per session). Each session consisted of two groups (4 subjects per group). The estimates were used to produce Table 9 in the main text. Significance levels: *** p<0.01, ** p<0.05, * p<0.10

Table B11: Frequency of Teaching in the *Tracked* vs. *Untracked* Treatments (Nonograms)

		Tracked		Difference (2)–(1)		Difference (3)–(1)	
	Untracked	Bottom Half	Top Half	OLS	Zero-Inflated Poisson	OLS	Zero-Inflated Poisson
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	4.14	5.75	4.00	1.61	1.61	-0.14	-0.14
	[5.26]	[7.13]	[2.98]	(2.65)	(1.82)	(2.04)	(1.54)
Vuong test of Zero-Inflated model vs. Standard Poisson					z= 2.64		z= 2.66
					p= 0.0041		p= 0.0039
# of Groups	14	8	8	22	22	22	22
# of Sessions	7	8		15	15	15	15

Notes: The unit of observation is a group. The sample is limited to the 15 *teaching* treatment sessions with 7 sessions for the *untracked* and 8 sessions for *tracked* treatments as there was one *untracked* session in which the audio recording did not work. For the *untracked* treatment, there are total of 14 groups (two groups for each session), while for the *tracked* treatment there are 8 groups each for subjects in bottom half (Group 2 in the *tracked* treatment in Figure 2) and for those in top half (Group 1 in the *tracked* treatment in Figure 2). Column (1) reports the mean number of teaching statements in the *untracked* treatment, and Columns (2) and (3) report the means for the *tracked* treatment for the bottom half group and the top half group, respectively. Standard deviations are reported in brackets. Columns (4) and (5) report the estimated difference between Columns (1) and (2) from OLS and zero-inflated Poisson (where the inflation equation includes just a dummy for whether the group was *tracked*) models, respectively, with standard errors in parentheses. Columns (6) and (7) report the corresponding estimated differences between Columns (1) and (3). The bottom half consists of those subjects ranked 5–8 in the Ability Block (T=0) and the top half consists of those subjects ranked 1–4. A teaching statement is defined to be any utterance in which subjects are engaged in trying to teach each other how to do Nonograms such as "You can't have a five there; there is already one in that column." The Vuong tests of the zero-inflated Poisson models against the standard Poisson models are reported with the z-scores and corresponding p-values; these tests support the use of the Zero-Inflated model. Significance levels for one-sided test: *** p<0.01, ** p<0.05, * p<0.10

Appendix C: Pooling Sudoku and Nonogram samples

Figure C1: Cumulative Distributions of Learning in the *No-teaching* and *Teaching* Treatments (Pooled)

A. Full Sample

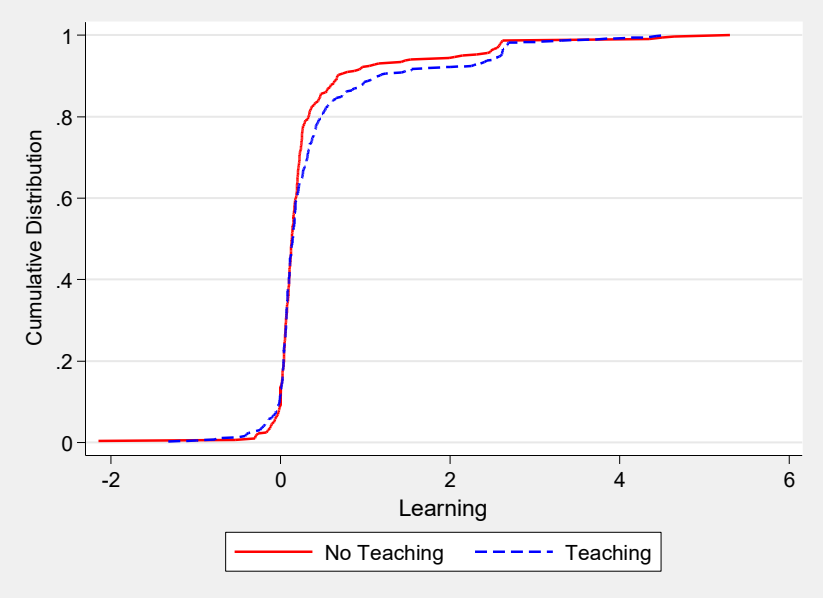

B. Among Subjects in the Bottom Half

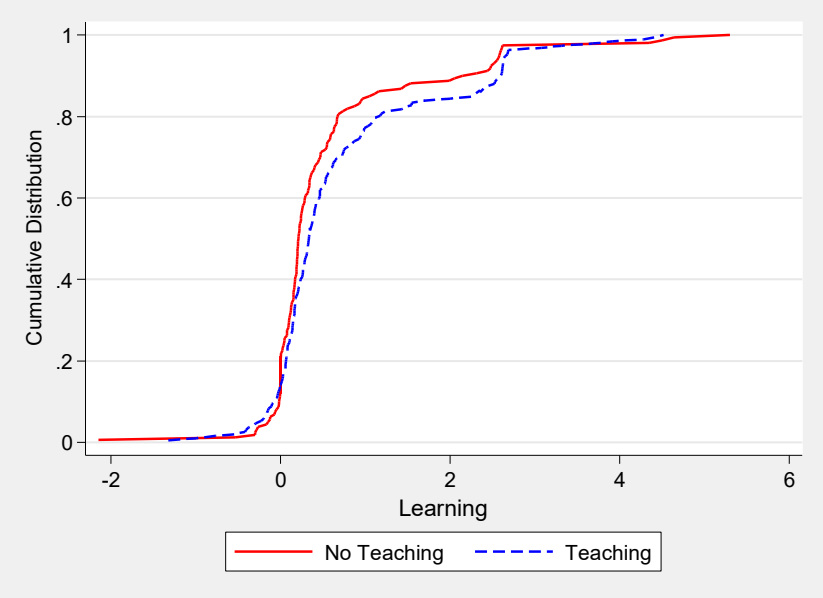

Notes: Kernel density plots of learning in the *no-teaching* and *teaching* treatments are displayed. Panel A is for the full sample while Panel B is restricted to subjects in the bottom half who ranked 5–8 in the Ability Block (T=0). Learning is calculated by subtracting the standardized average solving time in the Evaluation Block (T=1) from that in the Ability Block (T=0), so that higher values indicate *improvement* in average solving time. For Panel A, the p-value for the two-sample Kolmogorov-Smirnov test for the equality of the distributions between the *no-teaching* and *teaching* treatments is 0.013, while that for Panel B is 0.011. There are a total of 704 subjects (448 from Sudoku, and 256 from Nonograms) in Panel A and 352 subjects (224 from Sudoku, and 128 from Nonograms) in Panel B.

Figure C2: Cumulative Distributions of Learning in the *Tracked* and *Untracked* Treatments (Pooled)

A. In the *No-teaching* Treatment

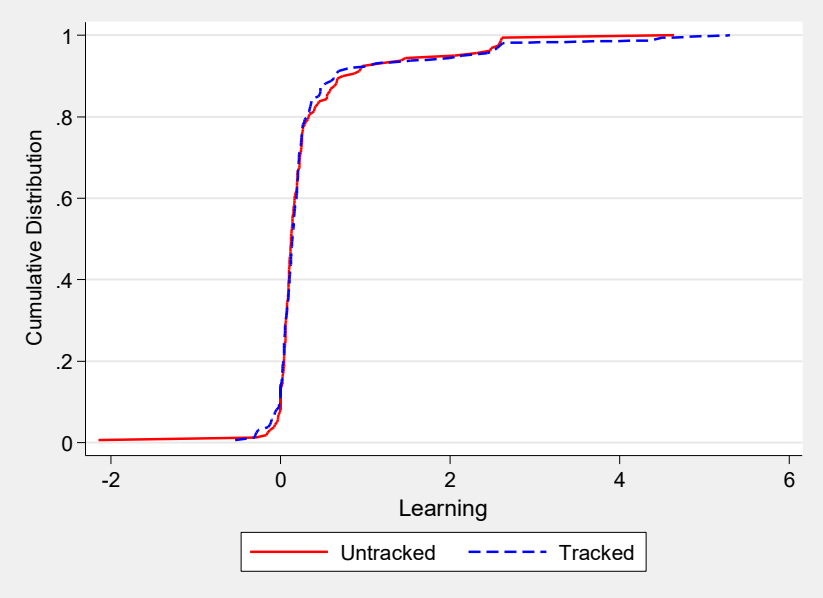

B. In the *Teaching* Treatment

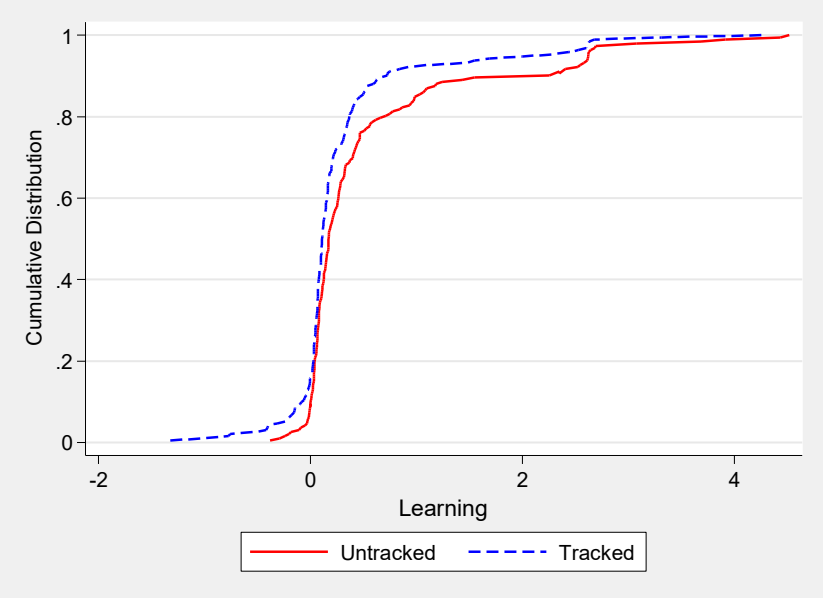

Notes: Cumulative distributions of learning for the *untracked* and *tracked* treatments are displayed. Learning is calculated by subtracting the standardized average solving time in the Evaluation Block (T=1) from that in the Ability Block (T=0), so that higher values indicate *improvement* in average solving time. Panel A plots the cumulative distributions in the *no-teaching* treatment where the p-value for a two-sample Kolmogorov-Smirnov test for equality of the distributions between the *untracked* and *tracked* treatments is 0.915. Panel B plots the cumulative distributions in the *teaching* treatment where the p-value for the Kolmogorov-Smirnov test between the *untracked* and *tracked* treatments is 0.010. There are 40 *no-teaching* sessions (24 from Sudoku and 16 from Nonograms) with 320 subjects, and 48 *teaching* sessions (32 from Sudoku and 16 from Nonograms) with 384 subjects.

Figure C3: Frequency of Teaching (Pooled)

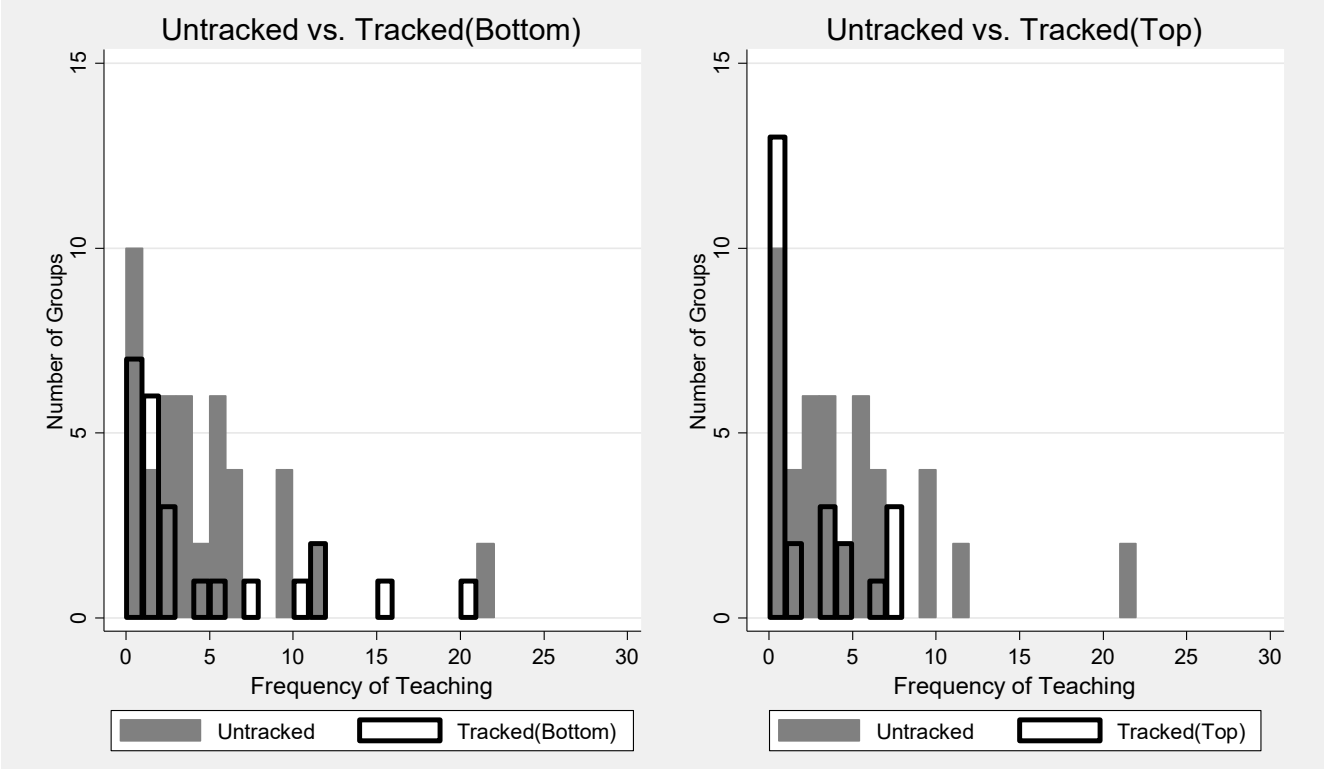

Notes: The unit of observation is a group. The sample is limited to the 48 *teaching* treatment sessions with 24 sessions each for the *untracked* and *tracked* treatments. In the *untracked* treatment, there are total of 48 groups (two groups for each session), while for the *tracked* treatment there are 24 groups each for subjects in the bottom half (Group 2 in the *tracked* treatment in Figure 2 consisting of subjects ranked 5–8 in the Ability Block (T=0)) and for subjects in the top half (Group 1 in the *tracked* treatment in Figure 2 consisting of subjects ranked 1-4 in the Ability Block (T=0)). The left graph plots the number of groups on the vertical axis exhibiting a given frequency of teaching on the horizontal axis for groups in the *untracked* treatment (N=46) and for groups consisting of subjects in the bottom half in the *tracked* treatment (N=24). The right graph similarly plots the number of groups by teaching frequency for groups in the *untracked* treatment (N=48) again and groups consisting of subjects in the top half in the *tracked* treatment (N=24). A teaching statements is defined to be any utterance in which subjects are engaged in trying to teach each other how to do Sudoku such as "You can't have a five there; there is already one in that column.

Table C1: Effect of *Teaching* on Learning (Pooled)

Outcome: Learning

	(1)	(2)	(3)	(4)
A. Overall				
Teaching	0.137***	0.142***		
	(0.040)	(0.040)		
B. Heterogeneity				
Teaching for Top half			-0.014	-0.005
			(0.018)	(0.017)
Teaching for Bottom half			0.280***	0.276***
			(0.078)	(0.077)
Controls	No	Yes	No	Yes

Notes: Each column reports the results from a different OLS regression. Columns (1) and (2) come from equation [2] with and without controls using the full sample. Columns (3) and (4) come from equation [3] with and without controls using full sample. Here, the control group is the *no-teaching* treatment. The estimated treatment effects and their standard errors reported in the table were computed using the *lincom* command in STATA. Standard errors clustered at the group level are reported in parentheses. The outcome is learning, which is calculated by subtracting the average solving time (AST) in the Evaluation Block (T=1) from that in the Ability Block (T=0), so that higher values indicate *improvement* in solving time. Note that AST in the Evaluation Block (T=1) and in the Ability Block (T=0) is standardized by the mean and standard deviation of raw AST at T=0 before taking the difference so that standardized AST at T=0 has a mean of zero and standard deviation of 1. The bottom half consists of those subjects ranked 5–8 and the top half those subjects ranked 1–4 at T=0. All regressions—even those labeled as including "no" controls—include a dummy for the eight subjects who could not solve any Sudoku or Nonograms puzzles in T=0. The controls further include a dummy for being male, a dummy for being experienced with Sudoku, risk attitudes (0–9), and prosociality (0–5). See Table 2 for definitions of each control variable. We also add a dummy variable for Sudoku. There were 88 sessions with 704 subjects (448 from Sudoku, and 256 from Nonograms). Each session consisted of two groups (4 subjects per group), and thus there were 176 groups. Significance levels for one-sided test: *** p<0.01, ** p<0.05, * p<0.10

Table C2: Effects of *Tracking* on Learning in the *No-teaching* and *Teaching* Treatments (Pooled)

Outcome: Learning

	In the *No-teaching* Treatment		In the *Teaching* Treatment	
	(1)	(2)	(3)	(4)
A. Overall				
Tracked	0.037		-0.129**	
	(0.053)		(0.061)	
B. Heterogeneity				
Tracked for Top half		0.015		-0.036
		(0.022)		(0.030)
Tracked for Bottom half		0.058		-0.244**
		(0.107)		(0.115)
Controls	Yes	Yes	Yes	Yes

Notes: The estimated treatment effects in Columns (1) and (3) come from equation [4], while the estimated treatment effects in Columns (2) and (4) come from equation [5]. The control treatment is the *untracked* treatment. The estimated treatment effects and their standard errors were computed using the *lincom* command in STATA. Standard errors clustered at the group level are reported in parentheses. The outcome is learning, which is calculated by subtracting the average solving time (AST) in the Evaluation Block (T=1) from that in the Ability Block (T=0), so that higher values indicate *improvement* in solving time. Note that AST in the Evaluation Block (T=1) and in the Ability Block (T=0) is standardized by the mean and standard deviation of raw AST at T=0 before taking the difference so that standardized AST at T=0 has a mean of zero and standard deviation of 1. The bottom half consists of those subjects ranked 5–8 in the Ability Block (T=0). See Figure 2 for details of the procedure used to assign subjects to groups in the *untracked* and *tracked* treatments. The controls include a dummy for being male, a dummy for being experienced with Sudoku, risk attitudes (0–9), prosociality (0–5), and a dummy for the eight subjects who could not solve any Sudoku or Nonograms puzzles in the Ability Block (T=0). We also add a dummy variable for Sudoku. There are 40 *no-teaching* sessions (24 from Sudoku and 16 from Nonograms) with 320 subjects, and 48 *teaching* sessions (32 from Sudoku and 16 from Nonograms) with 384 subjects. Significance levels: *** p<0.01, ** p<0.05, * p<0.10

Table C3: Frequency of Teaching in the *Tracked* vs. *Untracked* Treatments (Pooled)

Outcome: Frequency of Teaching

	Untracked	*Tracked*		Difference (2)–(1)		Difference (3)–(1)	
		Bottom Half	Top Half	OLS	Zero-Inflated Poisson	OLS	Zero-Inflated Poisson
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	4.59	4.00	1.96	-0.59	-0.59	-2.63**	-2.63***
	[6.17]	[5.64]	[2.69]	(1.51)	(0.91)	(1.32)	(0.81)
Vuong test of Zero-Inflated model vs. Standard Poisson					z= 4.09 p= 0.000		z= 4.34 p= 0.000
# of Groups	46	24	24	70	70	70	70
# of Sessions	23	24		47	47	47	47

Notes: The unit of observation is a group. The sample is limited to the 48 *teaching* treatment sessions with 24 sessions each for the *untracked* and *tracked* treatments. In the *untracked* treatment, there are 46 groups (two groups per session) as there was one *untracked* session in which the audio recording did not work, while in the *tracked* treatment there are 24 groups each for subjects in the bottom half (Group 2 in the *tracked* treatment in Figure 2) and for those in the top half (Group 1 in the *tracked* treatment in Figure 2). Column (1) reports the mean number of teaching statements exhibited by groups in the *untracked* treatment, and Columns (2) and (3) report them for the *tracked* treatment for the bottom half group and the top half group, respectively. Standard deviations are reported in brackets. Columns (4) and (5) report the estimated difference between Columns (1) and (2) from OLS and zero-inflated Poisson (where the inflation equation includes just a dummy for *tracked* sessions) models, respectively, with standard errors in parentheses. Columns (6) and (7) report the corresponding estimated differences between Columns (1) and (3). A teaching statements is defined to be any utterance in which subjects are engaged in trying to teach each other how to do Sudoku such as "You can't have a five there; there is already one in that column." Reported below Columns (5)–(6) and (8)–(9) are z-scores and p-values for the Vuong test of the Zero-Inflated Poisson model against a standard Poisson; these tests support the use of the Zero-Inflated model. Significance levels for one-sided test: *** p<0.01, ** p<0.05, * p<0.10

Appendix D: Experiment Instructions

Below we include the instructions for the experiment. Notes for the reader (not visible to subjects) indicate page breaks. These notes are contained in [brackets] and explain which instructions were viewed by subjects in each treatment. The instructions below are for the Sudoku experiment. The instructions for the Nonogram experiment were identical with "Nonogram" appearing everywhere "Sudoku" appears below; the only exception is that the portion of the instructions explaining the rules of the game varied accordingly, as indicated below.

[Introduction – All treatments]

You are now participating in a decision-making experiment. If you follow the instructions carefully, you can earn a considerable amount of money depending on your decisions and the decisions of the other participants. Your earnings will be paid to you in CASH at the end of the experiment.

Today's experiment will involve multiple tasks, and your total payment will be the sum of your payments from each task. Please do not talk to other subjects during the experiment. This set of instructions is for your private use only. In case of questions, please raise your hand. Then we will come to you and answer your questions privately.

[Dictator Game – All treatments]

In Part I of this experiment, you are matched with one other person in this room.

Each person begins with 5 Dollars. Each person chooses how to allocate this money between him/herself and the person he/she is paired with. To specify an allocation, please type the amount you want to allocate to yourself and the amount you want to allocate to the other person and then click "Next." The two amounts must sum up to 5 Dollars.

After everyone has chosen their allocation, the computer will randomly choose one person from each pair whose decision is implemented. Your payment for Part I will be based on the randomly chosen person's decisions.

For example, if your decision is randomly chosen to be implemented, then you will be paid according to your allocation. If instead, the other persons decision is chosen to be implemented, you will be paid according to their allocation.

This is the end of the instructions for Part I. If you have any questions, please raise your hand and an experimenter will answer them privately.

[Risk Preference – All treatments]

In this task, you will make a series of choices between two uncertain options. For each decision, all you have to do is indicate whether you prefer Option A or Option B and click the appropriate button with your mouse.

At the end, the computer will randomly pick ONE decision and then randomly draw a whole number between 1 and 100 to determine your payoff. Your payoff then depends on your choice of A or B in the randomly chosen decision and the number drawn by the computer.

If you choose option A, then numbers between 1 and 50 will always pay $1 and numbers between 51 and 100 will always pay $3.

If you choose option B, then the relationship between the number drawn and the payment is different in each decision, so make sure you carefully compare options A and B to make your choices.

In the first decision, Option B will pay $0 on numbers between 1 and 90 and $3 on numbers between 91 and 100. As you proceed through the decisions, Option B will change. The chance of receiving $3 will increase, while the chance of receiving $0 will decrease.

Remember, after you make all of your choices, the computer will randomly pick one of them to count and will draw a random number to determine your payment.

This means that each decision could be the decision-that-counts so it is in your interest to treat each decision as if it could be the one that determines your payment.

We will reveal the outcome of this task at the end of todays experiment, just before you are paid.

[Risk Preference Decision – All treatments]

Choose whether you prefer Option A or Option B.

At the end of the experiment, the computer will choose one of these decisions for actual payment and will then draw a random number between 1 and 100, which together with your choices, determines how much you get paid.

Remember, this could be the decision that counts, so consider your choice carefully!

[Sudoku General – All treatments]

The next task will involve solving puzzles known as Sudokus.

	6	4	3	2	
5	3			6	4
4					1
3					2
6	4			1	3
	1	3	4	5	

A 6×6 sudoku puzzle is a grid 6 squares wide and 6 squares deep. The lines of squares running horizontally are called rows, and the lines running vertically are called columns. The grid is further divided by the darker lines into six 2×3 rectangular boxes.

Some of the squares already have numbers in them. Your task is to fill in the blank squares.

Theres only one rule: Each row, column and box must end up containing all of the numbers from 1 to 6. This rule has an important side-effect, which is the basis of all solving techniques: Each number can only appear once in a row, column or box.

[Nonograms General – All treatments]

The next task will involve solving puzzles known as Nonograms.

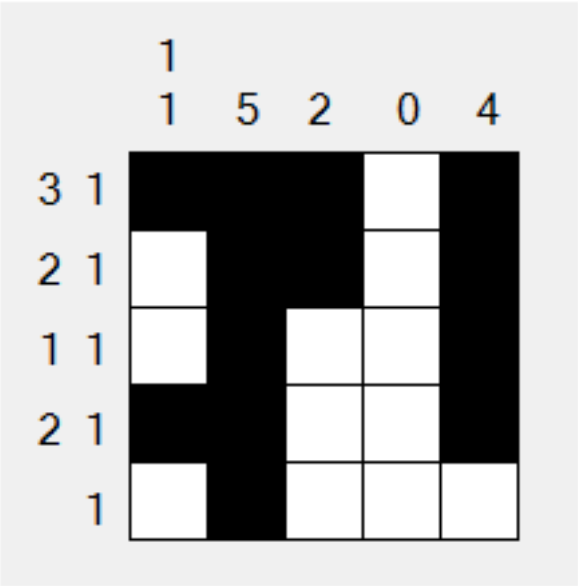

A 5×5 Nonogram puzzle is a grid 5 squares wide and 5 squares deep. The lines of squares running horizontally are called rows, and the lines running vertically are called columns.

Each square must be either filled in black or left blank (or, equivalently, marked with an X).

Beside each row of the grid are listed the lengths of the runs of black squares on that row.

For example, the top row of the grid says 3,1. This means the row must contain a run of 3 consecutive black squares and a single black square.

Above each column are listed the lengths of the runs of black squares in that column.

For example, the leftmost column says 1, 1. This means the column must contain two (and only two) single black squares that are not consecutive.

Your aim is to find all black squares. Left click on a square to make it black. Left click a second time to mark it with an X. Left click a third time to turn it white again.

The puzzle is correct when all the rows and columns have the correct number of black squares in the proper positions.

[Sudoku Video – All treatments]

Before the experiment begins, we will watch the following video. Please pay attention to the video as you may find its content useful for the next portion of the experiment for which you will be paid.

This video describes strategies for solving 9×9 Sudokus, but they apply equally to 6×6 Sudokus.

[Ability Block – All treatments]

You will now be given 20 Sudoku puzzles to solve. You will have 10:00 minutes. You will be paid $0.5 for each puzzle that you correctly complete and $0 for each puzzle you do not complete or complete incorrectly. There is a timer in the upper right hand portion of the screen for your reference. Use the mouse to select cells. Numbers may be entered by clicking the buttons on screen or with the numbers on the keyboard. If you finish before the time limit, please wait quietly until the next puzzle begins.

[Review - *tracked* and *piece-rate*]

Now we will begin task 4. You will participate in two periods where you will attempt to solve more sudoku puzzles. The first of these periods will be for practice and the second will be for payment which will be added to your earnings at the end of the experiment.

Based on your performance in Part I, you have been sorted into groups. The top half performers are in one group and the bottom half performers are in the other group. The highest performers are those who correctly completed the most puzzles, with ties broken in favor of those who completed them the fastest

In the second period, you will be paid $0.5 for each puzzle that you correctly complete and $0 for each puzzle you do not complete or complete incorrectly.

[Review - *tracked* and *tournament*]

Now we will begin task 4. You will participate in two periods where you will attempt to solve more sudoku puzzles. The first of these periods will be for practice and the second will be for payment which will be added to your earnings at the end of the experiment.

Based on your performance in Part I, you have been sorted into groups. The top half performers are in one group and the bottom half performers are in the other group. The highest performers are those who correctly completed the most puzzles, with ties broken in favor of those who completed them the fastest

In the second period, your payment will be determined by your RELATIVE performance in your group. In other words, at the end of the second period we will count the number of puzzles completed correctly by each person, and pay will be based on your rank.

The participant in rank 1 will be paid $20
The participant in rank 2 will be paid $10
The participant in rank 3 will be paid $5
The participant in rank 4 will be paid $0

Ties will be broken in favor of those who completed the puzzles the fastest

[Review - *untracked* and *piece-rate*]

Now we will begin task 4. You will participate in two periods where you will attempt to solve more sudoku puzzles. The first of these periods will be for practice and the second will be for payment which will be added to your earnings at the end of the experiment.

Based on your performance in Part I, you have been sorted into groups. These groups have been balanced so that they contain both high and low performers. The highest performers are those who correctly completed the most puzzles, with ties broken in favor of those who completed them the fastest.

In the second period, you will be paid $0.5 for each puzzle that you correctly complete and $0 for each puzzle you do not complete or complete incorrectly.

[Review - *untracked* and *tournament*]

Now we will begin task 4. You will participate in two periods where you will attempt to solve more sudoku puzzles. The first of these periods will be for practice and the second will be for payment which will be added to your earnings at the end of the experiment.

Based on your performance in Part I, you have been sorted into groups. These groups have been balanced so that they contain both high and low performers. The highest performers are those who correctly completed the most puzzles, with ties broken in favor of those who completed them the fastest.

In the second period, your payment will be determined by your RELATIVE performance in your group. In other words, at the end of the second period we will count the number of puzzles completed correctly by each person, and pay will be based on your rank.

The participant in rank 1 will be paid $20
The participant in rank 2 will be paid $10
The participant in rank 3 will be paid $5
The participant in rank 4 will be paid $0

Ties will be broken in favor of those who completed the puzzles the fastest.

[Practice Block - *no-teaching*]

You will now be given a single Sudoku puzzle to solve. You will not be paid for this puzzle. You will have 10:00 minutes to solve the puzzle.

[Practice Block - *teaching*]

You will now be given a single Sudoku puzzle to solve. You will not be paid for this puzzle. You will have 10:00 minutes to solve the puzzle.

You can complete this puzzle working with the people in your group. During this period, your microphone will be enabled and a voice chat room will be available in which you can discuss the puzzle you are working on. You may discuss any aspects of the experiment in the chat room, but you may not reveal your identity, make threats, or use inappropriate language (including shorthand like WTF). Other

participants will be identified by a number next to their mouse cursor. This is their rank within the group. Please only speak English.

[Evaluation Block - *piece-rate*]

You will now be given 30 Sudoku puzzles to solve. You will have 15:00 minutes.

Remember, you will be paid $0.5 for each puzzle that you correctly complete, and you will be paid $0 for each puzzle you do not complete or complete incorrectly. At the end of the 15:00 minutes, you will learn the number of puzzles you correctly completed.

You will also learn your rank among the other subjects in terms of the number of correctly completed items in the previous 15:00 minutes.

[Evaluation Block - *tournament*]

You will now be given 30 Sudoku puzzles to solve. You will have 15:00 minutes.

At the end of this task, you will learn your rank among the other subjects in terms of the number of correctly completed items. Remember, in this task, your payment will be determined by your RELATIVE performance in your group. In other words, at the end of this task we will count the number of puzzles completed correctly by each person, and pay will be based on your rank

The participant in rank 1 will be paid $20
The participant in rank 2 will be paid $10
The participant in rank 3 will be paid $5
The participant in rank 4 will be paid $0

Ties will be broken in favor of those who completed the puzzles the fastest

[At the conclusion, subjects observe their payment from each task and are asked to wait quietly to be called over for payment]

Appendix E: Screenshots

Screenshots from each stage of the Experiment. Nonograms screenshots included only where they differ from Sudoku.

Page 1 – Demographics

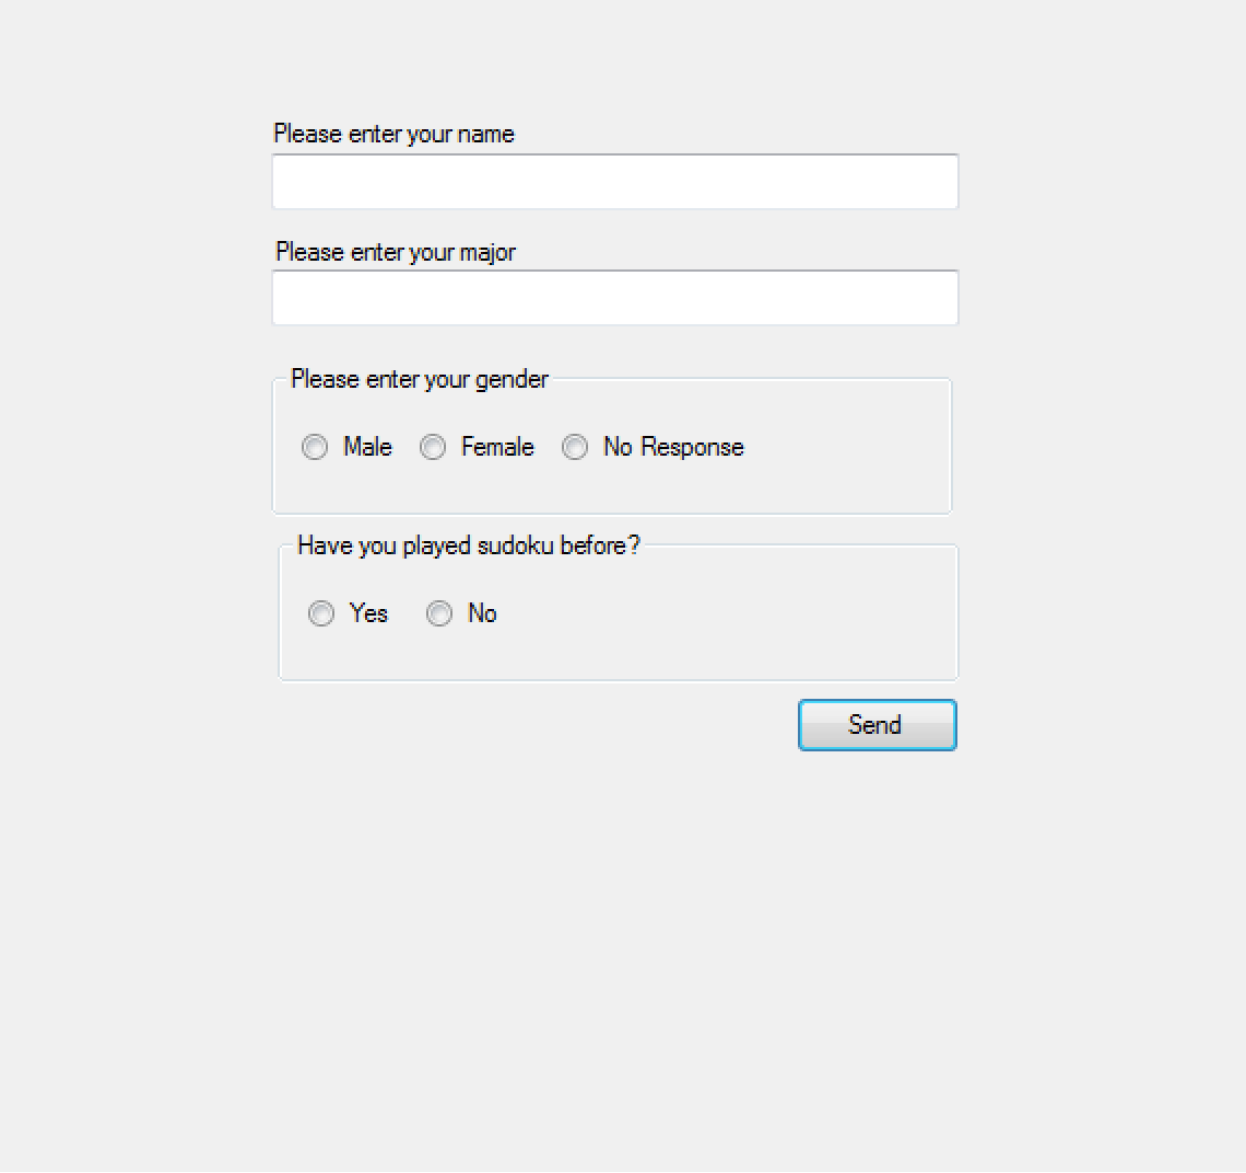

Page 2 – General Instructions

You are now participating in a decision-making experiment. If you follow the instructions carefully, you can earn a considerable amount of money depending on your decisions and the decisions of the other participants. Your earnings will be paid to you in CASH at the end of the experiment.

Today's experiment will involve multiple tasks, and your total payment will be the sum of your payments from each task. Please do not talk to other subjects during the experiment. This set of instructions is for your private use only. In case of questions, please raise your hand. Then we will come to you and answer your questions privately.

Next

Page 3 – Dictator Allocation

In Part I of this experiment, you are matched with one other person in this room.

Each person begins with 5 Dollars. Each person chooses how to allocate this money between him/herself and the person he/she is paired with. To specify an allocation, please type the amount you want to allocate to yourself and the amount you want to allocate to the other person and then click "Next." The two amounts must sum up to 5 Dollars.

After everyone has chosen their allocation, the computer will randomly choose one person from each pair whose decision is implemented. Your payment for Part I will be based on the randomly chosen person's decisions.

For example, if your decision is randomly chosen to be implemented, then you will be paid according to your allocation. If instead, the other person's decision is chosen to be implemented, you will be paid according to their allocation.

This is the end of the instructions for Part I. If you have any questions, please raise your hand and an experimenter will answer them privately.

You Receive	They Receive
$	$

Next

Page 4 – Risk Elicitation Outline

In this task, you will make a series of choices between two uncertain options. For each decision, all you have to do is indicate whether you prefer Option A or Option B and click the appropriate button with your mouse.

At the end, the computer will randomly pick ONE decision and then randomly draw a whole number between 1 and 100 to determine your payoff. Your payoff then depends on your choice of A or B in the randomly chosen decision and the number drawn by the computer.

If you choose option A, then numbers between 1 and 50 will always pay $1 and numbers between 51 and 100 will always pay $3.

If you choose option B, then the relationship between the number drawn and the payment is different in each decision, so make sure you carefully compare options A and B to make your choices.

In the first decision, Option B will pay $0 on numbers between 1 and 90 and $3 on numbers between 91 and 100. As you proceed through the decisions, Option B will change. The chance of receiving $3 will increase, while the chance of receiving $0 will decrease.

Remember, after you make all of your choices, the computer will randomly pick one of them to count and will draw a random number to determine your payment.

This means that each decision could be the decision-that-counts so it is in your interest to treat each decision as if it could be the one that determines your payment.

We will reveal the outcome of this task at the end of today's experiment, just before you are paid.

Next

Page 5 – Risk Elicitation Decision Screen (9 choices)

Choose whether you prefer Option A or Option B.

At the end of the experiment, the computer will choose one of these decisions for actual payment and will then draw a random number between 1 and 100, which together with your choices, determines how much you get paid.

Remember, this could be the decision that counts, so consider your choice carefully!

Option A	Option B
Numbers between 1 and 50 pay $1.	Numbers between 1 and 90 pay $0.
Numbers between 51 and 100 pay $3.	Numbers between 91 and 100 pay $3.
Select Option A	Select Option B

Next Decision

Current Decision: 1 / 9

Page 6a – Basic Sudoku Instructions

The next task will involve solving puzzles known as Sudokus.

	6	4	3	2	
5	3			6	4
4					1
3					2
6	4			1	3
	1	3	4	5	

A 6x6 sudoku puzzle is a grid 6 squares wide and 6 squares deep. The lines of squares running horizontally are called rows, and the lines running vertically are called columns. The grid is further divided by the darker lines into six 2 X 3 rectangluar 'boxes'.

Some of the squares already have numbers in them. Your task is to fill in the blank squares. There's only one rule:

- Each row, column and box must end up containing all of the numbers from 1 to 6.

This rule has an important side-effect, which is the basis of all solving techniques:

- Each number can only appear once in a row, column or box.

Next

Page 6b – Basic Nonograms Instructions

The next task will involve solving puzzles known as Nonograms.

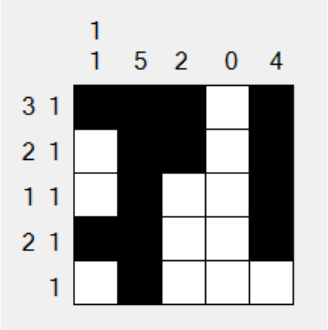

A 5x5 Nonogram puzzle is a grid 5 squares wide and 5 squares deep. The lines of squares running horizontally are called rows, and the lines running vertically are called columns.

Each square must be either filled in black or left blank (or, equivalently, marked with an X).

Beside each row of the grid are listed the lengths of the runs of black squares on that row.

- For example, the top row of the grid says 3,1. This means the row must contain a run of 3 consecutive black squares and a single black square.

Above each column are listed the lengths of the runs of black squares in that column.

- For example, the leftmost column says 1, 1. This means the column must contain two (and only two) single black squares that are not consecutive.

Your aim is to find all black squares. Left click on a square to make it black. Left click a second time to mark it with an X. Left click a third time to turn it white again.

The puzzle is correct when all the rows and columns have the correct number of black squares in the proper positions.

Next

Page 7a – Sudoku Instructional Video

Before the experiment begins, we will watch the following video. Please pay attention to the video as you may find its content useful for the next portion of the experiment for which you will be paid .

This video describes strategies for solving 9x9 sudokus, but they apply equally to 6x6 sudokus.

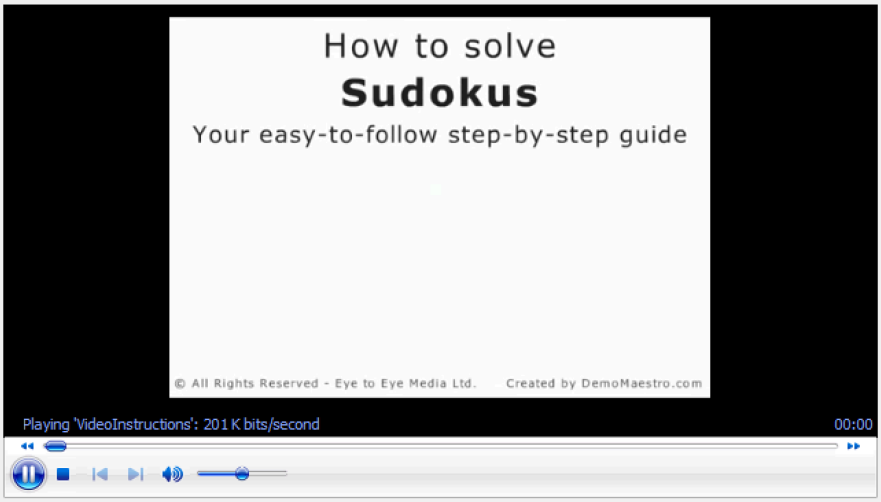

Please Watch The Video

Page 7b – Nonogram Instructional Video

Before the experiment begins, we will watch the following video. Please pay attention to the video as you may find its content useful for the next portion of the experiment for which you will be paid .

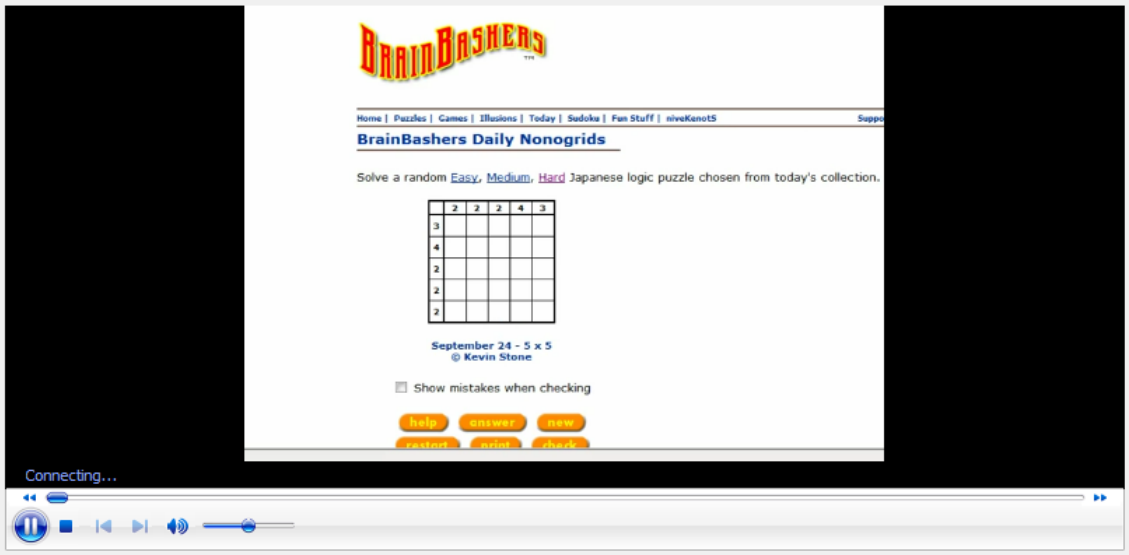

Please Watch The Video

Page 8 – Ability Block Instructions

You will now be given 20 Sudoku puzzles to solve. You will have 10:00 minutes. You will be paid $0.5 for each puzzle that you correctly complete and $0 for each puzzle you do not complete or complete incorrectly. There is a timer in the upper right hand portion of the screen for your reference. Use the mouse to select cells. Numbers may be entered by clicking the buttons on screen or with the numbers on the keyboard. If you finish before the time limit, please wait quietly until the next puzzle begins.

Click when ready to continue

Page 9 – Ability Block Screen

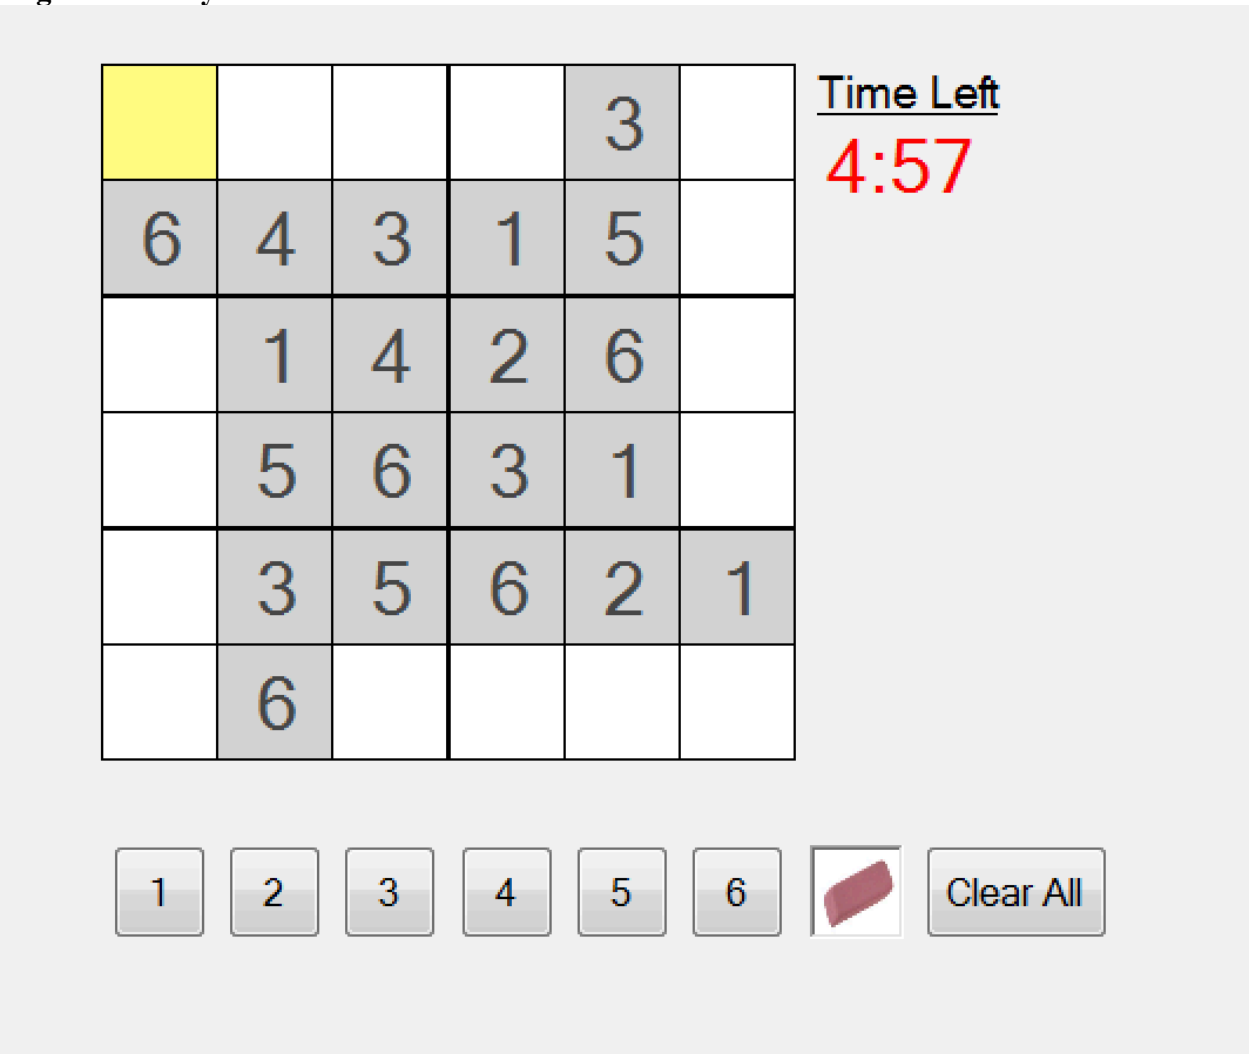

Page 10 – Ability Block Results and Tracking Information

Your Results		Your Group's Results		
Number of correct sudokus:	2	First place solved:	2	(You)
Your rank:	1	Second place solved:	1	
You have earned: $	1	Third place solved:	1	
		Fourth place solved:	1	

Now we will begin task 4. You will participate in two periods where you will attempt to solve more sudoku puzzles. The first of these periods will be for practice and the second will be for payment which will be added to your earnings at the end of the experiment.

Based on your performance in Part I, you have been sorted into groups. The top half performers are in one group and the bottom half performers are in the other group. The highest performers are those who correctly completed the most puzzles, with ties broken in favor of those who completed them the fastest

In the second period, you will be paid $0.5 for each puzzle that you correctly complete and $0 for each puzzle you do not complete or complete incorrectly.

Click when ready to continue

Page 11 – Practice Block Instructions (Chat Treatments)

You will now be given a single Sudoku puzzle to solve. You will not be paid for this puzzle. You will have 10:00 minutes to solve the puzzle.

You can complete this puzzle working with the people in your group. During this period, your microphone will be enabled and a voice chat room will be available in which you can discuss the puzzle you are working on. You may discuss any aspects of the experiment in the chat room, but you may not reveal your identity, make threats, or use inappropriate language (including shorthand like WTF). Other participants will be identified by a number next to their mouse cursor. This is their rank within the group. Please only speak english.

Click when ready to continue

Page 12 – Practice Block Screen (Chat Treatments)

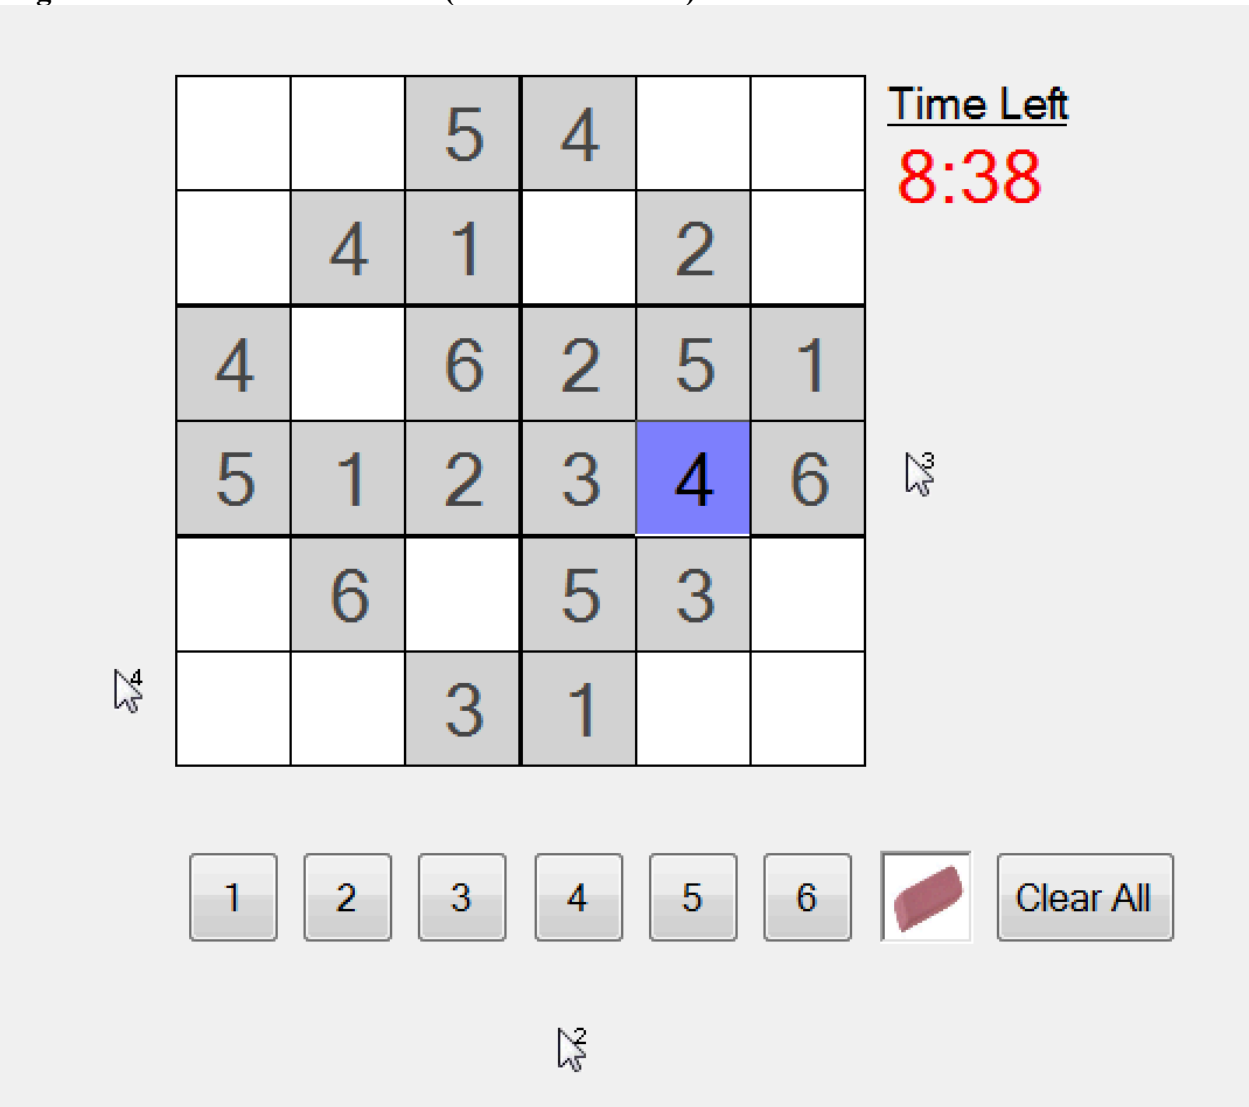

Page 13 – Practice Block Results

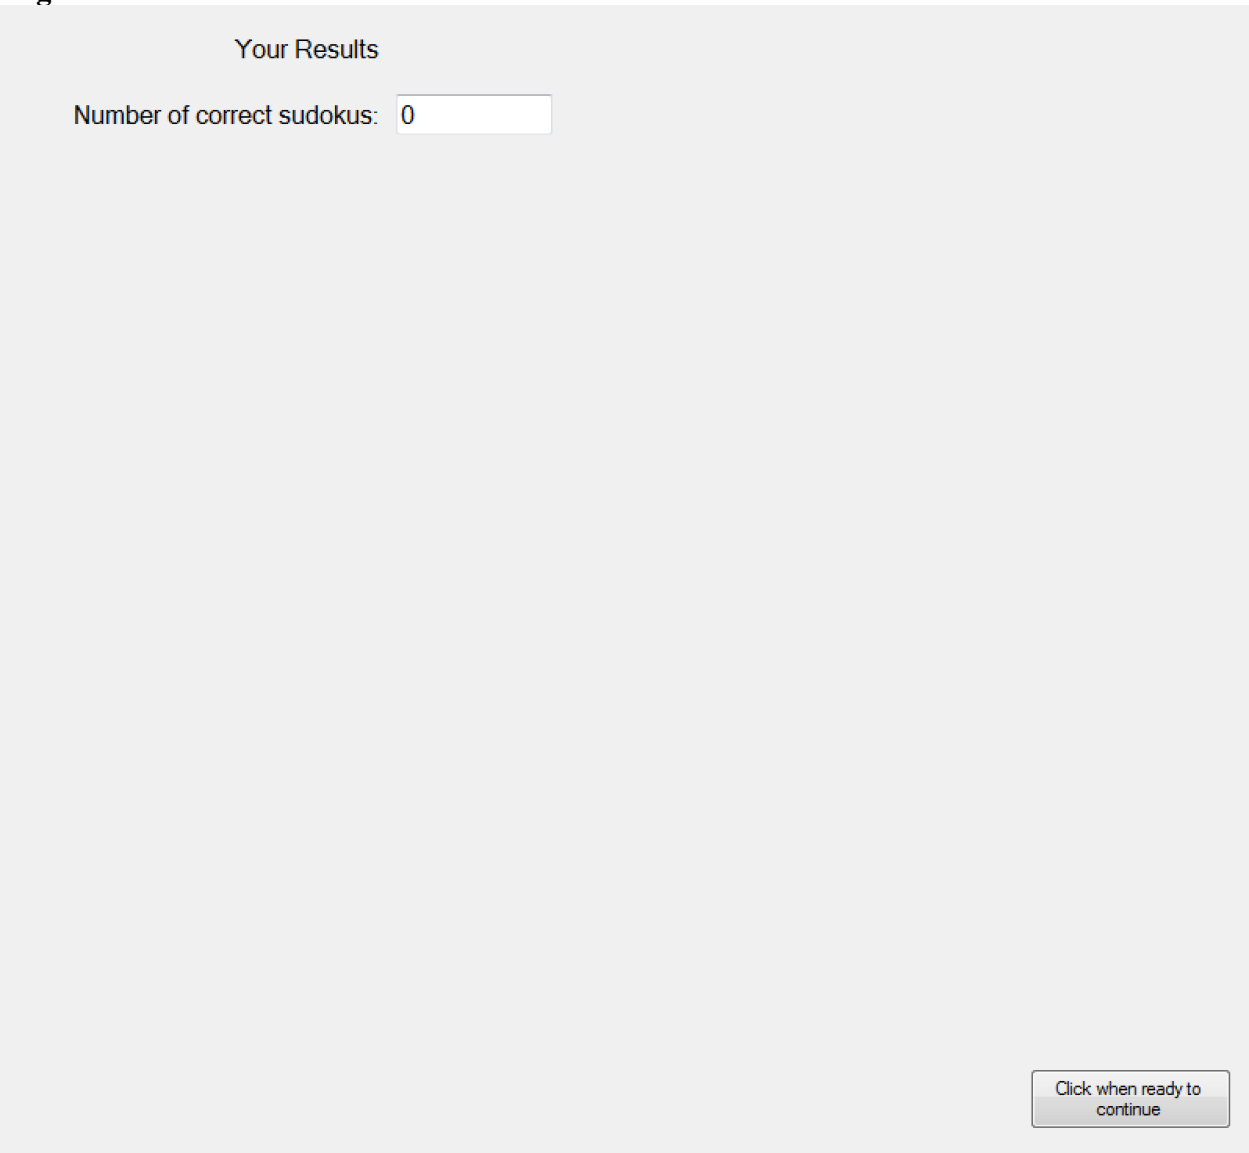

Page 14 – Evaluation Block Instructions (Piece Rate Treatment)

You will now be given 30 Sudoku puzzles to solve. You will have 15:00 minutes.

Remember, you will be paid $0.5 for each puzzle that you correctly complete, and you will be paid $0 for each puzzle you do not complete or complete incorrectly. At the end of the 15:00 minutes, you will learn the number of puzzles you correctly completed.

You will also learn your rank among the other subjects in terms of the number of correctly completed items in the previous 15:00 minutes.

Click when ready to continue

Page 15 – Payment Summary Screen

Thank you for participating in this research study

In task 1, your choice was selected	From task 1 you have earned $: 5
In task 2, decision number 1 was randomly selected. You selected choice A and recieved a roll of: 33	From task 2 you have earned $: 1
In task 3 you solved 2 sudokus correctly	From task 3 you have earned $: 1
For task 4, you solved 1 sudokus correctly You are ranked number 1 out of 4 participants	From task 4 you have earned $: 0.5

Total Earnings $: 7.5

Please Wait Quietly

Appendix F: Methods Used to Count Instances of Teaching in Audio Data

Two research assistants were given audio files for each group in each session of the *teaching* treatments. We explained that subjects were working together on a single Sudoku puzzle and that the chat transcript provided a record of this interaction. We instructed them to come up with a count of the number of instances of peer-to-peer teaching in the text. To define teaching, they adapted a scheme developed by Kline (2016) that identifies teaching with behavior related to the relevant task in which the following events occur:

- **+/- verbal feedback**- Actor gives positive or negative verbal appraisal of F's behavior with respect to, in F's presence.
- **+/- consequences**- Actor creates positive or negative consequences for F. Includes physical punishment or reward, or verbal description of promised punishment or reward.
- **teasing**- Actor threatens or punishes but in a joking manner, as indicated by smiling and/or laughing. Verbal or gestural.
- **warning (of danger)**- Verbal warning of danger, a separately coded subset of negative verbal appraisal.
- **command to stop**- Verbal command to focal to stop, a separately coded subset of negative verbal appraisal.
- **command to say/do [x]**- Verbal command to say a given phrase, or repeat a particular phrase or gesture, a separately coded subset of positive verbal appraisal.
- **direct attention to object-** Actor directs F's attention toward an object or location, verbally or through gesture
- **direct attention to person-** Actor directs F's attention toward another person, verbally or through gesture.
- **command to watch**- Actor directs F's attention toward watching the actor.
- **other-prompted behavior-** F undertaking an action or behavior after being commanded to do so by another actor.
- **other-assisted behavior**- F undertaking an action or behavior made possible by another actor's help.
- **abstract communication-** Actor gives a verbal explanation or states abstract information to F – including the statement of rules and what is "taboo" behavior.
- **demonstration**- Actor performs behavior or action conditioned on F's attention; may follow a request from F, or the Actor may first manipulate F's attention.

After coding each session independently, our research assistants met to resolve any inconsistencies and produced a count of teaching instances for each group in each session.

A few other notes:

1) A single instance of teaching would include both a question asked by a subject and the answer received.

2) If no answer was received to a question, this was *not* counted as an instance of teaching.
3) We had hoped to get individual level data on teaching, but unfortunately it was not consistently possible to identify specific individuals by their voices and match them to the data.

References:

Kline, M.A. 2016. "TEACH: An Ethogram-based Method to Observe and Record Teaching Behavior." *Field Methods.* doi:10.1177/1525822X16669282