Online Appendix

Stochastic Equilibria: Noise in Actions or Beliefs?

Evan Friedman

A Technical axioms: an example with three pure actions

Figure 10 plots the simplex that defines player k's action when she has three pure actions. Consider the sets A, B, and C. A is contained in the interior of the simplex, B is a subset of the edge $\{(p, 1-p, 0) : p \in (0,1)\}$, and $C = \{(1,0,0)\}$ is a corner. Further, A and B "overlap" in the sense that $cl(A) \cap B = B$.

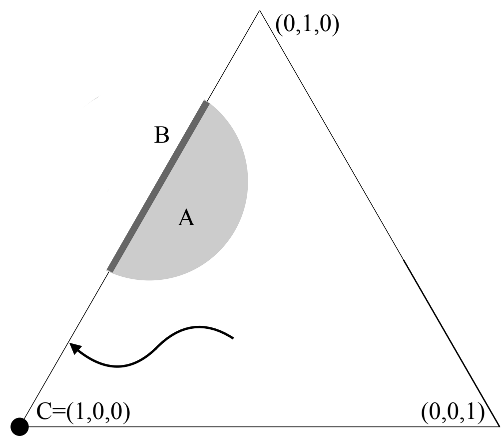

Figure 10: *Technical axioms: an example with three pure actions.*

Now consider the sequence $\{\sigma_k^t\}$ drawn as the black arrow, which starts from the interior and limits to $\sigma_k^\infty \in \{(p, 1-p, 0) : p \in (0,1)\}$. By (B1), for $t < \infty$, $\mu_k^i(A|\sigma_k^t) > 0$ and $\mu_k^i(B|\sigma_k^t) = \mu_k^i(C|\sigma_k^t) = 0$. Also by (B1), $\mu_k^i(A|\sigma_k^\infty) = 0$, $\mu_k^i(B|\sigma_k^\infty) > 0$, and $\mu_k^i(C|\sigma_k^t) = 0$. Thus, there is a discontinuity: $\mu_k^i(B|\sigma_k^t) = 0$ for $t < \infty$, but $\mu_k^i(B|\sigma_k^\infty) > 0$. However, by (B2), there is not a discontinuity in $\mu_k^i(A \cup B|\sigma_k^t)$ as $t \to \infty$ because A and B overlap in the sense of (B2)-*(ii)*.

B The NBE of generalized matching pennies

We derive the set of NBE (and hence QRE) attainable for arbitrary Γ^m.

Along the lines of Example 1, it is easy to show that the reactions for Γ^m depend only on the Nash equilibrium $(\sigma_U^{NE}, \sigma_L^{NE})$ and satisfy:

$$
\Psi_U(\sigma_L) \in \begin{cases} \{0\} & \sigma_L = 0 \\ (0, \frac{1}{2}) & \sigma_L \in (0, \sigma_L^{NE}) \\ \{\frac{1}{2}\} & \sigma_L = \sigma_L^{NE} \\ (\frac{1}{2}, 1) & \sigma_L \in (\sigma_L^{NE}, 1) \\ \{1\} & \sigma_L = 1 \end{cases} \qquad \Psi_L(\sigma_U) \in \begin{cases} \{1\} & \sigma_U = 0 \\ (\frac{1}{2}, 1) & \sigma_U \in (0, \sigma_U^{NE}) \\ \{\frac{1}{2}\} & \sigma_U = \sigma_U^{NE} \\ (0, \frac{1}{2}) & \sigma_U \in (\sigma_U^{NE}, 1) \\ \{0\} & \sigma_U = 1. \end{cases}
$$

The set of attainable NBE is given by $\{(\sigma_U, \sigma_L) | \sigma_U \in \Psi_U(\sigma_L),\ \sigma_L \in \Psi_L(\sigma_U)\}$ and consists of the union of one or two rectangles of positive measure, except when $(\sigma_U^{NE}, \sigma_L^{NE}) = (\frac{1}{2}, \frac{1}{2})$ in which case the unique NBE is $(\sigma_U, \sigma_L) = (\frac{1}{2}, \frac{1}{2})$.

B.1 Mean-unbiasedness

We derive the set of attainable NBE when *unbiasedness* (B4$'$) is replaced with *mean*-unbiasedness.

We first derive the upper and lower bounds on player 1's reaction function $\Psi_U(\sigma_L)$ under the restriction that belief-distributions are unbiased on mean. That is, we find $\bar{\Psi}_U(\sigma_L) = \underset{\sigma_L^* |\ \mathbb{E}(\sigma_L^*) = \sigma_L}{sup} \mathbb{P}(\sigma_L^*(\sigma_L) \geq \sigma_L^{NE})$ and $\underline{\Psi}_U(\sigma_L) = \underset{\sigma_L^* |\ \mathbb{E}(\sigma_L^*) = \sigma_L}{inf} \mathbb{P}(\sigma_L^*(\sigma_L) \geq \sigma_L^{NE})$. These bounds can be achieved through the following family of two-atom belief-distributions:

$$
\hat{\sigma}_L^*(\sigma_L) = \begin{cases} \underline{\sigma}_L(\sigma_L) & \text{w.p. } 1 - \alpha(\sigma_L) \\ \bar{\sigma}_L(\sigma_L) & \text{w.p. } \alpha(\sigma_L) \end{cases},
$$

where $\underline{\sigma}_L(\sigma_L) < \bar{\sigma}_L(\sigma_L)$ are realized beliefs and $\alpha(\sigma_L)$ is a probability of realizing the higher belief; and these terms depend on σ_L. This belief-map violates *continuity* (B1$'$), but it is clear that a continuous version can approximate arbitrarily well the reactions they induce, and hence it is sufficient to find $\bar{\Psi}_U(\sigma_L) = \underset{\underline{\sigma}_L, \bar{\sigma}_L, \alpha | (1-\alpha)\underline{\sigma}_L + \alpha\bar{\sigma}_L = \sigma_L}{sup} \mathbb{P}(\hat{\sigma}_L^*(\sigma_L) \geq \sigma_L^{NE})$ and $\underline{\Psi}_U(\sigma_L) = \underset{\underline{\sigma}_L, \bar{\sigma}_L, \alpha | (1-\alpha)\underline{\sigma}_L + \alpha\bar{\sigma}_L = \sigma_L}{inf} \mathbb{P}(\hat{\sigma}_L^*(\sigma_L) \geq \sigma_L^{NE})$.

Case 1: $\sigma_L \in (0, \sigma_L^{NE})$. It is obvious that $\underline{\Psi}_U(\sigma_L) = 0$. It is easy to check that $\bar{\Psi}_U(\sigma_L)$ is achieved when $\underline{\sigma}_L = 0$ and $\bar{\sigma}_L = \sigma_L^{NE}$, and thus α is determined by the constraint $(1-\alpha)\underline{\sigma}_L + \alpha\bar{\sigma}_L = \alpha\sigma_L^{NE} = \sigma_L$, which implies $\bar{\Psi}_U(\sigma_L) = \alpha = \frac{\sigma_L}{\sigma_L^{NE}}$.

Case 2: $\sigma_L \in (\sigma_L^{NE}, 1)$. It is obvious that $\bar{\Psi}_U(\sigma_L) = 1$. It is easy to check that $\underline{\Psi}_U(\sigma_L)$ is achieved when $\underline{\sigma}_L = \sigma_L^{NE}$ and $\bar{\sigma}_L = 1$, and thus α is determined by the constraint

$(1-\alpha)\underline{\sigma}_L + \alpha\bar{\sigma}_L = (1-\alpha)\sigma_L^{NE} + \alpha = \sigma_L$, which implies $\underline{\Psi}_U(\sigma_L) = \alpha = \frac{\sigma_L - \sigma_L^{NE}}{1-\sigma_L^{NE}}$.

The knife's edge cases $\sigma_L \in \{0, \sigma_L^{NE}, 1\}$ are obvious. Summarizing, and giving the analogue for player 2:

$$\Psi_U^{mean}(\sigma_L) \in \begin{cases} \{0\} & \sigma_L = 0 \\ (0, \frac{\sigma_L}{\sigma_L^{NE}}) & \sigma_L \in (0, \sigma_L^{NE}) \\ (0,1) & \sigma_L = \sigma_L^{NE} \\ (\frac{\sigma_L - \sigma_L^{NE}}{1-\sigma_L^{NE}}, 1) & \sigma_L \in (\sigma_L^{NE}, 1) \\ \{1\} & \sigma_L = 1 \end{cases} \qquad \Psi_L^{mean}(\sigma_U) \in \begin{cases} \{1\} & \sigma_U = 0 \\ (1 - \frac{\sigma_U}{\sigma_U^{NE}}, 1) & \sigma_U \in (0, \sigma_U^{NE}) \\ (0,1) & \sigma_U = \sigma_U^{NE} \\ (0, \frac{1-\sigma_U}{1-\sigma_U^{NE}}) & \sigma_U \in (\sigma_U^{NE}, 1) \\ \{0\} & \sigma_U = 1. \end{cases}$$

The set of attainable "mean NBE" is given by $\{(\sigma_U, \sigma_L) | \sigma_U \in \Psi_U^{mean}(\sigma_L),\ \sigma_L \in \Psi_L^{mean}(\sigma_U)\}$ and consists of a single diamond region which has positive measure for all $(\sigma_U^{NE}, \sigma_L^{NE})$ including the case that $(\sigma_U^{NE}, \sigma_L^{NE}) = (\frac{1}{2}, \frac{1}{2})$. The set always contains an open ball around the NE, and hence has a non-trivial intersection with the set of NBE/QRE.

C Generalizing Lemma 3

For the statement of the lemma and its proof, we let $v_i \in \mathbb{R}_{++}^{J(i)}$ be a utility vector with strictly positive components, where, without loss, $v_{i1} \geq v_{i2} \geq \cdots \geq v_{iJ(i)}$. Let $J^+(v_i) = \{j :\ v_{ij} \geq v_{ik}\ \forall k\}$ and $J^-(v_i) = \{j :\ v_{ij} \leq v_{ik}\ \forall k\}$ be the indices corresponding to the highest and lowest payoff components respectively. Note that $J^+(v_i) \cap J^-(v_i) = \emptyset$ if and only if $v_{ij} \neq v_{ik}$ for some j and k.

Lemma 4. *Let $v_i \in \mathbb{R}_{++}^{J(i)}$ be such that $J^+(u_i') \cap J^-(u_i') = \emptyset$.*
(i) Let Q_i be translation invariant and weakly substitutable, and let $\beta > 1$. $Q_i(\beta v_i) = Q_i(\tilde{v}_i(\beta))$

for some $\tilde{v}_i(\beta)$ such that $\tilde{v}_{il}(\beta) = v_{il} + \delta_l(\beta)$ where $\delta_l(\beta) = 0$ if $l \in J^-(v_i)$, $\delta_l(\beta) > 0$ and $\delta_l(\beta) \to \infty$ if $l \notin J^-(v_i)$, and $\delta_j(\beta) - \delta_k(\beta) \to \infty$ if $v_{ij} > v_{ik}$. $Q_{ij}(\beta v_i) > Q_{ij}(v_i)$ for all $j \in J^+(v_i)$ and $Q_{ik}(\beta v_i) < Q_{ik}(v_i)$ for all $k \in J^-(v_i)$.

(ii) Let Q_i be scale invariant and weakly substitutable, and let $\gamma > 0$. $Q_i(v_i + \gamma e_{J(i)}) = Q_i(\tilde{v}_i(\gamma))$

for some $\tilde{v}_i(\gamma)$ such that $\tilde{v}_{il}(\gamma) = v_{il} + \delta_l(\gamma)$ where $\delta_l(\gamma) = 0$ if $l \in J^+(v_i)$, $\delta_l(\gamma) > 0$ and $\delta_l(\gamma) \to v_{i1} - v_{il}$ if $l \notin J^+(v_i)$. $Q_{ij}(v_i + \gamma e_{J(i)}) < Q_{ij}(v_i)$ for all $j \in J^+(v_i)$ and $Q_{ik}(v_i + \gamma e_{J(i)}) > Q_{ik}(v_i)$ for all $k \in J^-(v_i)$.

Proof. *(i)*: Fix v_i with $J^+(v_i) \cap J^-(v_i) = \emptyset$ and let $\beta > 1$. It is easy to show that $\bar{\gamma}(\beta) \equiv (\beta - 1)v_{i1} > 0$ and $\underline{\gamma}(\beta) \equiv (\beta - 1)v_{iJ(i)} > 0$ satisfy $\beta v_{ij} - \bar{\gamma}(\beta) = v_{ij}$ for all $j \in J^+(v_i)$ and $\beta v_{ik} - \underline{\gamma}(\beta) = v_{ik}$ for all $k \in J^-(v_i)$. Notice that $v_{i1} > v_{ik} \iff \beta v_{ik} - \bar{\gamma}(\beta) < v_{ik}$ and $v_{iJ(i)} < v_{ij} \iff \beta v_{ij} - \underline{\gamma}(\beta) > v_{ij}$ and thus $\beta v_{ik} - \bar{\gamma}(\beta) < v_{ik}$ for all $k \notin J^+(v_i)$ and $\beta v_{ij} - \underline{\gamma}(\beta) > v_{ij}$ for all $j \notin J^-(v_i)$.

By translation invariance, $Q_i(\beta v_i) = Q_i(\beta v_i - \underline{\gamma}(\beta)e_{J(i)}) = Q_i(\beta v_i - \bar{\gamma}(\beta)e_{J(i)})$. Since $Q_i(\beta v_i) = Q_i(\beta v_i - \underline{\gamma}(\beta)e_{J(i)})$, we have that $Q_i(\beta v_i) = Q_i(\tilde{v}(\beta))$ where $\tilde{v}_{il}(\beta) \equiv \beta v_{il} - \underline{\gamma}(\beta) = \beta v_{il} - (\beta - 1)v_{iJ(i)} = v_{il} + \delta_l(\beta)$ where $\delta_l(\beta) = (\beta - 1)(v_{il} - v_{iJ(i)})$.

By weak substitutability, $Q_{ij}(\beta v_i - \bar{\gamma}(\beta)e_{J(i)}) > Q_{ij}(v_i)$ for all $j \in J^+(v_i)$ and $Q_{ik}(\beta v_i - \underline{\gamma}(\beta)e_{J(i)}) < Q_{ik}(v_i)$ for all $k \in J^-(v_i)$. Therefore, $Q_{ij}(\beta v_i) > Q_{ij}(v_i)$ for all $j \in J^+(v_i)$ and $Q_{ik}(\beta v_i) > Q_{ik}(v_i)$ for all $k \in J^-(v_i)$.

(ii): Fix v_i with $J^+(v_i) \cap J^-(v_i) = \emptyset$ and let $\gamma > 0$. It is easy to show that $\bar{\beta}(\gamma) \equiv \frac{v_{i1}}{v_{i1}+\gamma} \in (0,1)$ and $\underline{\beta}(\gamma) \equiv \frac{v_{iJ(i)}}{v_{iJ(i)}+\gamma} \in (0,1)$ satisfy $\bar{\beta}(\gamma)(v_{ij} + \gamma) = v_{ij}$ for all $j \in J^+(v_i)$ and $\underline{\beta}(\gamma)(v_{ik} + \gamma) = v_{ik}$ for all $k \in J^-(v_i)$. Notice that $v_{i1} > v_{ik} \iff \bar{\beta}(\gamma)(v_{ik} + \gamma) > v_{ik}$ and $v_{iJ(i)} < v_{ij} \iff \underline{\beta}(\gamma)(v_{ij} + \gamma) < v_{ij}$ and thus $\bar{\beta}(\gamma)(v_{ik} + \gamma) > v_{ik}$ for all $k \notin J^+(v_i)$ and $\underline{\beta}(\gamma)(v_{ij} + \gamma) < v_{ij}$ for all $j \notin J^-(v_i)$.

By scale invariance, $Q_i(v_i + \gamma e_{J(i)}) = Q_i(\underline{\beta}(\gamma)(v_i + \gamma e_{J(i)})) = Q_i(\bar{\beta}(\gamma)(v_i + \gamma e_{J(i)}))$. Since $Q_i(v_i + \gamma e_{J(i)}) = Q_i(\underline{\beta}(\gamma)(v_i + \gamma e_{J(i)}))$, we have that $Q_i(v_i + \gamma e_{J(i)}) = Q_i(\tilde{v}(\gamma))$ where $\tilde{v}_{il}(\gamma) \equiv \bar{\beta}(\gamma)(v_{il} + \gamma) = \frac{v_{i1}}{v_{i1}+\gamma}(v_{il} + \gamma) = v_{il} + \delta_l(\gamma)$ where $\delta_l(\gamma) = \frac{\gamma(v_{i1} - v_{il})}{v_{i1}+\gamma}$.

By weak substitutability, $Q_{ij}(\bar{\beta}(\gamma)(v_i + \gamma e_{J(i)})) < Q_{ij}(v_i)$ for all $j \in J^+(v_i)$ and $Q_{ik}(\underline{\beta}(\gamma)(v_i + \gamma e_{J(i)})) > Q_{ik}(v_i)$ for all $k \in J^-(v_i)$. Therefore, $Q_{ij}(v_i + \gamma e_{J(i)}) < Q_{ij}(v_i)$ for all $j \in J^+(v_i)$ and $Q_{ik}(v_i + \gamma e_{J(i)}) > Q_{ik}(v_i)$ for all $k \in J^-(v_i)$. □

D Example: QRE with Translation or Scale invariance

In the matching pennies game of Table 11, parameter $Y > 0$ scales player 2's payoffs and hence indexes games in the same scale family. For any fixed Y, the sets of attainable NBE and QRE are identical (Theorem 3). However, while NBE is invariant to an increase in Y, the comparative static for QRE is ambiguous. If QRE is augmented with scale invariance, the QRE prediction coincides with that of NBE trivially. If QRE is augmented with translation invariance, the predictions diverge. Of course, by Lemma 3, one cannot impose both scale and translation invariance.

Example. Let $Y > 0$ and consider the game in Table 11.

(i) Fix σ^*. In the NBE, σ_U and σ_L are constant in Y.

	L	**R**
U	$9, 0$	$0, Y$
D	$0, Y$	$1, 0$

Table 11: *Matching pennies Y.*

(ii) Fix scale invariant Q. In the QRE, σ_U and σ_L are constant in Y.

(iii) Fix translation invariant Q. In the QRE, σ_U and σ_L are strictly decreasing in Y.

Proof.

(i) and *(ii)*: These follow directly from Theorems 4 and 5. *(iii)*: Suppose Q is translation invariant. Any QRE of this game is given as the unique fixed point

$$\sigma_U = Q_U(9\underset{+}{\sigma_L}, 1 - \underset{-}{\sigma_L}) \tag{7}$$

$$\sigma_L = Q_L((1 - \underset{+}{\sigma_U})Y, \underset{-}{\sigma_U}Y). \tag{8}$$

As Y increases, it must be from (7) that either σ_U and σ_L remain constant, σ_U and σ_L increase, or σ_U and σ_L decrease. The first case is impossible since if σ_U were constant, an increase in Y (a scale increase) would change σ_L (by Lemma 3) from (8). The second case is also impossible since $\sigma_U > \frac{1}{2}$ for all $Y > 0$ (as is easy to show along the lines of Example 1), and thus an increase in σ_U and Y must increase $\sigma_U Y$ by more than $(1 - \sigma_U)Y$ increases, which implies a decrease in σ_L from (8) by translation invariance. $\square$

E QRE in sets of binary-action games: necessary conditions

Theorem 6. *Fix dataset $\{\mathcal{G}, \hat{\sigma}, \hat{u}\}$ where $\mathcal{G} = \{g^1, ...g^m, ..., g^M\}$ is a set of games that differ only in payoffs with $J(i) = 2$ for all i, $\hat{\sigma} = \{\hat{\sigma}^m_{ij}\}_{ijm}$ are action frequencies, and $\hat{u} = \{\hat{u}^m_{ij}\}_{ijm}$ are expected utilities, i.e. $\hat{u}^m_{ij} = \bar{u}^m_{ij}(\hat{\sigma}^m_{-i})$. Without loss, relabel all actions so that $\hat{u}^m_{i1} \geq \hat{u}^m_{i2}$ for all m and i.*

(i) $\{\mathcal{G}, \hat{\sigma}, \hat{u}\}$ is consistent with translation invariant QRE only if, for all i:

$$\hat{u}^m_{i1} - \hat{u}^m_{i2} \geq \hat{u}^{m'}_{i1} - \hat{u}^{m'}_{i2} \iff \hat{\sigma}^m_{i1} \geq \hat{\sigma}^{m'}_{i1} \ \forall m, m'$$
$$\textit{and } \hat{\sigma}^m_{i1} \geq \frac{1}{2} \ \forall m.$$

(ii) $\{\mathcal{G}, \hat{\sigma}, \hat{u}\}$ is consistent with scale invariant QRE only if, for all i:

$$\hat{u}_{i1}^m/\hat{u}_{i2}^m \geq \hat{u}_{i1}^{m'}/\hat{u}_{i2}^{m'} \iff \hat{\sigma}_{i1}^m \geq \hat{\sigma}_{i1}^{m'} \ \forall m, m'$$
$$and\ \hat{\sigma}_{i1}^m \geq \frac{1}{2}\ \forall m.$$

Proof. (i): In order for the dataset to be generated by translation invariant QRE, it must be that $\hat{\sigma}_{i1}^m = Q_{i1}(\hat{u}_{i1}^m, \hat{u}_{i2}^m)$ and $\hat{\sigma}_{i1}^{m'} = Q_{i1}(\hat{u}_{i1}^{m'}, \hat{u}_{i2}^{m'})$ for some Q satisfying *responsiveness* (A3), *monotonicity* (A4), and translation invariance. By translation invariance, $Q_{i1}(\hat{u}_{i1}^m, \hat{u}_{i2}^m) = Q_{i1}(\hat{u}_{i1}^m - \hat{u}_{i2}^m, 0)$ and $Q_{i1}(\hat{u}_{i1}^{m'}, \hat{u}_{i2}^{m'}) = Q_{i1}(\hat{u}_{i1}^{m'} - \hat{u}_{i2}^{m'}, 0)$ for all m and m'. Therefore, by (A3), it must be that $\hat{\sigma}_{i1}^m = Q_{i1}(\hat{u}_{i1}^m, \hat{u}_{i2}^m) = Q_{i1}(\hat{u}_{i1}^m - \hat{u}_{i2}^m, 0) \geq Q_{i1}(\hat{u}_{i1}^{m'} - \hat{u}_{i2}^{m'}, 0) = Q_{i1}(\hat{u}_{i1}^{m'}, \hat{u}_{i2}^{m'}) = \hat{\sigma}_{i1}^{m'} \iff \hat{u}_{i1}^m - \hat{u}_{i2}^m \geq \hat{u}_{i1}^{m'} - \hat{u}_{i2}^{m'}$ for all m and m'. $\hat{u}_{i1}^m \geq \hat{u}_{i2}^m$, and hence by (A4), it must be that $\hat{\sigma}_{i1}^m = Q_{i1}(\hat{u}_{i1}^m, \hat{u}_{i2}^m) \geq \frac{1}{2}$ for all m.

(ii): In order for the dataset to be generated by scale invariant QRE, it must be that $\hat{\sigma}_{i1}^m = Q_{i1}(\hat{u}_{i1}^m, \hat{u}_{i2}^m)$ and $\hat{\sigma}_{i1}^{m'} = Q_{i1}(\hat{u}_{i1}^{m'}, \hat{u}_{i2}^{m'})$ for some Q satisfying *responsiveness* (A3), *monotonicity* (A4), and scale invariance. By scale invariance, $Q_{i1}(\hat{u}_{i1}^m, \hat{u}_{i2}^m) = Q_{i1}(\hat{u}_{i1}^m/\hat{u}_{i2}^m, 1)$ and $Q_{i1}(\hat{u}_{i1}^{m'}, \hat{u}_{i2}^{m'}) = Q_{i1}(\hat{u}_{i1}^{m'}/\hat{u}_{i2}^{m'}, 1)$ for all m and m'. Therefore, by (A3), it must be that $\hat{\sigma}_{i1}^m = Q_{i1}(\hat{u}_{i1}^m, \hat{u}_{i2}^m) = Q_{i1}(\hat{u}_{i1}^m/\hat{u}_{i2}^m, 1) \geq Q_{i1}(\hat{u}_{i1}^{m'}/\hat{u}_{i2}^{m'}, 1) = Q_{i1}(\hat{u}_{i1}^{m'}, \hat{u}_{i2}^{m'}) = \hat{\sigma}_{i1}^{m'} \iff \hat{u}_{i1}^m/\hat{u}_{i2}^m \geq \hat{u}_{i1}^{m'}/\hat{u}_{i2}^{m'}$ for all m and m'. $\hat{u}_{i1}^m \geq \hat{u}_{i2}^m$, and hence by (A4), it must be that $\hat{\sigma}_{i1}^m \equiv Q_{i1}(\hat{u}_{i1}^m, \hat{u}_{i2}^m) \geq \frac{1}{2}$ for all m. □

F Logit transform NBE in normal form games

For arbitrary normal form games, we generalize logit transform NBE by parametrizing player i's belief-map over action j of player k as

$$\sigma_{kj}^{i*}(p_k;\tau) = \frac{exp\left(ln\left(\frac{\sigma_{kj}}{1-\sigma_{kj}}\right) + \tau\varepsilon_{kj}^i\right)}{1 + exp\left(ln\left(\frac{\sigma_{kj}}{1-\sigma_{kj}}\right) + \tau\varepsilon_{kj}^i\right)} \cdot \left(\sum_{l=1}^{J(k)} \frac{exp\left(ln\left(\frac{\sigma_{kl}}{1-\sigma_{kl}}\right) + \tau\varepsilon_{kl}^i\right)}{1 + exp\left(ln\left(\frac{\sigma_{kl}}{1-\sigma_{kl}}\right) + \tau\varepsilon_{kl}^i\right)}\right)^{-1},$$

where $\varepsilon_{kj}^i \sim_{iid} \mathcal{N}(0,1)$, and $\tau \in (0,\infty)$ determines the noisiness of beliefs. This belief-map is derived through the following procedure:

1. Map each $\sigma_{kj} \in [0,1]$ to the extended real line via the logit transform $\mathcal{L}(\sigma_{kj}) = ln\left(\frac{\sigma_{kj}}{1-\sigma_{kj}}\right)$, using the convention that $\mathcal{L}(0) = -\infty$ and $\mathcal{L}(1) = \infty$.

2. Add $\tau\varepsilon_{kj}^i$ to each $\mathcal{L}(\sigma_{kj})$.

3. Map each $\mathcal{L}(\sigma_{kj}) + \tau\varepsilon^i_{kj}$ back to $[0, 1]$ via the inverse logit transform

$$\mathcal{L}^{-1}(\mathcal{L}(\sigma_{kj}) + \tau\varepsilon^i_{kj}) = \frac{exp\left(ln\left(\frac{\sigma_{kj}}{1-\sigma_{kj}}\right) + \tau\varepsilon^i_{kj}\right)}{1 + exp\left(ln\left(\frac{\sigma_{kj}}{1-\sigma_{kj}}\right) + \tau\varepsilon^i_{kj}\right)}.$$

4. Normalize the set of $\{\mathcal{L}^{-1}(\mathcal{L}(\sigma_{kl}) + \tau\varepsilon^i_{kl})\}_{l=1}^{J(k)}$ so that they sum to 1 by dividing each $\mathcal{L}^{-1}(\mathcal{L}(\sigma_{kj}) + \tau\varepsilon^i_{kj})$ by the sum

$$\sum_{l=1}^{J(k)} \frac{exp\left(ln\left(\frac{\sigma_{kl}}{1-\sigma_{kl}}\right) + \tau\varepsilon^i_{kl}\right)}{1 + exp\left(ln\left(\frac{\sigma_{kl}}{1-\sigma_{kl}}\right) + \tau\varepsilon^i_{kl}\right)}.$$

This belief-map does not satisfy *unbiasedness* (B4) exactly, but simulations (unreported) suggest that the bias is small for low τ.

G Logit transform NBE with binary actions

For the binary-action case, we derive the CDF and PDF of the logit transform model (5), and show that it satisfies axioms (B1$'$)-(B4$'$). Figure 11 plots the CDF and PDF of belief-distributions for $\tau = 0.5$ and different values of r.

Fact 1. *$r^*(r;\tau)$ has CDF*

$$F^i_k(\bar{r}|r;\tau) = \Phi\left(\frac{1}{\tau}\left[ln\left(\frac{\bar{r}}{1-\bar{r}}\right) - ln\left(\frac{r}{1-r}\right)\right]\right)$$

and PDF

$$f^i_k(\bar{r}|r;\tau) = \phi\left(\frac{1}{\tau}\left[ln\left(\frac{\bar{r}}{1-\bar{r}}\right) - ln\left(\frac{r}{1-r}\right)\right]\right)\frac{1}{\tau}\left(\frac{1}{\bar{r}(1-\bar{r})}\right).$$

.

Proof.

$$F_k^i(\bar{r}|r;\tau) = \mathbb{P}(r^*(r;\tau) \leq \bar{r}) = \mathbb{P}\left(\frac{exp\left(ln\left(\frac{r}{1-r}\right) + \tau\varepsilon_i\right)}{1+exp\left(ln\left(\frac{r}{1-r}\right) + \tau\varepsilon_i\right)} \leq \bar{r}\right)$$
$$= \mathbb{P}\left(ln\left(\frac{r}{1-r}\right) + \tau\varepsilon_i \leq ln\left(\frac{\bar{r}}{1-\bar{r}}\right)\right)$$
$$= \mathbb{P}\left(\varepsilon_i \leq \frac{1}{\tau}\left[ln\left(\frac{\bar{r}}{1-\bar{r}}\right) - ln\left(\frac{r}{1-r}\right)\right]\right)$$
$$= \Phi\left(\frac{1}{\tau}\left[ln\left(\frac{\bar{r}}{1-\bar{r}}\right) - ln\left(\frac{r}{1-r}\right)\right]\right).$$

Notice that F_k^i is differentiable in $\bar{r}$ for $\bar{r}, r \in (0,1)$. Hence, the PDF is easily derived as $f_k^i(\bar{r}|r;\tau) = \frac{\partial F_k^i(\bar{r}|r;\tau)}{\partial \bar{r}}|_{\bar{r},r\in(0,1)}$ using the chain rule. □

Fact 2. *$r^*(r;\tau)$ satisfies (B1$'$)-(B4$'$).*

Proof.
(B1) For any $r \in (0,1)$, $F_k^i(\bar{r}|r;\tau)$ is strictly increasing and continuous in $\bar{r} \in [0,1]$; $r^*(0;\tau) = 0$ and $r^*(1;\tau) = 1$:

That $r^*(0;\tau) = 0$ and $r^*(1;\tau) = 1$ is obvious from the definition of $r^*(\cdot;\tau)$ and the convention that $\mathcal{L}(0) = -\infty$ and $\mathcal{L}(1) = \infty$ where $\mathcal{L}(r) = ln\left(\frac{r}{1-r}\right)$. It is also obvious that, for any $r \in (0,1)$, $F_k^i(\bar{r}|r;\tau)$ is continuous in $\bar{r} \in [0,1]$. For all $r \in (0,1)$, $F_k^i(0|r;\tau) = 0$ and $F_k^i(1|r;\tau) = 1$ (from inspecting $F_k^i(\cdot|\cdot;\tau)$). All we need to show is that $F_k^i(\bar{r}|r;\tau)$ is strictly increasing in $\bar{r} \in (0,1)$ for all $r \in (0,1)$. Notice that $\frac{\partial F_k^i(\bar{r}|r;\tau)}{\partial \bar{r}}|_{\bar{r},r\in(0,1)} = \phi\left(\frac{1}{\tau}\left[ln\left(\frac{\bar{r}}{1-\bar{r}}\right) - ln\left(\frac{r}{1-r}\right)\right]\right)\frac{1}{\tau}\left(\frac{1}{\bar{r}(1-\bar{r})}\right) > 0$ since $\phi(\cdot) > 0$, and we are done.

(B2$'$) For any $\bar{r} \in (0,1)$, $F_k^i(\bar{r}|r;\tau)$ is continuous in $r \in [0,1]$:

We show something stronger, that $F_k^i(\bar{r}|r;\tau)$ is jointly continuous in $(\bar{r}, r) \in (0,1)\times[0,1]$. $F_k^i(\bar{r}|r;\tau) = \Phi\left(\frac{1}{\tau}\left[ln\left(\frac{\bar{r}}{1-\bar{r}}\right) - ln\left(\frac{r}{1-r}\right)\right]\right)$ is obviously continuous for every $(\bar{r}, r) \in (0,1) \times (0,1)$. $F_k^i(\bar{r}|r;\tau)$ is also continuous at all points $(\bar{r}, r) \in (0,1) \times \{0,1\}$. To see this, notice that $F_k^i(\bar{r}|0;\tau) = 1$ for all $\bar{r} \in (0,1)$ and $lim_{r\to 0^+} F_k^i(\bar{r}|r;\tau) = 1$ for all $\bar{r} \in (0,1)$, showing continuity at $(\bar{r}, r) \in (0,1)\times\{0\}$. A similar argument shows continuity at $(\bar{r}, r) \in (0,1)\times\{1\}$.

(B3$'$) For all $r < r' \in [0,1]$, $F_k^i(\bar{r}|r';\tau) \leq F_k^i(\bar{r}|r;\tau)$ for $\bar{r} \in [0,1]$ and $F_k^i(\bar{r}|r';\tau) < F_k^i(\bar{r}|r;\tau)$ for $\bar{r} \in (0,1)$:

(i) If $r' > r \in (0,1)$:

(a) If $\bar{r} \in (0,1)$,

$$\begin{aligned} F_k^i(\bar{r}|r^{'};\tau) &< F_k^i(\bar{r}|r;\tau) \iff \\ \Phi\left(\frac{1}{\tau}\left[ln\left(\frac{\bar{r}}{1-\bar{r}}\right) - ln\left(\frac{r^{'}}{1-r^{'}}\right)\right]\right) &< \Phi\left(\frac{1}{\tau}\left[ln\left(\frac{\bar{r}}{1-\bar{r}}\right) - ln\left(\frac{r}{1-r}\right)\right]\right) \iff \\ \frac{1}{\tau}\left[ln\left(\frac{\bar{r}}{1-\bar{r}}\right) - ln\left(\frac{r^{'}}{1-r^{'}}\right)\right] &< \frac{1}{\tau}\left[ln\left(\frac{\bar{r}}{1-\bar{r}}\right) - ln\left(\frac{r}{1-r}\right)\right] \iff \\ r^{'} &> r. \end{aligned}$$

(b) If $\bar{r} = 0$, $F_k^i(\bar{r}|r;\tau) = F_k^i(\bar{r}|r^{'};\tau) = 0$ (from inspecting $F_k^i(\cdot|\cdot;\tau)$).

(c) If $\bar{r} = 1$, $F_k^i(\bar{r}|r;\tau) = F_k^i(\bar{r}|r^{'};\tau) = 1$ (from inspecting $F_k^i(\cdot|\cdot;\tau)$).

(ii) If $1 = r^{'} > r > 0$, $F_k^i(\bar{r}|r^{'};\tau) = \mathbf{1}_{\{\bar{r}=1\}} \leq F_k^i(\bar{r}|r;\tau)$ for $\bar{r} \in [0,1]$ (using $r^*(1;\tau) = 1$ w.p. 1).

(iii) If $1 > r^{'} > r = 0$, $F_k^i(\bar{r}|r^{'};\tau) \leq F_k^i(\bar{r}|r;\tau) = 1$ for $\bar{r} \in [0,1]$ (using $r^*(0;\tau) = 0$ w.p. 1).

Finally, $\frac{\partial F_k^i(\bar{r}|r;\tau)}{\partial r}|_{\bar{r},r\in(0,1)} = -\phi\left(\frac{1}{\tau}\left[ln\left(\frac{\bar{r}}{1-\bar{r}}\right) - ln\left(\frac{r}{1-r}\right)\right]\right)\frac{1}{\tau}\left(\frac{1}{r(1-r)}\right) < 0$ since $\phi(\cdot) > 0$.

(B4$'$) $F_k^i(r|r;\tau) = \frac{1}{2}$ for $r \in (0,1)$:

For $r \in (0,1)$, $F_k^i(r|r;\tau) = \Phi\left(\frac{1}{\tau}\left[ln\left(\frac{r}{1-r}\right) - ln\left(\frac{r}{1-r}\right)\right]\right) = \Phi(0) = \frac{1}{2}$. □

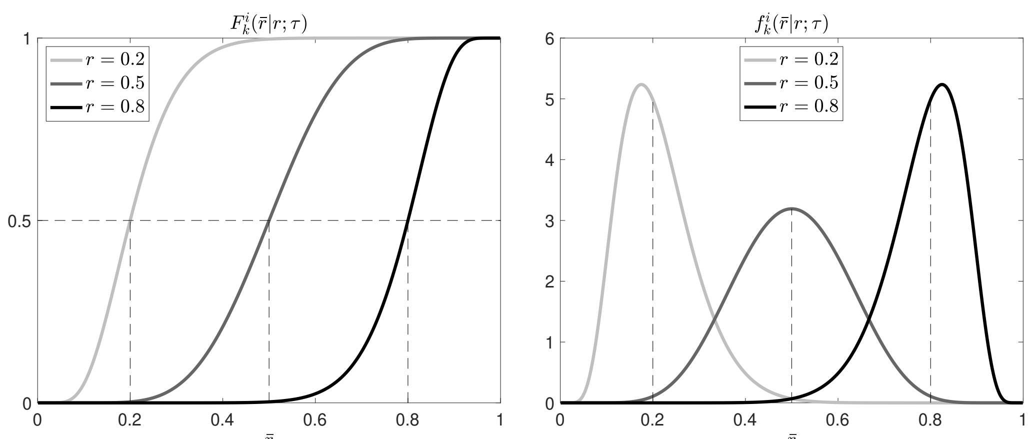

Figure 11: *Logit transform belief-distributions.* This figure plots the CDFs and PDFs of player i's logit transform belief-distributions for noise parameter $\tau = 0.5$ and player k's action $r \in \{0.2, 0.5, 0.8\}$.

H Out-of-sample performance in the McKelvey et al. 2000 data

In Section VI, it was established that NBE and QRE make similar predictions in-sample, but very different predictions out-of-sample across games that differ in scale in the McKelvey et al. [2000] data. We now quantify these effects by examining the prediction error of the parametric models in making out-of-sample predictions across games. For each game $x \in \{A, B, C, D\}$, we take the in-sample estimates $\hat{\tau}^x$ and $\hat{\lambda}^x$ of logit transform NBE and logit QRE from Table 7, and use these to make out-of-sample predictions for game $y \in \{A, B, C, D\}$. We define the *xy-squared distance* for model M with parameter θ by $\mathcal{D}^{xy}(M) = (\sigma_U^y(\hat{\theta}^x)^2 - \hat{\sigma}_U^y)^2 + (\sigma_L^y(\hat{\theta}^x)^2 - \hat{\sigma}_L^y)^2$, where $\hat{\theta}^x$ is the parameter estimated in-sample for game x, $\{\sigma_U^y(\hat{\theta}^x), \sigma_L^y(\hat{\theta}^x)\}$ is the corresponding out-of-sample prediction for game y, and $\{\hat{\sigma}_U^y, \hat{\sigma}_L^y\}$ is the observed action frequency for game y. The *xy-difference in prediction error* between QRE and NBE is given by

$$\Delta\mathcal{D}^{xy} = \mathcal{D}^{xy}(QRE) - \mathcal{D}^{xy}(NBE), \tag{9}$$

which we use to populate the matrix in Table 12. The diagonal entries are in-sample, the off-diagonal entries are out-of-sample, and positive entries indicate that NBE outperforms QRE. From the table, it is clear that the models perform similarly well in-sample, but NBE outperforms QRE in 11 of 12 out-of-sample comparisons and in most cases by a wide margin: the average out-of-sample difference in prediction error (0.0165) is large compared to the in-sample difference (0.0003).

	A	B	C	D
A	0.0067	0.0119	0.0015	0.0034
B	0.0585	−0.0071	0.0277	0.0185
C	0.0377	−0.0004	−0.0004	0.0159
D	0.0062	0.0145	0.0028	0.0022

Table 12: *Out-of-sample differences in prediction error.* The xy-th entry corresponds to $\Delta\mathcal{D}^{xy}$ as in (9) for games $x, y \in \{A, B, C, D\}$ and gives the difference in prediction error between the two models using the data from game x (row) to make predictions about game y (column). Positive (negative) entries indicate that NBE performs better than (worse than) QRE.

I Data and log-likelihoods from existing studies

Table 13 reports empirical frequencies and sample sizes from McKelvey et al. [2000], Selten and Chmura [2008], and Melo et al. [2018], along with the maximized log-likelihoods from fitting logit transform NBE and logit QRE.

Study	Game	Data			NBE	QRE	Vuong
		σ_U	σ_L	N	$-ln(L)$		
mpw2000	A	0.64	0.24	1800	2,216.3	2,286.0	4.4***
	B	0.63	0.24	1200	1,499.5	1,477.9	−6.3***
	C	0.59	0.26	1200	1,563.6	1,603.7	2.1**
	D	0.55	0.33	600	810.0	817.3	1.0
	Pooled	–	–	–	6,090.5	6,285.0	5.9***
sc2008	1	0.08	0.69	9600	8,627.4	8,769.3	11.5***
	2	0.22	0.53	9600	11,713.4	11,712.5	−0.4
	3	0.16	0.79	9600	9,198.1	9,192.7	−0.9
	4	0.29	0.74	9600	11,299.9	11,290.2	−3.3***
	5	0.33	0.66	9600	12,208.3	12,206.5	−3.6***
	6	0.45	0.60	9600	13,088.5	13,087.6	−3.5***
	7	0.14	0.56	4800	5,342.6	5,334.8	−1.2
	8	0.25	0.59	4800	6,072.0	6,091.1	3.9***
	9	0.25	0.83	4800	5,034.1	4,981.6	−9.2***
	10	0.37	0.70	4800	6,225.5	6,192.9	−14.5***
	11	0.33	0.65	4800	6,158.1	6,156.9	−3.4***
	12	0.44	0.60	4800	6,517.9	6,517.6	−2.3**
	Pooled	–	–	–	102,024.1	102,149.3	4.3***
mps2018		*Player 1*	*Player 2*				
		1: 0.25	0.36				
	2	*2*: 0.30	0.44	825	1,756.9	1,757.7	0.6
		J: 0.44	0.20				
		1: 0.34	0.26				
	3	*2*: 0.46	0.32	470	1,015.2	1,009.2	−3.9***
		3: 0.20	0.42				
		1: 0.47	0.49				
	4	*2*: 0.22	0.15	150	315.1	315.0	−0.2
		J: 0.31	0.37				
	Pooled	–	–	–	3,087.6	3,084.5	−1.4*

Table 13: *Data and log-likelihoods from existing studies.*

J Statistical tests: details and robustness

For 2×2 games, and similarly for other games, data is given by counts of each action: N_U, N_D, N_L, and N_R. Logit transform NBE predicts $\sigma_U(\tau)$, $\sigma_D(\tau)$, $\sigma_L(\tau)$, and $\sigma_R(\tau)$; and logit QRE predicts $\sigma_U(\lambda)$, $\sigma_D(\lambda)$, $\sigma_L(\lambda)$, and $\sigma_R(\lambda)$. Let the maximized log-likelihoods be

Study	Game	Vuong				Model	Likelihood Ratio			
		$\rho = 0$	$\rho = 0.50$	$\rho = 0.75$	$\rho = 0.90$		$\rho = 0$	$\rho = 0.50$	$\rho = 0.75$	$\rho = 0.90$
mpw2000	A	4.4***	3.1***	2.2**	1.4*					
	B	−6.3***	−4.4***	−3.1***	−2.0**	NBE	2.1	1.1	0.5	0.2
	C	2.1**	1.5*	1.0	0.6					
	D	1.0	0.7	0.5	0.3	QRE	200.2***	100.1***	50.0***	20.0***
	Pooled	5.9***	4.3***	3.0***	1.9**					
sc2008	1	11.5***	8.1***	5.8***	3.6***					
	2	−0.4	−0.3	−0.2	−0.1	NBE	1,076.4***	538.2***	269.1***	107.6***
	3	−0.9	−0.6	−0.4	−0.3					
	4	−3.3***	−2.4***	−1.7**	−1.1	QRE	1,231.0***	615.5***	307.8***	123.1***
	5	−3.6***	−2.6***	−1.8**	−1.1					
	6	−3.5***	−2.5***	−1.8**	−1.1					
	7	−1.2	−0.8	−0.6	−0.4					
	8	3.9***	2.8***	2.0**	1.2					
	9	−9.2***	−6.5***	−4.6***	−2.9***					
	10	−14.5***	−10.2***	−7.2***	−4.6***					
	11	−3.4***	−2.4***	−1.7**	−1.1					
	12	−2.3**	−1.6*	−1.2	−0.7					
	Pooled	4.3***	3.0***	2.1**	1.3*					
mps2018	2	0.6	0.4	0.3	0.2					
	3	−3.9***	−2.7***	−1.9**	−1.2	NBE	0.9	0.5	0.2	0.1
	4	−0.2	−0.1	−0.1	−0.1					
	Pooled	−1.4*	−1.0	−0.7	−0.5	QRE	5.0*	2.5	1.3	0.5

Table 14: *Robustness of statistical tests to within-subject correlation.* We report Vuong statistics comparing performance of logit transform NBE and logit QRE; and likelihood ratio statistics for tests of the restriction that a model's parameter is fixed across the games within a study. These are presented for different values of ρ following online Appendix J.

$ln(L(\hat{\tau})) = \sum_a N_a ln(\sigma_a(\hat{\tau}))$ and $ln(L(\hat{\lambda})) = \sum_a N_a ln(\sigma_a(\hat{\lambda}))$, respectively, where a sums over the four pure actions and $\hat{\tau}$ and $\hat{\lambda}$ are the MLE-parameters. The log-likelihood ratio is $LR = ln(L(\hat{\tau})) - ln(L(\hat{\lambda}))$.

Let the total number of observations be $\tilde{N} = N_U + N_D + N_L + N_R$ (this is two times N, the number of experimental rounds). Let $l_a = ln(\sigma_a(\hat{\tau})) - ln(\sigma_a(\hat{\lambda}))$ be the log-likelihood ratio of observed action a.

The Vuong statistic comparing model performance is given by $z = \frac{LR}{\sqrt{\tilde{N}}\omega}$, where $\omega^2 = \frac{1}{\tilde{N}} \sum_a N_a l_a^2$. If the observations are independent, then under the null that the models perform equally well, z is distributed as a standard normal. However, since each subject plays the same game several times, there may be some degree of within-subject correlation, and ignoring this will overstate differences in model performance. If there is no auto-correlation, then the original test is fine. If there is perfect auto-correlation, each subject always takes the same action, and the effective sample size shrinks from $\tilde{N}$ to the total number of subjects. In reality, the truth is somewhere in between.

To account for some degree of auto-correlation, we consider smaller effective sample sizes by "throwing out" a fraction ρ of the data. The Vuong statistic becomes $z(\rho) = \frac{(1-\rho)LR}{\sqrt{(1-\rho)\tilde{N}\omega}}$, and $|z(\rho)|$ decreases in ρ as expected. Similarly, we can modify the likelihood ratio statistics for tests of the restriction that a model's parameter is fixed across the games within a study by multiplying by $(1-\rho)$. The results are presented in Table 14 for $\rho \in \{0, 0.50, 0.75, 0.90\}$. All results that are highly significant for $\rho = 0$ remain significant for ρ of at least 0.75, so we conclude that the results are largely robust.

K Risk aversion

Goeree et al. [2003] construct "game 4" in Table 15 to "exaggerate the effects of possible risk aversion" by giving each player a "safe" option with payoffs of 200 and 160 and a "risky" option with payoffs of 370 and 10.[1] Goeree et al. [2003] show that, under risk neutrality, the data $(\hat{\sigma}_U, \hat{\sigma}_L) = (0.53, 0.65)$ is inconsistent with any QRE, and hence NBE also by Theorem 3. With risk aversion, however, both models can rationalize the data.

	L	**R**
U	$370, 200$	$10, 370$
D	$200, 160$	$160, 10$

Table 15: *Matching pennies with safe and risky decisions from* Goeree et al. [2003].

Goeree et al. [2003] fit logit QRE to game 4 and games *A*-*D* from McKelvey et al. [2000] by jointly estimating λ and risk aversion parameter r, where the utility function takes the constant relative risk aversion (CRRA) form:

$$u_r(x) = \frac{x^{1-r} - 10^{1-r}}{370^{1-r} - 10^{1-r}}.$$

Note that utility is normalized so that $u_r(10) = 0$ and $u_r(370) = 1$. To make monetary payoffs comparable across game 4 and games *A*-*D*, the payoffs of *A*-*D* given in Table 4 are first multiplied by 10 before the models are fit. Table 16 is essentially a replication of

[1] Relative to how the matrix is given in Goeree et al. [2003], we have switched the rows so that the game has the form of Table 1.

Table 3 of Goeree et al. [2003], but includes NBE for comparison. We find that the fit for game 4 is statistically the same for both models. However, for games A-D, NBE significantly outperforms QRE. Interestingly, the estimated risk aversion parameters are extremely stable, both across games as well as across models.

Study	Game	Data			NBE		QRE		NBE		QRE		NBE	QRE	Vuong
		σ_U	σ_L	N	σ_U	σ_L	σ_U	σ_L	$\hat{r}$	$\hat{r}$	$\hat{\lambda}$	$\hat{r}$	$-ln(L)$		
ghp2003	4	0.53	0.65	340	0.53	0.67	0.53	0.67	1.10	0.45	6.65	0.45	455.2	455.2	0.0
mpw2000	A	0.64	0.24	1800	0.62	0.25	0.65	0.25	0.72	0.39	23.91	0.43	$5,915.9$	$5,925.7$	1.8^{**}
	B	0.63	0.24	1200	0.62	0.25	0.57	0.25							
	C	0.59	0.26	1200	0.62	0.25	0.58	0.25							
	D	0.55	0.33	600	0.58	0.33	0.59	0.33							

Table 16: *Summary of estimates with risk aversion.*

Finally, we show that for CRRA, but not for general utility functions, NBE predictions are invariant to scaling the monetary payoffs. We think this is potentially important as it provides a robustness argument for the prediction of scale invariance in the presence of risk aversion. For an arbitrary normal form game, we now reinterpret $u_i(a_{ij}, a_{-i})$ as the monetary payoff to player i of taking action a_{ij} given that the opponents' play a_{-i}. The corresponding utility payoff is simply $u_r(u_i(a_{ij}, a_{-i}))$. After a β-scaling of monetary payoffs, the utility payoff becomes $u_r(\beta u_i(a_{ij}, a_{-i})) = \beta^{1-r}(u_i(a_{ij}, a_{-i}))$. Using this, it is clear that β drops out of the expression for the ij-response set (1) for all i and j, from which the result is immediate.

A consequence of this is that when fitting logit transform NBE with risk aversion to games A-D pooled together in Table 16, the predictions are the same for each of A-C, consistent with the fact that scale invariance cannot be rejected statistically (see Table 5).

References

Jacob Goeree, Charles Holt, and Thomas Palfrey. Risk averse behavior in generalized matching pennies games. *Games and Economic Behavior*, 2003. K, 15, K, 1

Richard McKelvey, Thomas Palfrey, and Roberto Weber. The effects of payoff magnitude and heterogeneity on behavior in 2x2 games with unique mixed strategy equilibria. *Journal of Economic Behavior and Organization*, 2000. H, I, K

Emerson Melo, Kiril Pogorelskiy, and Matthew Shum. Testing the quantal response hypothesis. *International Economic Review*, 2018. I

Reinhard Selten and Thorsten Chmura. Stationary concepts for experimental 2x2 games. *American Economic Review*, 2008. I