A. Theoretical Details

A.1. Representing Distributions of Play via Extensive-Form Correlation Plans

As mentioned in the body, every distribution over *randomized* stratregy profiles for the team members is equivalent to a different distribution over *deterministic* strategy profiles by means of Kuhn's theorem (Kuhn, 1953), one of the most fundamental results about extensive-form game playing. Specifically, given two independent mixed strategies $\boldsymbol{y}_{\mathsf{T1}} \in \mathcal{Y}_{\mathsf{T1}}$ and $\boldsymbol{y}_{\mathsf{T2}} \in \mathcal{Y}_{\mathsf{T2}}$ for the team members, let μ_{T1} and μ_{T2} be the distributions over normal-form plans $\Pi_{\mathsf{T1}}, \Pi_{\mathsf{T2}}$ equivalent to $\boldsymbol{y}_{\mathsf{T1}}$ and $\boldsymbol{y}_{\mathsf{T2}}$, respectively. Then, the distribution over reandomized strategy profiles that assignes probability 1 to $(\boldsymbol{y}_{\mathsf{T1}}, \boldsymbol{y}_{\mathsf{T2}})$ is equivalent to the product distribution of μ_{T1} and μ_{T2}, that is, the distirbution over $\Pi_{\mathsf{T1}} \times \Pi_{\mathsf{T2}}$ that picks a generic profile $(\pi_{\mathsf{T1}}, \pi_{\mathsf{T2}})$ with probability $\pi_{\mathsf{T1}}(\pi_{\mathsf{T1}}) \times \pi_{\mathsf{T2}}(\pi_{\mathsf{T2}})$. The reverse is also true: a product distribution over $\Pi_{\mathsf{T1}} \times \Pi_{\mathsf{T2}}$ is equivalent to a distribution over randomized profiles that picks exactly one profile with probability 1.

We now show that a similar result holds when the distribution over normal-form plans is represented as an extensive-form correlation plan. First, we introduce the notion of *product* correlation plan.

Definition 3. *Let $\boldsymbol{\xi}_{\mathsf{T}} \in \mathcal{V}$ be a vector in the von Stengel-Forges polytope. We say that $\boldsymbol{\xi}_{\mathsf{T}}$ is a* product *correlation plan if*

$$\xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, \sigma_{\mathsf{T2}}] = \xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, \varnothing] \cdot \xi_{\mathsf{T}}[\varnothing, \sigma_{\mathsf{T2}}]$$

for all $(\sigma_{\mathsf{T1}}, \sigma_{\mathsf{T2}}) \in \Sigma_{\mathsf{T1}} \bowtie \Sigma_{\mathsf{T2}}$.

Lemma 1. *A product correlation plan is always an element of Ξ_{T}.*

Proof. Let $\boldsymbol{\xi}_{\mathsf{T}}$ be a product correlation plan. Since by definition, $\boldsymbol{\xi}_{\mathsf{T}} \in \mathcal{V}$, the vectors $\boldsymbol{y}_{\mathsf{T1}}, \boldsymbol{y}_{\mathsf{T2}}$ indexed over Σ_{T1} and Σ_{T2}, repsectively, and defined as

$$y[\sigma_{\mathsf{T1}}] = \xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, \varnothing], y[\sigma_{\mathsf{T2}}] = \xi_{\mathsf{T}}[\varnothing, \sigma_{\mathsf{T2}}]$$

are sequence-form strategies. By Kuhn's theorem, there exist distributions $\mu_{\mathsf{T1}}, \mu_{\mathsf{T2}}$ over Π_{T1} and Π_{T2}, respectively, such that

$$y[\sigma_{\mathsf{T1}}] = \sum_{\pi_{\mathsf{T1}} \in \Pi_{\mathsf{T1}}(\sigma_{\mathsf{T1}})} \mu_{\mathsf{T1}}[\pi_{\mathsf{T1}}] \qquad \forall \sigma_{\mathsf{T1}} \in \Sigma_{\mathsf{T1}}, \qquad (3)$$

$$y[\sigma_{\mathsf{T2}}] = \sum_{\pi_{\mathsf{T2}} \in \Pi_{\mathsf{T2}}(\sigma_{\mathsf{T2}})} \mu_{\mathsf{T2}}[\pi_{\mathsf{T2}}] \qquad \forall \sigma_{\mathsf{T2}} \in \Sigma_{\mathsf{T2}}. \qquad (4)$$

Consider the distribution μ_{T} over $\Pi_{\mathsf{T1}} \times \Pi_{\mathsf{T2}}$ defined as the product distribution $\mu_{\mathsf{T1}} \otimes \mu_{\mathsf{T2}}$, that is,

$$\mu_{\mathsf{T}}[\sigma_{\mathsf{T1}}, \sigma_{\mathsf{T2}}] := \mu_{\mathsf{T1}}[\pi_{\mathsf{T1}}] \cdot \mu_{\mathsf{T2}}[\pi_{\mathsf{T2}}]$$

for all $(\pi_{\mathsf{T1}}, \pi_{\mathsf{T2}}) \in \Pi_{\mathsf{T1}} \times \Pi_{\mathsf{T2}}$. We will show that is the extensive-form correlation plan corresponding to μ_{T} according to (2), that is,

$$\xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, \sigma_{\mathsf{T2}}] := \sum_{\substack{\pi_{\mathsf{T1}} \in \Pi_{\mathsf{T1}}(\sigma_{\mathsf{T1}}) \\ \pi_{\mathsf{T2}} \in \Pi_{\mathsf{T2}}(\sigma_{\mathsf{T2}})}} \mu_{\mathsf{T}}[\pi_{\mathsf{T1}}, \pi_{\mathsf{T2}}]$$

for all $(\sigma_{\mathsf{T1}}, \sigma_{\mathsf{T2}}) \in \Sigma_{\mathsf{T1}} \bowtie \Sigma_{\mathsf{T2}}$. Indeed, using the fact that $\boldsymbol{\xi}_{\mathsf{T}}$ is a *product* correlation plan together with (3) and (4):

$$\begin{aligned}
\xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, \sigma_{\mathsf{T2}}] &= \xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, \varnothing] \cdot \xi_{\mathsf{T}}[\varnothing, \sigma_{\mathsf{T2}}] \\
&= y_{\mathsf{T1}}[\sigma_{\mathsf{T1}}] \cdot y_{\mathsf{T2}}[\sigma_{\mathsf{T2}}] \\
&= \left(\sum_{\pi_{\mathsf{T1}} \in \Pi_{\mathsf{T1}}(\sigma_{\mathsf{T1}})} \mu_{\mathsf{T1}}[\pi_{\mathsf{T1}}] \right) \left(\sum_{\pi_{\mathsf{T2}} \in \Pi_{\mathsf{T2}}(\sigma_{\mathsf{T2}})} \mu_{\mathsf{T2}}[\pi_{\mathsf{T2}}] \right) \\
&= \sum_{\substack{\pi_{\mathsf{T1}} \in \Pi_{\mathsf{T1}}(\sigma_{\mathsf{T1}}) \\ \pi_{\mathsf{T2}} \in \Pi_{\mathsf{T2}}(\sigma_{\mathsf{T2}})}} \mu_{\mathsf{T1}}[\pi_{\mathsf{T1}}] \cdot \mu_{\mathsf{T2}}[\pi_{\mathsf{T2}}] \\
&= \sum_{\substack{\pi_{\mathsf{T1}} \in \Pi_{\mathsf{T1}}(\sigma_{\mathsf{T1}}) \\ \pi_{\mathsf{T2}} \in \Pi_{\mathsf{T2}}(\sigma_{\mathsf{T2}})}} \mu_{\mathsf{T}}[\pi_{\mathsf{T1}}, \pi_{\mathsf{T2}}].
\end{aligned}$$

This concludes the proof. □

Lemma 2. *An extensive-form correlation plan is equivalent to a distribution of play for the team that picks one profile of randomized strategies $(\boldsymbol{y}_{\mathsf{T1}}, \boldsymbol{y}_{\mathsf{T2}}) \in \mathcal{Y}_{\mathsf{T1}} \times \mathcal{Y}_{\mathsf{T2}}$ if and only if $\boldsymbol{\xi}_{\mathsf{T}}$ is a product correlation plan. Furthermore, when that is the case, $y_{\mathsf{T1}}[\sigma_{\mathsf{T1}}] = \xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, \varnothing], y_{\mathsf{T2}}[\sigma_{\mathsf{T2}}] = \xi_{\mathsf{T}}[\varnothing, \sigma_{\mathsf{T2}}]$ for all $\sigma_{\mathsf{T1}} \in \Sigma_{\mathsf{T1}}, \sigma_{\mathsf{T2}} \in \Sigma_{\mathsf{T2}}$.*

Proof. The proof of Lemma 1 already shows that when $\boldsymbol{\xi}_{\mathsf{T}}$ is a product correlation plan, it is equivalent to playing according to the distribution of play for the team with singleton support $(\boldsymbol{y}_{\mathsf{T1}}, \boldsymbol{y}_{\mathsf{T2}})$, where $y_{\mathsf{T1}}[\sigma_{\mathsf{T1}}] = \xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, \varnothing], y_{\mathsf{T2}}[\sigma_{\mathsf{T2}}] = \xi_{\mathsf{T}}[\varnothing, \sigma_{\mathsf{T2}}]$ for all $\sigma_{\mathsf{T1}} \in \Sigma_{\mathsf{T1}}, \sigma_{\mathsf{T2}} \in \Sigma_{\mathsf{T2}}$. So, the only statement that remains to prove is that distributions μ_{T} over randomized strategy profiles for the team members with a singleton support are mapped (Eq. (2)) to product correlation plans.

Let $\{(\boldsymbol{y}_{\mathsf{T1}}, \boldsymbol{y}_{\mathsf{T2}})\} \subseteq \mathcal{Y}_{\mathsf{T1}} \times \mathcal{Y}_{\mathsf{T2}}$ be the (singleton) support of μ_{T}, and let $\mu_{\mathsf{T1}}, \mu_{\mathsf{T2}}$ be distributions over Π_{T1} and Π_{T2}, respectively, equivalent to $\boldsymbol{y}_{\mathsf{T1}}$ and $\boldsymbol{y}_{\mathsf{T2}}$. Then,

$$y[\sigma_{\mathsf{T1}}] = \sum_{\pi_{\mathsf{T1}} \in \Pi_{\mathsf{T1}}(\sigma_{\mathsf{T1}})} \mu_{\mathsf{T1}}[\pi_{\mathsf{T1}}] \qquad \forall \sigma_{\mathsf{T1}} \in \Sigma_{\mathsf{T1}}, \qquad (5)$$

$$y[\sigma_{\mathsf{T2}}] = \sum_{\pi_{\mathsf{T2}} \in \Pi_{\mathsf{T2}}(\sigma_{\mathsf{T2}})} \mu_{\mathsf{T2}}[\pi_{\mathsf{T2}}] \qquad \forall \sigma_{\mathsf{T2}} \in \Sigma_{\mathsf{T2}}. \qquad (6)$$

Since by assumption the two team members sample strategies independently, their equivalent distribution of play over determinitic strategies is the product distribution $\mu_{\mathsf{T}} :=$

$\mu_{\mathsf{T1}} \otimes \mu_{\mathsf{T2}}$. Using (2), μ_{T} has a representation as extensive-form correlation plan given by

$$\begin{aligned}\xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, \sigma_{\mathsf{T2}}] &= \sum_{\substack{\pi_{\mathsf{T1}} \in \Pi_{\mathsf{T1}}(\sigma_{\mathsf{T1}}) \\ \pi_{\mathsf{T2}} \in \Pi_{\mathsf{T2}}(\sigma_{\mathsf{T2}})}} \mu_{\mathsf{T}}[\pi_{\mathsf{T1}}, \pi_{\mathsf{T2}}] \\ &= \sum_{\substack{\pi_{\mathsf{T1}} \in \Pi_{\mathsf{T1}}(\sigma_{\mathsf{T1}}) \\ \pi_{\mathsf{T2}} \in \Pi_{\mathsf{T2}}(\sigma_{\mathsf{T2}})}} \mu_{\mathsf{T1}}[\pi_{\mathsf{T1}}] \cdot \mu_{\mathsf{T2}}[\pi_{\mathsf{T2}}] \\ &= \left(\sum_{\pi_{\mathsf{T1}} \in \Pi_{\mathsf{T1}}(\sigma_{\mathsf{T1}})} \mu_{\mathsf{T1}}[\pi_{\mathsf{T1}}]\right)\left(\sum_{\pi_{\mathsf{T2}} \in \Pi_{\mathsf{T2}}(\sigma_{\mathsf{T2}})} \mu_{\mathsf{T2}}[\pi_{\mathsf{T2}}]\right) \\ &= y_{\mathsf{T1}}[\sigma_{\mathsf{T1}}] \cdot y_{\mathsf{T2}}[\sigma_{\mathsf{T2}}] \end{aligned} \tag{7}$$

for all $(\sigma_{\mathsf{T1}}, \sigma_{\mathsf{T2}}) \in \Sigma_{\mathsf{T1}} \times \Sigma_{\mathsf{T2}}$. In particular, choosing $\sigma_{\mathsf{T2}} = \varnothing$ in (7), and using the fact that $y_{\mathsf{T2}}[\varnothing] = 1$, we obtain

$$\xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, \varnothing] = y_{\mathsf{T1}}[\sigma_{\mathsf{T1}}] \qquad \forall\, \sigma_{\mathsf{T1}} \in \Sigma_{\mathsf{T1}}.$$

Similarly,

$$\xi_{\mathsf{T}}[\varnothing, \sigma_{\mathsf{T2}}] = y_{\mathsf{T2}}[\sigma_{\mathsf{T2}}] \qquad \forall\, \sigma_{\mathsf{T2}} \in \Sigma_{\mathsf{T2}}.$$

Substituting the last two equalities into (7) we can write

$$\xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, \sigma_{\mathsf{T2}}] = \xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, \varnothing] \cdot \xi_{\mathsf{T}}[\varnothing, \sigma_{\mathsf{T2}}]$$

for all $(\sigma_{\mathsf{T1}}, \sigma_{\mathsf{T2}}) \in \Sigma_{\mathsf{T1}} \times \Sigma_{\mathsf{T2}}$. That, together with the inclusion $\Xi_{\mathsf{T}} \subseteq \mathcal{V}$, shows that $\boldsymbol{\xi}_{\mathsf{T}}$ is a product correlation plan. ☐

Semi-randomized correlation plans are product plans In the body we mentioned that semi-randomized correlation plans correspond to a distribution of play where one team member plays a deterministic strategy and the other team member plays a randomized strategy. We now give more formal grounding that that assertion.

Lemma 3. *Let $\boldsymbol{\xi}_{\mathsf{T}} \in \Xi^*_{\mathsf{T1}} \cup \Xi^*_{\mathsf{T2}}$ be a semi-randomized plan. Then, $\boldsymbol{\xi}_{\mathsf{T}}$ is a product plan.*

We reuse some ideas that already appeared in Farina & Sandholm (2020) to prove Lemma 3. In particular, in the proof we will make use of the following lemma.

Lemma 4 (Farina & Sandholm (2020, Lemma 6))**.** *Let $\boldsymbol{\xi}_{\mathsf{T}} \in \mathcal{V}$. For all $\sigma_{\mathsf{T1}} \in \Sigma_{\mathsf{T1}}$ such that $\xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, \varnothing] = 0$, $\xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, \sigma_{\mathsf{T2}}] = 0$ for all $\sigma_{\mathsf{T2}} \in \Sigma_{\mathsf{T2}} : \sigma_{\mathsf{T1}} \bowtie \sigma_{\mathsf{T2}}$. Similarly, for all $\sigma_{\mathsf{T2}} \in \Sigma_{\mathsf{T2}}$ such that $\xi_{\mathsf{T}}[\varnothing, \sigma_{\mathsf{T2}}] = 0$, $\xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, \sigma_{\mathsf{T2}}] = 0$ for all $\sigma_{\mathsf{T1}} \in \Sigma_{\mathsf{T1}} : \sigma_{\mathsf{T1}} \bowtie \sigma_{\mathsf{T2}}$.*

Proof of Lemma 3. We will only show the proof for the case $\boldsymbol{\xi}_{\mathsf{T}} \in \Xi^*_{\mathsf{T1}}$. The other case ($\boldsymbol{\xi}_{\mathsf{T}} \in \Xi^*_{\mathsf{T2}}$) is symmetric.

To show that

$$\xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, \sigma_{\mathsf{T2}}] = \xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, \varnothing] \cdot \xi_{\mathsf{T}}[\varnothing, \sigma_{\mathsf{T2}}]$$

for all $(\sigma_{\mathsf{T1}}, \sigma_{\mathsf{T2}}) \in \Sigma_{\mathsf{T1}} \bowtie \Sigma_{\mathsf{T2}}$, we perform induction on the depth of the sequence σ_{T2}. The depth $\mathrm{depth}(\sigma_{\mathsf{T2}})$ of a generic sequence $\sigma_{\mathsf{T2}} = (J, b) \in \Sigma_{\mathsf{T2}}$ of Player i is defined as the number of actions that Player T2 plays on the path from the root of the tree down to action b at information set J, included. Conventionally, we let the depth of the empty sequence be 0.

The base case for the induction proof corresponds to the case where σ_{T2} has depth 0, that is, $\sigma_{\mathsf{T2}} = \varnothing$. In that case, the theorem is clearly true, because $\xi_{\mathsf{T}}[\varnothing, \varnothing] = 1$ as part of the von Stengel-Forges constraints (Definition 1).

Now, suppose that the statement holds as long as $\mathrm{depth}(\sigma_{\mathsf{T2}}) \le d$. We will show that the statement will hold for any $(\sigma_{\mathsf{T1}}, \sigma_{\mathsf{T2}}) \in \Sigma_{\mathsf{T1}} \bowtie \Sigma_{\mathsf{T2}}$ such that $\mathrm{depth}(\sigma_{\mathsf{T2}}) \le d + 1$. Indeed, consider $(\sigma_{\mathsf{T1}}, \sigma_{\mathsf{T2}}) \in \Sigma_{\mathsf{T1}} \bowtie \Sigma_{\mathsf{T2}}$ such that $\sigma_{\mathsf{T2}} = (J, b)$ with $\mathrm{depth}(\sigma_{\mathsf{T2}}) = d + 1$.

There are only two possible cases:

- Case 1: $\xi_{\mathsf{T}}[\varnothing, \sigma_{\mathsf{T2}}] = 0$. From Lemma 4, $\xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, \sigma_{\mathsf{T2}}] = 0$ and the statement holds.
- Case 2: $\xi_{\mathsf{T}}[\varnothing, \sigma_{\mathsf{T2}}] = 1$. From the von Stengel-Forges constraints, $\xi_{\mathsf{T}}[\varnothing, \sigma(J)] = \sum_{b' \in A_J} \xi_{\mathsf{T}}[\varnothing, (J, b')] = 1 + \sum_{b' \in A_J, b' \neq b} \xi_{\mathsf{T}}[\varnothing, (J, b')] \ge 1$. Hence, because all entries of $\xi_{\mathsf{T}}[\varnothing, \sigma_2]$ are in $\{0, 1\}$ by definition of Ξ^*_{T1}, it must be $\xi_{\mathsf{T}}[\varnothing, \sigma(J)] = 1$ and $\xi_{\mathsf{T}}[\varnothing, (J, b')] = 0$ for all $b' \in A_J, b' \neq b$.

  Using the inductive hypothesis, we have that

  $$\xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, \sigma(J)] = \xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, \varnothing] \cdot \xi_{\mathsf{T}}[\varnothing, \sigma(J)] = \xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, \varnothing] \tag{8}$$

  for all $\sigma_{\mathsf{T1}} \in \Sigma_{\mathsf{T1}}, \sigma_{\mathsf{T1}} \bowtie \sigma(J)$. On the other hand, since $\xi_{\mathsf{T}}[\varnothing, (J, b')] = 0$ for all $b' \in A_J, b' \neq b$, from Lemma 4 we have that

  $$\xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, (J, b')] = 0 \quad \forall \sigma_{\mathsf{T1}} \bowtie J, b' \neq b. \tag{9}$$

  Hence, summing over all $b' \in A_J$ and using the von Stengel-Forges constraints, we get

  $$\begin{aligned}\xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, \varnothing] \cdot \xi_{\mathsf{T}}[\varnothing, \sigma_{\mathsf{T2}}] &= \xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, \sigma(J)] \\ &= \sum_{b' \in A_J} \xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, (J, b')] \\ &= \xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, (J, b)] = \xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}, \sigma_{\mathsf{T2}}]\end{aligned}$$

  for all $\sigma_{\mathsf{T1}} \bowtie (J, b)$. This concludes the proof by induction. ☐

So, from Lemma 2 it follows that semi-randomized plans correspond to distributions of play over randomized profiles with the singleton support $(\boldsymbol{y}_{\mathsf{T1}}, \boldsymbol{y}_{\mathsf{T2}}) \in \mathcal{Y}_{\mathsf{T1}} \times \mathcal{Y}_{\mathsf{T2}}$. Furthermore, because of the second part of Lemma 2, when $\boldsymbol{\xi}_{\mathsf{T}} \in \Xi^*_{\mathsf{T1}}$, $\boldsymbol{y}_{\mathsf{T2}}[\sigma_{\mathsf{T2}}] \in \{0, 1\}$ for all $\sigma_{\mathsf{T2}} \in \Sigma_{\mathsf{T2}}$, which means that $\boldsymbol{y}_{\mathsf{T2}}$ is a deterministic strategy for Player T2 (a similar statement holds for $\boldsymbol{\xi}_{\mathsf{T}} \in \Xi^*_{\mathsf{T2}}$).

Convex combinations of product plans Both of the algorithms we presented in the paper ultimately produce an extensive-form correlation plan $\boldsymbol{\xi}_{\mathsf{T}}$ that is a convex combination of semi-randomized plans $\boldsymbol{\xi}_{\mathsf{T}}^{(1)}, \dots, \boldsymbol{\xi}_{\mathsf{T}}^{(n)}$, that is, of the form

$$\boldsymbol{\xi}_{\mathsf{T}} = \lambda^{(1)}\boldsymbol{\xi}_{\mathsf{T}}^{(1)} + \cdot + \lambda^{(n)}\boldsymbol{\xi}_{\mathsf{T}}^{(n)}$$

for $\lambda^{(i)} \geq 0$ such that $\lambda^{(1)} + \cdots + \lambda^{(n)} = 1$. Since semi-randomized correlation plans are product correlation plans (Lemma 3), from Lemma 2 each $\boldsymbol{\xi}_{\mathsf{T}}^{(i)}$ is equivalent to the team playing a single profile of randomized strategies $(\boldsymbol{y}_{\mathsf{T1}}^{(i)}, \boldsymbol{y}_{\mathsf{T2}}^{(i)}) \in \mathcal{Y}_{\mathsf{T1}} \times \mathcal{Y}_{\mathsf{T2}}$ with probability 1. By linearity, it is immediate to show that $\boldsymbol{\xi}_{\mathsf{T}}$ is equivalent to playing according to the distribution over randomized strategies for the team that picks $(\boldsymbol{y}_{\mathsf{T1}}^{(i)}, \boldsymbol{y}_{\mathsf{T2}}^{(i)})$ with probability $\lambda^{(i)}$.

A.2. TMECor Formulation Based on Extensive-Form Correlation Plans

Proposition 1. *An extensive-form correlation plan $\boldsymbol{\xi}_{\mathsf{T}}$ is a TMECor if and only if it is a solution to the LP*

$$\begin{cases} \underset{\boldsymbol{\xi}_{\mathsf{T}}}{\arg\max} \quad v_\varnothing, \quad \text{subject to:} \\ \text{①}\ v_I - \displaystyle\sum_{\substack{I' \in \mathcal{I}_{\mathsf{O}} \\ \sigma_{\mathsf{O}}(I') = (I,a)}} v_{I'} \leq \sum_{\substack{z \in Z \\ \sigma_{\mathsf{O}}(z) = (I,a)}} \hat{u}_{\mathsf{T}}(z)\xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}(z), \sigma_{\mathsf{T2}}(z)] \qquad \forall\, (I,a) \in \Sigma_{\mathsf{O}} \setminus \{\varnothing\} \\ \text{②}\ v_\varnothing - \displaystyle\sum_{\substack{I' \in \mathcal{I}_{\mathsf{O}} \\ \sigma_{\mathsf{O}}(I') = \varnothing}} v_{I'} \leq \sum_{\substack{z \in Z \\ \sigma_{\mathsf{O}}(z) = \varnothing}} \hat{u}_{\mathsf{T}}(z)\xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}(z), \sigma_{\mathsf{T2}}(z)] \\ \text{③}\ v_\varnothing \text{ free}, v_I \text{ free} \quad \forall\, I \in \mathcal{I}_{\mathsf{O}} \\ \text{④}\ \boldsymbol{\xi}_{\mathsf{T}} \in \Xi_{\mathsf{T}}. \end{cases}$$

Proof. We follow the steps mentioned in the body, starting from the bilinear saddle point problem formulation of the problem of computing a TMECor strategy for the team:

$$\underset{\boldsymbol{\xi}_{\mathsf{T}} \in \Xi_{\mathsf{T}}}{\arg\max}\ \min_{\boldsymbol{y}_{\mathsf{O}} \in \mathcal{Y}_{\mathsf{O}}} \sum_{z \in Z} \hat{u}_{\mathsf{T}}(z)\xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}(z), \sigma_{\mathsf{T2}}(z)]y[\sigma_{\mathsf{O}}(z)].$$

Expanding the constraint $\boldsymbol{y}_{\mathsf{O}} \in \mathcal{Y}_{\mathsf{O}}$ using the *sequence-form constraints* (Koller et al., 1996; von Stengel, 1996), the inner minimization problem is

$$(P): \begin{cases} \displaystyle\min_{\boldsymbol{y}_{\mathsf{O}}} \quad \sum_{z \in Z} \hat{u}_{\mathsf{T}}(z)\xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}(z), \sigma_{\mathsf{T2}}(z)]y[\sigma_{\mathsf{O}}(z)] \\ \text{①}\ -y[\sigma(I)] + \displaystyle\sum_{a \in A_I} y_{\mathsf{O}}[(I,a)] = 0 \quad \forall I \in \mathcal{I}_{\mathsf{O}} \\ \text{②}\ y_{\mathsf{O}}[\varnothing] = 1 \\ \text{③}\ y_{\mathsf{O}}[\sigma_{\mathsf{O}}] \geq 0 \qquad \forall\, \sigma_{\mathsf{O}} \in \Sigma_{\mathsf{O}}. \end{cases}$$

Introducing the free dual variables $\{v_I : I \in \mathcal{I}_{\mathsf{O}}\}$ for Constraint ①, and the free dual variable $v_\varnothing$ for Constraint ②, we obtain the dual linear program

$$(D): \begin{cases} \displaystyle\max_{v_I, v_\varnothing} \quad v_\varnothing, \quad \text{subject to:} \\ \text{①}\ v_I - \displaystyle\sum_{\substack{I' \in \mathcal{I}_{\mathsf{O}} \\ \sigma_{\mathsf{O}}(I') = (I,a)}} v_{I'} \leq \sum_{\substack{z \in Z \\ \sigma_{\mathsf{O}}(z) = (I,a)}} \hat{u}_{\mathsf{T}}(z)\xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}(z), \sigma_{\mathsf{T2}}(z)] \quad \forall\, (I,a) \in \Sigma_{\mathsf{O}} \setminus \{\varnothing\} \\ \text{②}\ v_\varnothing - \displaystyle\sum_{\substack{I' \in \mathcal{I}_{\mathsf{O}} \\ \sigma_{\mathsf{O}}(I') = \varnothing}} v_{I'} \leq \sum_{\substack{z \in Z \\ \sigma_{\mathsf{O}}(z) = \varnothing}} \hat{u}_{\mathsf{T}}(z)\xi_{\mathsf{T}}[\sigma_{\mathsf{T1}}(z), \sigma_{\mathsf{T2}}(z)] \\ \text{③}\ v_\varnothing \text{ free}, v_I \text{ free} \quad \forall\, I \in \mathcal{I}_{\mathsf{O}}. \end{cases}$$

So, $\boldsymbol{\xi}_{\mathsf{T}}$ is a TMECor if and only if it is a solution of $\arg\max_{\boldsymbol{\xi}_{\mathsf{T}} \in \Xi_{\mathsf{T}}}(D)$, which is exactly the statement. □

A.3. Semi-Randomized Correlation Plans

Proposition 3. *In every game, Ξ_{T} is the convex hull of the set Ξ^*_{T1}, or equivalently of the set Ξ^*_{T2}. Formally, $\Xi_{\mathsf{T}} = \operatorname{co}\Xi^*_{\mathsf{T1}} = \operatorname{co}\Xi^*_{\mathsf{T2}} = \operatorname{co}(\Xi^*_{\mathsf{T1}} \cup \Xi^*_{\mathsf{T2}})$.*

Proof. We will show that $\Xi_{\mathsf{T}} = \operatorname{co}\Xi^*_{\mathsf{T1}}$. The proof that $\Xi_{\mathsf{T}} = \operatorname{co}\Xi^*_{\mathsf{T2}}$ is symmetric.

We will break the proof of $\Xi_{\mathsf{T}} = \operatorname{co}\Xi^*_{\mathsf{T1}}$ into two parts:

($\subseteq$) In the first part of the proof, we argue that $\Xi^*_{\mathsf{T1}} \subseteq \Xi_{\mathsf{T}}$. This is straightforward: from Lemma 3 we know that all elements of Ξ^*_{T1} are product correlation plans (Definition 3), which implies that $\Xi^*_{\mathsf{T1}} \subseteq \Xi_{\mathsf{T}}$ by Lemma 1. Since convex hulls preserve inclusions, we have

$$\operatorname{co}\Xi^*_{\mathsf{T1}} \subseteq \operatorname{co}\Xi_{\mathsf{T}},$$

which is exactly the statement $\Xi^*_{\mathsf{T1}} \subseteq \Xi_{\mathsf{T}}$ upon using the known fact that Ξ_{T} is a convex polytope and therefore $\operatorname{co}\Xi_{\mathsf{T}} = \Xi_{\mathsf{T}}$.

($\supseteq$) To complete the proof, we now argue that the reverse inclusion, namely $\Xi_{\mathsf{T}} \subseteq \operatorname{co}\Xi^*_{\mathsf{T1}}$, also holds. Let $f : \mu_{\mathsf{T}} \mapsto \xi_{\mathsf{T}}$ be the mapping from the distribution of play $\mu_{\mathsf{T}} \in \Delta(\Pi_{\mathsf{T1}} \times \Pi_{\mathsf{T2}})$ to its corresponding extensive-form correlation plan defined in Eq. (2). By definition, $\Xi_{\mathsf{T}} = f(\Delta(\Pi_{\mathsf{T1}} \times \Pi_{\mathsf{T2}}))$. Let $\mathbb{1}_{(\pi_{\mathsf{T1}}, \pi_{\mathsf{T2}})}$ denote the distribution of play with singleton support $(\pi_{\mathsf{T1}}, \pi_{\mathsf{T2}})$, that is, the distribution of play that assigns the deterministic strategy profile $(\pi_{\mathsf{T1}}, \pi_{\mathsf{T2}})$ for the team with probability 1. Since f is linear, and since

$$\Delta(\Pi_{\mathsf{T1}} \times \Pi_{\mathsf{T2}}) = \operatorname{co}\{\mathbb{1}_{(\pi_{\mathsf{T1}}, \pi_{\mathsf{T2}})} : \pi_{\mathsf{T1}} \in \Pi_{\mathsf{T1}}, \pi_{\mathsf{T2}} \in \Pi_{\mathsf{T2}}\},$$

we have

$$\Xi_{\mathsf{T}} = \operatorname{co}\{f(\mathbb{1}_{(\pi_{\mathsf{T1}},\pi_{\mathsf{T2}})}) : \pi_{\mathsf{T1}} \in \Pi_{\mathsf{T1}}, \pi_{\mathsf{T2}} \in \Pi_{\mathsf{T2}}\}.$$

Hence, to conclude the proof of this part, it will be enough to show that for each $\pi_{\mathsf{T1}} \in \Pi_{\mathsf{T1}}, \pi_{\mathsf{T2}} \in \Pi_{\mathsf{T2}}$, it holds that $f(\mathbb{1}_{(\pi_{\mathsf{T1}},\pi_{\mathsf{T2}})}) \in \Xi^*_{\mathsf{T1}}$. Since $\mathbb{1}_{(\pi_{\mathsf{T1}},\pi_{\mathsf{T2}})}$ assigns probability 1 to one profile and 0 to all other profiles, $f(\mathbb{1}_{(\pi_{\mathsf{T1}},\pi_{\mathsf{T2}})})$ is an extensive-form correlation plan whose entris are all in $\{0,1\}$. So, in particular, $f(\mathbb{1}_{(\pi_{\mathsf{T1}},\pi_{\mathsf{T2}})}) \in \Xi^*_{\mathsf{T1}}$. This concludes the proof of the inclusion $\Xi_{\mathsf{T}} \subseteq \operatorname{co}\Xi^*_{\mathsf{T1}}$.

Together, the two statements that we just prove show that $\Xi_{\mathsf{T}} = \operatorname{co}\Xi^*_{\mathsf{T1}}$.

Finally, using the fact that unions and convex hulls commute, we have

$$\operatorname{co}(\Xi^*_{\mathsf{T1}} \cup \Xi^*_{\mathsf{T2}}) = (\operatorname{co}\Xi^*_{\mathsf{T1}}) \cup (\operatorname{co}\Xi^*_{\mathsf{T2}}) = \Xi_{\mathsf{T}} \cup \Xi_{\mathsf{T}} = \Xi_{\mathsf{T}},$$

thereby concluding the proof. □

B. Game Instances

The size of the parametric instances we use as benchmark is described in Table 1. In the following, we provide a detailed explanation of the rules of each game.

Kuhn poker Two-player Kuhn poker was originally proposed by Kuhn (1950). We employ the three-player variation described in Farina et al. (2018). In a three-player Kuhn poker game with rank r there are r possible cards. At the beginning of the game, each player pays one chip to the pot, and each player is dealt a single private card. The first player can check or bet, i.e., putting an additional chip in the pot. Then, the second player can check or bet after a first player's check, or fold/call the first player's bet. If no bet was previously made, the third player can either check or bet. Otherwise, the player has to fold or call. After a bet of the second player (resp., third player), the first player (resp., the first and the second players) still has to decide whether to fold or to call the bet. At the showdown, the player with the highest card who has not folded wins all the chips in the pot.

Goofspiel This bidding game was originally introduced by Ross (1971). We use a 3-rank variant, that is, each player has a hand of cards with values $\{-1, 0, 1\}$. A third stack of cards with values $\{-1, 0, 1\}$ is shuffled and placed on the table. At each turn, a prize card is revealed, and each player privately chooses one of his/her cards to bid, with the highest card winning the current prize. In case of a tie, the prize is split evenly among the winners. After 3 turns, all the prizes have been dealt out and the payoff of each player is computed as follows: each prize card's value is equal to its face value and the players' scores are computed as the sum of the values of the prize cards they have won.

Goofspiel with limited information This is a variant of Goofspiel introduced by Lanctot et al. (2009). In this variation, in each turn the players do not reveal the cards that they have played. Rather, they show their cards to a fair umpire, which determines which player has played the highest card and should therefore received the prize card. In case of tie, the umpire directs the players to split the prize evenly among the winners, just like in the Goofspiel game. This makes the game strategically more challenging as players have less information regarding previous opponents' actions.

Leduc poker We use a three-player version of the classical Leduc hold'em poker introduced by Southey et al. (2005). We employ game instances of rank 3, in which the deck consists of three suits with 3 cards each. Our instances are parametric in the maximum number of bets, which in limit hold'em is not necessarely tied to the number of players. The maximum number of raise per betting round can be either 1, 2 or 3. As the game starts players pay one chip to the pot. There are two betting rounds. In the first one a single private card is dealt to each player while in the second round a single board card is revealed. The raise amount is set to 2 and 4 in the first and second round, respectively.

Liar's dice Liar's dice is another standard benchmark introduced by Lisý et al. (2015). In our three-player implementation, at the beginning of the game each of the three players privately rolls an unbiased k-face die. Then, the three players alternate in making (potentially false) claims about their toss. The first player begins bidding, announcing any face value up to k and the minimum number of dice that the player believes are showing that value among the dice of all the players. Then, each player has two choices during their turn: to make a higher bid, or to challenge the previous bid by declaring the previous bidder a "liar". A bid is higher than the previous one if either the face value is higher, or the number of dice is higher. If the current player challenges the previous bid, all dice are revealed. If the bid is valid, the last bidder wins and obtains a reward of +1 while the challenger obtains a negative payoff of -1. Otherwise, the challenger wins and gets reward +1, and the last bidder obtains reward of -1. All the other players obtain reward 0. We test our algorithms on Liar's dice instances with $k = 3$ and $k = 4$.

C. Additional Experimental Results

All experiments were run 10 times, and the experimental tables show average run times. We always use the same random seed to sample no-regret strategies for the team members in the seeding phase of our column-generation algorithm. The seed was never changed, and we don't treat it as a hyperparameter. So, all algorithms are deterministic, and the only source of randomness in the run time is due to system load. Consequently, we observed small standard deviations in the run times, less than 10% in all cases.

We used the same time limit for FTP that was found to be beneficial by the original authors (Farina et al., 2018), namely 15 seconds. For FTP and CG-18, we used the original implementations, with permission from the authors. In all algorithms, we observed that the majority of time is spent within Gurobi.

Table 4 and Table 5 show the comparison between our column-generation algorithm, FTP, and CG-18 when the opponent plays as the first and as the second player, respectively.

Comparison between the Algorithm of Section 5 and the Prior State of the Art

Depending on the cap n on the number or semi-randomized correlation plans, the algorithm we describe in Section 5 might not reach the optimal TMECor value for the team (although, as we argue in Section 7, a very small n already guarantees a large fraction of the optimal value empirically).

For completeness, we report the run time of the algorithm for a sample instance. We employ instance [H] with $\bigcirc = 3$ as it is has a good trade-off between dimensions and manageability. When $n = 1$ the algorithm reaches an optimal solution in 9.21s. The optimal solution with $n = 1$ achieves 15% of the optimal utility with no restrictions on the number of plans. With $n = 2$ the run time is 12m 05s and the solution reaches 90% of the optimal value. With $n = 3$ the run time is 1h 53m and the solution guarantees 92% of the optimal value.

The column-generation algorithm has better run time performances and guarantees to reach an optimal solution without having to pick the right support size. However, we observe that the algorithm of Section 5 already outperforms FTP and CG-18. Specifically, FTP cannot reach a strategy guaranteeing 50% of the optimal utility within the time limit, while our algorithm guarantees 90% of the optimal value within roughly 10 minutes. On the other hand, CG-18 cannot complete even a single iteration within the time limit. This confirms the our pricing formulation is significantly tighter than previous formulations.

Game	Ours		Fictitious Team Play (FTP)			CG-18	Pricers		Team utility after seeding			TMECor value
	Seeded	Not seed.	$\epsilon = 50\%$	$\epsilon = 10\%$	$\epsilon = 1\%$		Relax.	MIP	$m = 0$	100	1000	
[A]	2ms	2ms	0ms†	15.00s†	2m 35s†	66ms	5	0	–0.481	–0.133	–0.133	0
[B]	17ms	27ms	0ms	16m 39s	$>$ 6h	1.01s	0	3	–0.307	0.037	0.037	0.038
[C]	4.67s	6.86s	7m 36s	$>$ 6h	$>$ 6h	$>$ 6h	7	38	–0.381	0.055	0.058	0.066
[D]	302ms	654ms	2ms	$>$ 6h	$>$ 6h	1m 56s	17	0	–1.000	0.239	0.251	0.252
[E]	821ms	1.52s	6ms	$>$ 6h	$>$ 6h	23m 17s	31	0	–1.110	0.242	0.249	0.253
[F]	5.73s	15.34s	19m 25s	$>$ 6h	$>$ 6h	$>$ 6h	1	0	–0.926	0	0	0
[G]	8.08s	1h 09m	$>$ 6h	$>$ 6h	$>$ 6h	$>$ 6h	0	2	–0.696	0.039	0.063	0.063
[H]	32.12s	40.16s	2h 49m	$>$ 6h	$>$ 6h	$>$ 6h	1	90	–2.000	0.001	0.173	0.199
[I]	14m 51s	15m 12s	$>$ 6h	$>$ 6h	$>$ 6h	$>$ 6h	0	232	–2.000	–0.020	0.105	0.186
[J]	11m 23s	8m 24s	$>$ 6h	$>$ 6h	$>$ 6h	$>$ 6h	1042	123	–2.000	–0.433	0.395	0.549

(a) — Comparison of run times **(b)** **(c)**

Table 4: **Results for O = 1. (a)** Runtime comparison between our column generation algorithm, FTP, and CG-18. The seeded version of our algorithm runs $m = 1000$ iterations of CFR+ (Section 6.2), while the non seeded version runs $m = 0$. '†': since the TMECor value for the game is exactly zero, we measure how long it took the algorithm to find a distribution with expected value at least $-\epsilon/10$ for the team. **(b)** Number of times the pricing problem for our column-generation algorithm was solved to optimality by the linear relaxation ('Relax') and by the MIP solver ('MIP') when using our column-generation algorithm (seeded version with $m = 1000$). **(c)** Quality of the initial strategy of the team obtained for varying sizes of S compared to the expected utility of the team at the TMECor.

Game	Ours		Fictitious Team Play (FTP)			CG-18	Pricers		Team utility after seeding			TMECor value
	Seeded	Not seed.	$\epsilon = 50\%$	$\epsilon = 10\%$	$\epsilon = 1\%$		Relax.	MIP	$m = 0$	100	1000	
[A]	0ms	3ms	0ms†	19s†	3m 09s†	147ms	1	0	–0.667	0	0	0
[B]	0ms	12ms	1m 39s	$>$ 6h	$>$ 6h	7.53s	1	0	–0.281	0.027	0.027	0.027
[C]	4.01s	4.22s	48m 08s	$>$ 6h	$>$ 6h	$>$ 6h	6	26	–0.636	0.018	0.027	0.038
[D]	221ms	696ms	1ms	$>$ 6h	$>$ 6h	1m 46s	13	0	–1.000	0.247	0.252	0.252
[E]	1.11s	1.37s	1.39s	$>$ 6h	$>$ 6h	12m 30s	42	0	–1.110	0.241	0.252	0.253
[F]	20.72s	1m 03s	1h 30m	$>$ 6h	$>$ 6h	$>$ 6h	38	2	–0.778	0.246	0.256	0.256
[G]	5h 48m	5h 44m	$>$ 6h	$>$ 6h	$>$ 6h	$>$ 6h	140	17	–0.969	0.260	0.264	—
[H]	2m 02s	2m 44s	$>$ 6h	$>$ 6h	$>$ 6h	$>$ 6h	24	168	–1.679	0.061	0.164	0.253
[I]	27m 17s	27m 50s	$>$ 6h	$>$ 6h	$>$ 6h	$>$ 6h	6	553	–1.911	–0.011	0.076	0.183
[J]	28m 05s	11m 21s	$>$ 6h	$>$ 6h	$>$ 6h	$>$ 6h	4600	254	–2.000	0.190	0.392	0.628

(a) — Comparison of run times **(b)** **(c)**

Table 5: **Results for O = 2. (a)** Runtime comparison between our column generation algorithm, FTP, and CG-18. The seeded version of our algorithm runs $m = 1000$ iterations of CFR+ (Section 6.2), while the non seeded version runs $m = 0$. '†': since the TMECor value for the game is exactly zero, we measure how long it took the algorithm to find a distribution with expected value at least $-\epsilon/10$ for the team. **(b)** Number of times the pricing problem for our column-generation algorithm was solved to optimality by the linear relaxation ('Relax') and by the MIP solver ('MIP') when using our column-generation algorithm (seeded version with $m = 1000$). **(c)** Quality of the initial strategy of the team obtained for varying sizes of S compared to the expected utility of the team at the TMECor. 'oom': out of memory.