Online Appendices B–G for

M Equilibrium: A Theory of Beliefs and Choices in Games

Jacob K. Goeree and Philippos Louis

Appendix B. Variational Inequality Characterization of M Equilibrium

Here we present a characterization based on variational inequalities akin to those characterizing Nash equilibrium that is well-suited for computation. Recall that $\sigma \in \Sigma$ is a Nash equilibrium if $\pi_i(\sigma_i, \sigma_{-i}) \geq \pi_i(s_i, \sigma_{-i})$ for all $s_i \in S_i$, $i \in N$. We will rewrite these inequalities as variational inequalities. For $i \in N$, let ρ_i denote a permutation of the set $\{1, 2, \ldots, K_i\}$ and for $v \in \mathbb{R}^{K_i}$ let $\rho_i(v) = (v_{\rho_i(1)}, \ldots, v_{\rho_i(K_i)})$. For $v, w \in \mathbb{R}^{K_i}$ let $\langle v|w\rangle = \sum_{k=1}^{K_i} v_k w_k$ denote the usual inner product. The inequalities defining Nash equilibrium can be written as

$$\langle \pi_i(\sigma_{-i}) \,|\, \sigma_i - \rho_i(\hat{\mu}_i)\rangle \;\geq\; 0$$

for all permutations ρ_i and $i \in N$. (Here $\hat{\mu}_i = (1, 0, \ldots, 0)$ so the set of all permutations of $\hat{\mu}_i$ is S_i.) These variational inequalities have an intuitive geometric interpretation. For instance, if σ_i is a pure strategy corresponding to a vertex of Σ_i then the differences $\sigma_i - \rho_i(\hat{\mu}_i)$ define the normal cone at σ_i, see the left panel of Figure 1. The above inequality requires $\pi_i(\sigma_{-i})$ to lie in this normal cone. And if σ_i is totally mixed, the normal cone is one-dimensional and generated by $e = (1, 1, \ldots, 1)$. The requirement that $\pi_i(\sigma_{-i})$ lies in this normal cone means that all expected payoffs are equal.

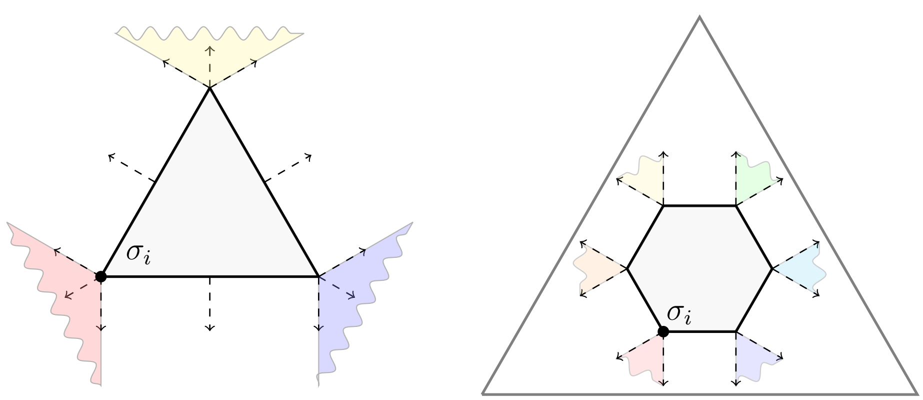

Figure 1: The left panel shows a probability simplex and its normal fan, i.e. the collection of normal cones. The right panel shows a permutahedron inside the simplex and its normal fan.

Now consider, for $i \in N$, the variational inequalities

$$\begin{aligned} \langle \pi_i(\sigma_{-i}) \,|\, \sigma_i - \rho_i(\sigma_i) \rangle &\geq \ 0 \\ \langle \pi_i(\omega_i) \,|\, \sigma_i - \rho_i(\sigma_i) \rangle &\geq \ 0 \end{aligned} \tag{B.1}$$

The interpretation of the top inequality is that σ is a Nash equilibrium when, for $i \in N$, strategies are restricted to the permutahedron generated by σ_i, see the right panel of Figure 1. This permutahedron has fewer than $K_i!$ vertices when $\rho_i(\sigma_i) = \sigma_i$, i.e. when σ_i belongs to one or more simplex diagonals. The contrapositives of the conditions in Definition 1 then imply

$$\begin{aligned} \rho_i(\sigma_i) \ = \ \sigma_i \Rightarrow \ \rho_i(\pi_i(\sigma_{-i})) \ &= \ \pi_i(\sigma_{-i}) \\ \rho_i(\sigma_i) \ = \ \sigma_i \Rightarrow \ \rho_i(\pi_i(\omega_i)) \ &= \ \pi_i(\omega_i) \end{aligned} \tag{B.2}$$

Proposition 1 *For any normal-form game G:*

$$\bigcup_{M \in \mathcal{M}(G)} M^c \times M^b \ = \ \{(\sigma, \omega) \in \Sigma_{int} \times \ \Omega \,|\, \text{(B.1) and (B.2) hold } \forall \rho_i,\ i \in N\}$$

Proof. For the choices these conditions are equivalent to $(\sigma_{ik} - \sigma_{ij})(\pi_{ik}(\sigma) - \pi_{ij}(\sigma)) \geq 0$ and $\sigma_{ik} = \sigma_{ij} \Rightarrow \pi_{ik}(\sigma) = \pi_{ij}(\sigma)$ for $i \in N$ and $1 \leq j, k \leq K_i$. These imply $\pi_{ik}(\sigma) > \pi_{ij}(\sigma) \Rightarrow \sigma_{ik} > \sigma_{ij}$ and $\sigma_{ik} = \sigma_{ij} \Rightarrow \pi_{ik}(\sigma) = \pi_{ij}(\sigma)$. This is Definition 1 restricted to choices. Conversely, Definition 1 implies $(\sigma_{ik} - \sigma_{ij})(\pi_{ik}(\sigma) - \pi_{ij}(\sigma)) \geq 0$ and $\pi_{ij}(\sigma) = \pi_{ik}(\sigma)$ if $\sigma_{ij} = \sigma_{ik}$. This is Proposition 1 restricted to choices. The proof for the beliefs, ω, is similar. ■

This proposition determines all (σ, ω) pairs that belong to some $M \in \mathcal{M}(G)$. This union can be broken down in individual M equilibria by coloring them with the expected payoffs associated with σ or ω.

Appendix C. Belief Elicitation

In the method we use to elicit beliefs, subjects submit the "percentage chance" with which they believe their opponent chooses each action by moving a single slider between endpoints labeled A and B. Any point on the slider corresponds to a unique chance of A being played (and B with complementary chance) with the A endpoint (B endpoint) corresponding to the belief that the opponent chooses A (B) for sure. The point chosen by the subject is then compared with two computer-generated random points on the slider. If the chosen point is closer to the opponent's actual choice (one of the endpoints) than at least one of the two randomly-drawn points then the subject receives a fixed prize. The next proposition generalizes this method so that it can be used for general normal-form games.

Proposition 2 *Consider the elicitation mechanism where player $i \in N$ uses $S_i = \sum_{j \neq i} K_j$ sliders, labeled S_{jk} for $j \neq i$ and $k = 1, \ldots, K_j$, to report her beliefs, q_{jk}, about player j's choice. For each slider, two (uniform) random numbers are drawn and player i's belief for that slider is "correct" if the reported belief is closer to the actual outcome (0 or 1) than at least one of the random draws. If players are risk-neutral then the elicitation mechanism is incentive compatible when a prize, $P \geq 0$, is paid for all correct beliefs for any randomly selected subset of sliders. If players are not risk neutral then the elicitation mechanism is incentive compatible if a prize is paid when the stated belief is correct for a single randomly selected slider.*

Proof of Prop 2. Let $u_i(x)$ denote player i's utility of being paid a prize amount x (with $u_i(0) = 0$) and let p and q denote the concatenations of player i's true and reported beliefs respectively. Player i wins a prize P for slider S_{jk} with chance

$$P_{jk} \;=\; p_{jk}(1 - (1 - q_{jk})^2) + (1 - p_{jk})(1 - q_{jk}^2)$$

and gets 0 with complementary probability. If all correct beliefs for a random subset $S \subseteq S_i$ of sliders pay a prize P then player i's expected utility of reporting q when her true beliefs are p is

$$U_i(p, q) \;=\; \sum_{W \subseteq S} u_i(P|W|)\binom{|S|}{|W|} \prod_{S_{jk} \in W} P_{jk} \prod_{S_{jk} \notin W} 1 - P_{jk} \tag{C.1}$$

where W is the subset of selected sliders for which player i wins the prize P. If player i is risk neutral, i.e. $u_i(x) = x$, then this reduces to the expected number of wins

$$U_i(p, q) \;=\; \frac{P\,|S|}{|S_i|} \sum_{S_{jk} \in S_i} P_{jk}$$

and optimizing with respect to q yields

$$\nabla_q U_i(p, q) \;=\; 2\,\frac{P\,|S|}{|S_i|}\,(p - q)$$

so truthful reporting is optimal. If player i is not risk neutral then there can only be two possible payoff outcomes, 0 and P, for the elicitation mechanism to be incentive compatible. In other words, $|S| = 1$ and (C.1) reduces to

$$U_i(p, q) \;=\; \frac{u_i(P)}{|S_i|} \sum_{S_{jk} \in S_i} P_{jk}$$

and truthful reporting is again optimal. ■

Appendix D. Empirical Support

Support for Result 1.

(i) As can be seen in the top left graph of Figure 9, there is very little overlap between the confidence intervals for the average choice frequencies in the five games. Formal statistical testing confirms this. The left panel of Table 1 reports the p-values from a Hotelling's t-test in pairwise comparisons of the average choice frequency in different games. In most cases, pairwise comparisons show differences that are significant at the 5% level. The null hypothesis that average choice frequencies in all games are jointly equal can be rejected at the 1% level.[1] This remains true even when excluding game 4 from the test, as it is the single most different game.

(ii) Again, one can see in the top right graph of Figure 9 that there is little overlap of the confidence intervals for average beliefs for the five games. The results of Hotelling's t-tests for the comparison of means confirms what is observed in the graph. The right panel of table 1 reports the p-values from the test. Average beliefs are statistically different across games for any reasonable significance level.

(iii) A careful inspection of the two graphs in Figure 9 reveals that for all five games, average choice frequencies lie far from the corresponding average reported beliefs. Again, using Hotelling's t-test we confirm this finding. The p-values from comparing actions to beliefs in each of the games are very close to zero in all five cases.

(iv) Figure 2 shows the empirical cumulative density functions (CDFs) for the average choice by Row and Column in each game. If actions were homogeneous, then these should look like the cdf for the observed success rate in a binomial distribution with 8 trials and a true success rate in each trial equal to the average choice across all individuals. We use a Monte Carlo simulation with 1000 repetitions to produce this cdf and compare the two. Using a Kolmogorov-Smirnov test we can reject that the two are the same for all cases at the 1% level for all games except game 2. For game 2 we can reject at the 5% level for Row (p-value is 0.046) but not for Column (p-value is 0.076). This result is further supported by a Cochran's Q test comparing individual choices in each game. The test rejects those being the same for both Row and Column in all games for any level of significance.

[1] For the joint test we use the Bonferroni correction. The null of simultaneous equality is rejected at a level of significance α by comparing the lowest p-value in all m pairwise comparisons to $\frac{\alpha}{m}$ and rejecting whenever it is lower. This ensures that the family-wise error rate (FWER), which is the probability that at least one of the averages is not equal to the others, remains at the desired level α. This is the most conservative test based on the FWER, in the sense that it rejects the null less often.

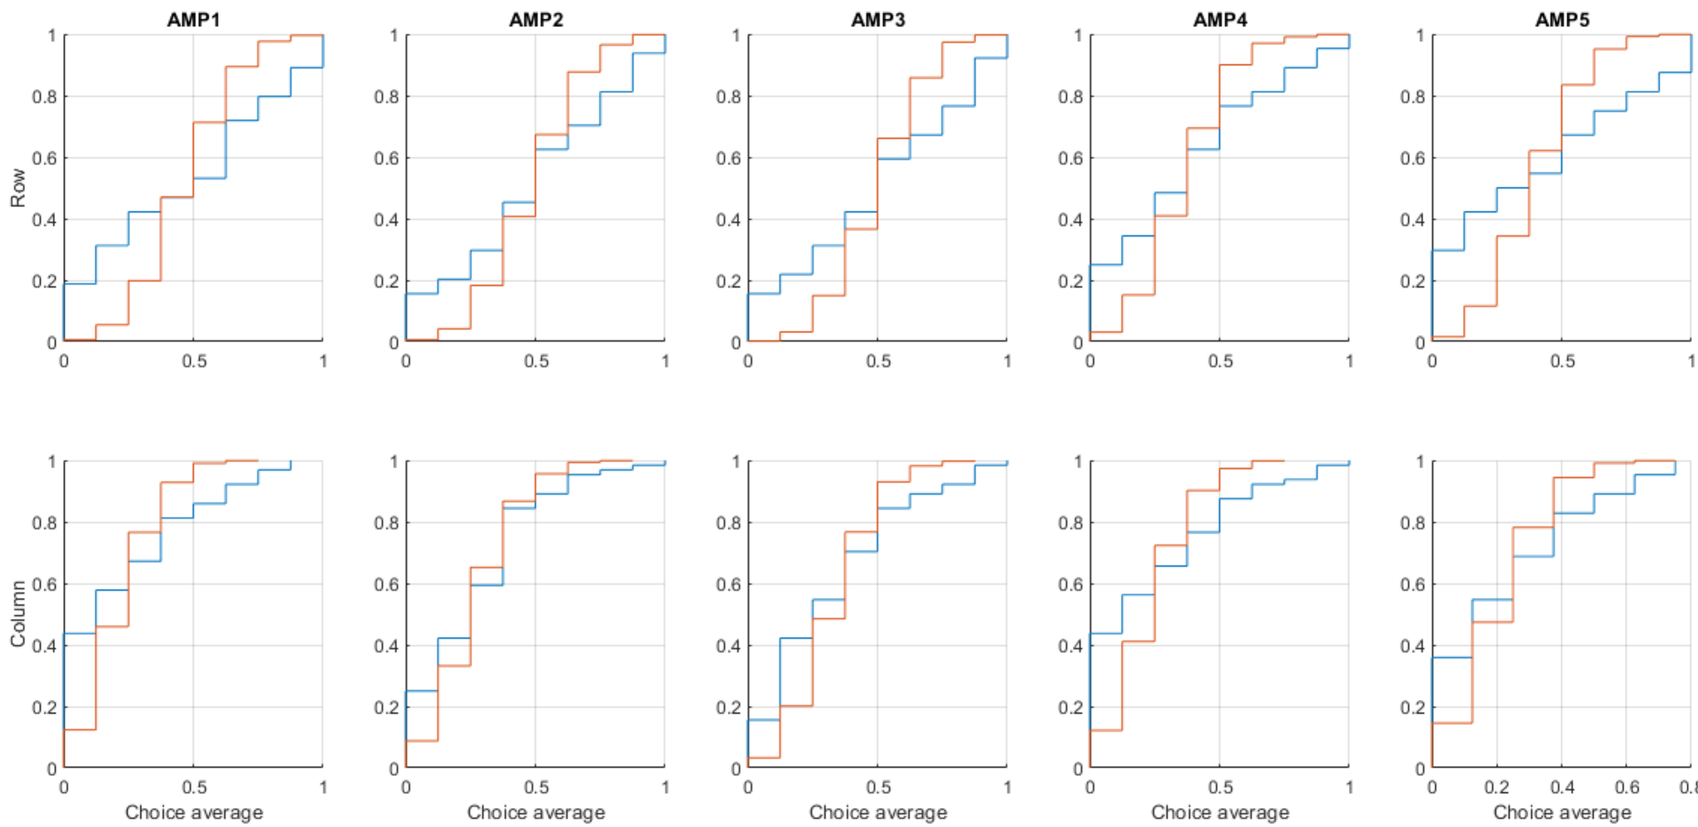

Figure 2: Heterogeneity of actions in Experiment 1. Empirical CFDs for subjects' actions (blue) and CDFs of the observed success rate in MC simulations of binomial distributions with 8 trials and true success rate equal to subjects' overall average choice (orange).

actions	2	3	4	5
1	.2344	.0011	.0045	.0863
2		.1374	.0004	.0052
3			.0000	.0000
4				.0491

beliefs	2	3	4	5
1	.0000	.0000	.0045	.6374
2		.0489	.0974	.0000
3			.0001	.0000
4				.0002

Table 1: Experiment 1: p-values from pairwise comparisons between games.

(v) We use Friedman's test to compare reported beliefs across individuals in each game, separately for Row and Column. In all cases we find there is heterogeneity that is significant at any reasonable significance level.

Appendix E. Cluster Analysis With the k-Means Algorithm

As mentioned in the main text, we classify our data of elicited beliefs into different clusters. This is done in a theory-free way using the k-means algorithm (MacQueen, 1967). Given k, the number of clusters one wants to use, k random 'centroids' are chosen, which are points in the same space as that of the data. In the next step, each observation is matched to its closest centroid (in terms of Euclidean distance). Next, new centroids are calculated by taking the mean of the observations in each cluster. The algorithm proceeds iteratively until it converges to a stable set of clusters.

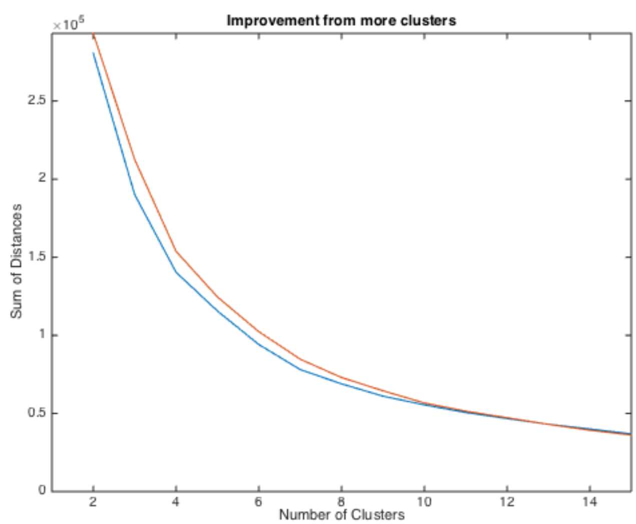

Figure 3: K means in the DS games. The curves show the sum of distances for different values of k for $DS1$ (blue line) and $DS2$ (red line).

While convergence is guaranteed, the resulting clusters will depend on the random initialisation of the algorithm. To address the issue we take the standard approach which is to repeat the analysis for a large number of times (5000 in our case) and choose the outcome with the smallest sum of errors (distance of each observation from its cluster centroid). The final issue to address is the number of clusters, k. We approach the problem using what is dubbed the 'elbow method'. We run the analysis for $k \in \{2, ..., 15\}$ and calculate the sum of errors in each case. This is then plotted against k, giving a convex, decreasing curve. For each game we choose the k that is at the 'elbow' of the curve. Figure 3 shows the curves for the case of the two DS games. In the analysis we use $k = 7$ for these games.

Appendix F. Derivation of μ Equilibria

In this appendix, we provide explicit formulae for the μ equilibria of the 3×3 games in Table 3 using $\mu_i(\varepsilon) = (1, \varepsilon, \varepsilon^2, \ldots, \varepsilon^{K_i})$ (appropriately normalized) for $\varepsilon \in [0, 1]$ and $i \in N$.

For the DS1 and DS2 games, the μ equilibrium is given by

$$\sigma(\varepsilon) = \frac{1}{(1+\varepsilon+\varepsilon^2)} \begin{cases} (\varepsilon, 1, \varepsilon^2) & \text{if} \quad 0.146 \le \varepsilon \le 1 \\ (\frac{1}{8}(1+\varepsilon+\varepsilon^2), 1, \frac{1}{8}(-1+7\varepsilon+7\varepsilon^2) & \text{if} \quad 0.146 \le \varepsilon \le 0.456 \\ (\varepsilon^2, 1, \varepsilon) & \text{if} \quad 0.333 \le \varepsilon \le 0.456 \\ (\varepsilon^2, 2+2\varepsilon-15\varepsilon^2, -1-\varepsilon+15\varepsilon^2) & \text{if} \quad 0.333 \le \varepsilon \le 0.4 \\ (\varepsilon^2, \varepsilon, 1) & \text{if} \quad 0.354 \le \varepsilon \le 0.4 \\ (\frac{1}{8}, -\frac{1}{8}+\varepsilon+\varepsilon^2, 1) & \text{if} \quad 0.125 \le \varepsilon \le 0.354 \\ (\varepsilon, \varepsilon^2, 1) & \text{if} \quad 0 \le \varepsilon \le 0.125 \end{cases}$$

see the colored curve in the upper-left panel of Figure 13, which also shows the greyed out M sets. The bottom-left panel shows the logit-QRE curve in black. The latter is unique for every value of λ. Note, however, that multiple μ equilibria may exist for certain ranges of ε.

For the "no logit" game, the μ equilibrium is given by

$$\sigma(\varepsilon) = \frac{1}{(1+\varepsilon+\varepsilon^2)} \begin{cases} (\varepsilon, 1, \varepsilon^2) & \text{if} \quad 0.053 \le \varepsilon \le 0.947 \\ (\frac{1}{40}+\frac{1}{2}\varepsilon+\frac{1}{2}\varepsilon^2, 1, -\frac{1}{40}+\frac{1}{2}\varepsilon+\frac{1}{2}\varepsilon^2) & \text{if} \quad 0.053 \le \varepsilon \le 0.947 \\ (\varepsilon^2, 1, \varepsilon) & \text{if} \quad 0.252 \le \varepsilon \le 1 \\ (\varepsilon^2, \frac{1}{3}(10+10\varepsilon-150\varepsilon^2), \frac{1}{3}(-7-7\varepsilon+150\varepsilon^2)) & \text{if} \quad 0.252 \le \varepsilon \le 0.283 \\ (\varepsilon^2, \varepsilon, 1) & \text{if} \quad 0.178 \le \varepsilon \le 0.283 \\ (\frac{1}{45}(1+2\varepsilon+2\varepsilon^2), \frac{1}{45}(-1+43\varepsilon+43\varepsilon^2), 1) & \text{if} \quad 0.024 \le \varepsilon \le 0.178 \\ (\varepsilon, \varepsilon^2, 1) & \text{if} \quad 0 \le \varepsilon \le 0.024 \end{cases}$$

see the colored curve in the upper-middle panel of Figure 13. The colored curve does enter the lower M-choice set unlike the black logit-QRE curve shown in the bottom-middle panel.

Finally for the KM game, the μ equilibrium is given by

$$\sigma(\varepsilon) \;=\; \frac{(1,\varepsilon^2,\varepsilon)}{(1+\varepsilon+\varepsilon^2)}$$

for $\varepsilon \in [0,1]$, see the orange curve in the upper-right panel of Figure 13, which is similar to the black logit-QRE curve shown in the bottom-right panel.

Except for the KM game, the above formulae may appear somewhat involved. However, they are relatively easy to determine. First, we check for "pure" strategies on each of the six parts of the barycentric division of the simplex, see the middle-left panel of Figure 2. For instance, for the part with $\sigma_R > \sigma_B > \sigma_Y$, we check if the expected payoffs based on the belief $\omega = (1,\varepsilon,\varepsilon^2)/(1+\varepsilon+\varepsilon^2)$ are ranked $\pi_R > \pi_B > \pi_Y$. For the DS and "no logit" games, this yields the first, third, fifth, and seventh cases in the formulae above. The remaining (second, fourth, and sixth) cases simply follow by connecting the pure strategy cases across the different parts of the barycentric division using the payoff indifference curves, see the top panels of Figure 13.

Appendix G. Instructions for the Experiment

The following is the translation of the instructions used in the experiment. The original text in Greek is available from the authors upon request. Participants were recruited from the standard UCY - LExEcon subject pool with no additional restrictions.

Thank you for participating in today's session. Please remain quiet. The entire experiment will be conducted through your computer terminal and you will only interact with other participants via the computer. Please do not talk or make other noises during the experiment. The use of mobile phones and similar devices is not allowed.

General instructions: During the experiment you can earn points. At the end of the experiment you will receive 1 euro for every 250 points you earned. The amount earned may be different for each participant and depends on his/her decisions, the decisions of other participants and luck. In addition, you will receive 5 euros as a show-up fee. Payment in cash will take place privately, immediately after the experiment.

The games: The experiment consists of a series of different games. You will play each game for 15 or 8 periods (this will be announced beforehand). In each period you will be randomly matched with another participant. The participant you are matched with will be different in each period. On your screen you will see a table like the one below:

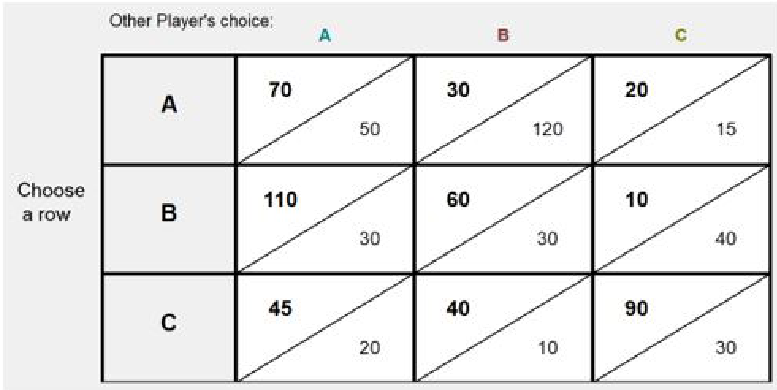

(This is just an example. The numbers in the tables in the experiment will be different.)

Each player selects A, B or C by clicking the corresponding row. Each player wins the points indicated in the cell that corresponds to these choices. In the example above, if you choose C and the other player chooses B, you win 40 points and the other player wins 10 points.

Notice that the other player, exactly like yourself, sees the table with the points he/she may win depending on both of your choices, as well as the points you may win.

Belief task: In each period you are also asked to state your beliefs about the percentage chance that the other player chooses A, B or C.

On the screen, above the game you will see three bars like in the picture below. You can indicate what you believe to be the percentage chance that the other player chooses A, B and C by clicking and dragging the corresponding bar. The sum must always be equal to 100%.

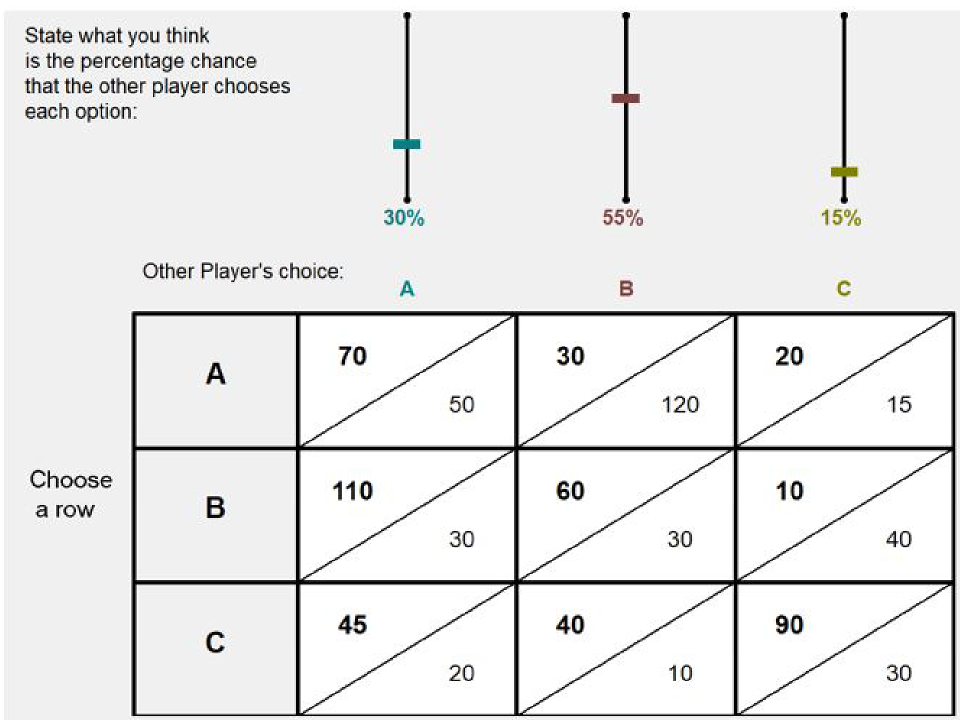

Each time, the computer chooses randomly one of the three bars and on that bar it chooses randomly two points. If your stated belief in that bar is closer to the other player's actual choice than at least one of these two points, then you may win 60 points.

Example 1

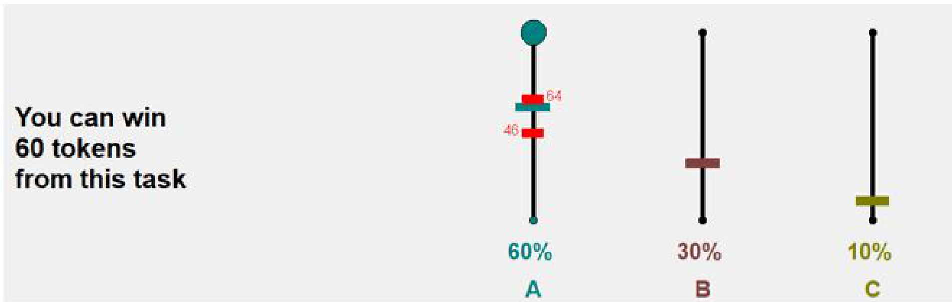

Example 2

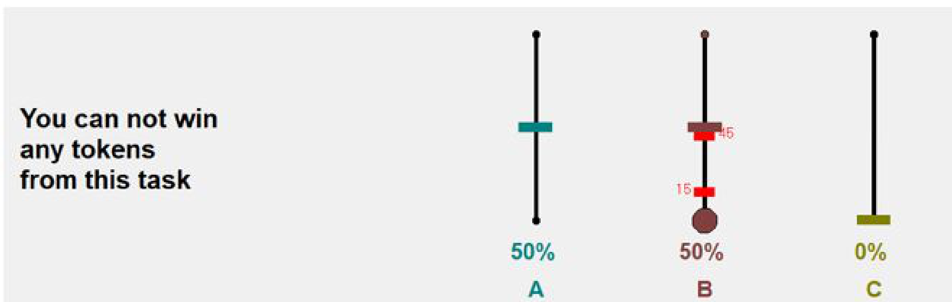

To maximize your chance of winning these points you should state your true beliefs about the other player's percentage chance of making each choice.

Notice:

- You may change your choice as many times as you like (both for the game and the beliefs) When you are ready, press "Submit."
- For the first period of the game you have to wait 30 seconds before you are able to submit your choices. Use that time to familiarize yourself with the new game table.
- Please submit your choices within a reasonable time period. All participants will need to submit their choices before the experiment can move on to the next period.
- After submitting your choices, you will be informed about the other player's choice and your possible point earnings from the game. You will also be informed about the bar and points selected by the computer and whether you may earn points from the belief task.
- To determine your earnings in each period, the computer will randomly choose either the game or the belief task (both with equal probability) and you will earn the points you were awarded in the corresponding task.

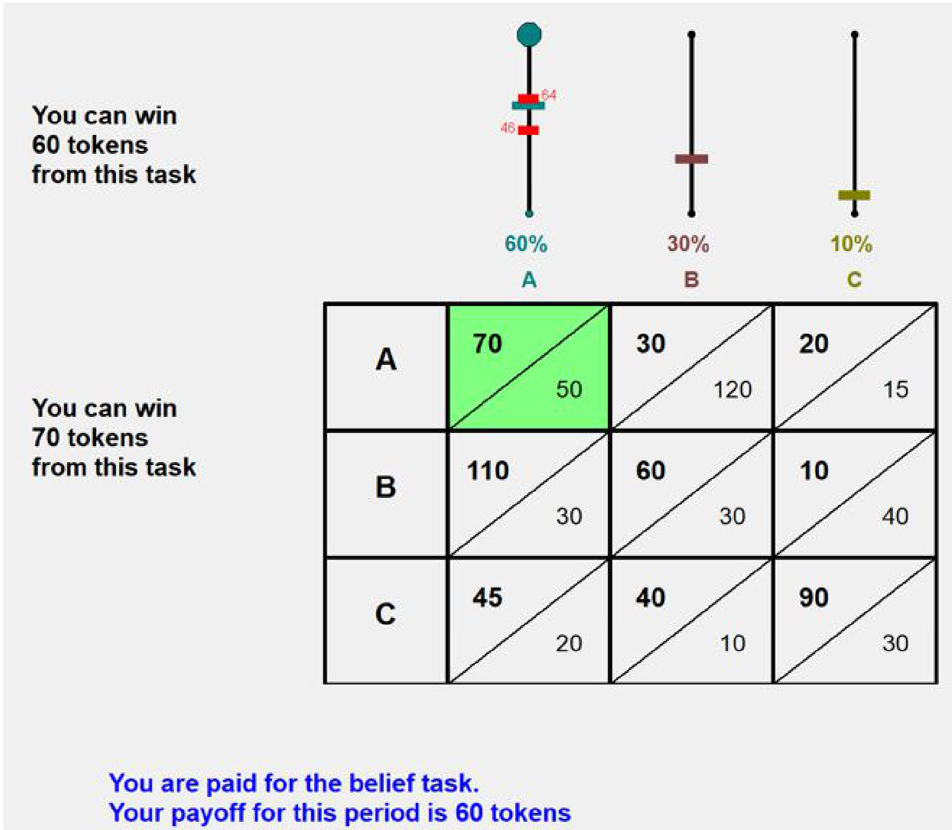
You can win
60 tokens
from this task
64
46
60%
30%
10%
A
B
C
You can win
70 tokens
from this task
A
B
C
70
50
30
120
20
15
110
30
60
30
10
40
45
20
40
10
90
30
You are paid for the belief task.
Your payoff for this period is 60 tokens