Coarse Correlation in Extensive-Form Games

Gabriele Farina,[1] **Tommaso Bianchi,**[2] **Tuomas Sandholm**[1,3,4,5,]
[1]Computer Science Department, Carnegie Mellon University, [2]DEIB, Politecnico di Milano
[3]Strategic Machine, Inc., [4]Strategy Robot, Inc., [5]Optimized Markets, Inc.
gfarina@cs.cmu.edu, tommaso4.bianchi@mail.polimi.it, sandholm@cs.cmu.edu

Abstract

Coarse correlation models strategic interactions of rational agents complemented by a *correlation device* which is a mediator that can recommend behavior but not enforce it. Despite being a classical concept in the theory of *normal-form* games since 1978, not much is known about the merits of coarse correlation in *extensive-form* settings. In this paper, we consider two instantiations of the idea of coarse correlation in extensive-form games: normal-form coarse-correlated equilibrium (NFCCE), already defined in the literature, and *extensive-form coarse-correlated equilibrium (EFCCE)*, a new solution concept that we introduce. We show that EFCCEs are a subset of NFCCEs and a superset of the related extensive-form correlated equilibria. We also show that, in n-player extensive-form games, social-welfare-maximizing EFCCEs and NFCEEs are bilinear saddle points, and give new efficient algorithms for the special case of two-player games with no chance moves. Experimentally, our proposed algorithm for NFCCE is two to four orders of magnitude faster than the prior state of the art.

Introduction

As a generic term, *correlated equilibrium* denotes a family of solution concepts whereby a mediator that can recommend behavior, but not enforce it, complements the interaction of rational agents. Before the game starts, the mediator—also called a *correlation device*—samples a tuple of normal-form plans (one for each player) from a publicly known correlated distribution. The mediator then proceeds to privately ask each player whether they would like to commit to playing according to the plan that was sampled for them. Being part of an *equilibrium*, the correlated distribution must be such that no player can benefit from not following the recommendations, assuming all other players follow. Example games where a correlation device is natural include traffic, congestion control, load balancing (Ashlagi, Monderer, and Tennenholtz 2008), and carbon abatement (Moulin, Ray, and Gupta 2014).

In the context of extensive-form (that is, tree-form) games, two different instantiations of the idea of correlated equilibrium are known in the literature: *normal-form correlated equilibrium (NFCE)* (Aumann 1974; Gilboa and Zemel 1989) and *extensive-form correlated equilibrium (EFCE)* (von Stengel and Forges 2008). The two solution concepts differ in what the mediator reveals to the players. In an NFCE, the mediator privately reveals to each player, just before the game starts, the (whole) normal-form plan that was sampled for the player. Players are then free to either play according to the plan, or play any other strategy that they desire. In an EFCE, the mediator does not reveal the whole plan to the players before the game starts. Instead, the mediator incrementally reveals the plan by recommending individual *moves*. Each recommended move is only revealed when the player reaches the decision point for which the recommendation is relevant. Each player is free to play a move different than the recommended one, but doing so comes at the cost of future recommendations, as the mediator will immediately stop issuing recommendations to any player who did not follow all the recommendations so far. Because of this deterrent, and because players have to decide whether to follow recommendations knowing less about the sampled normal-form plan than in NFCE, a social-welfare-maximizing EFCE always achieves social welfare equal or higher than any NFCE.

Coarse correlated equilibrium differs from correlated equilibrium in that players must decide whether or not to commit to playing according to the recommendations of the mediator *before* observing such recommendations. *Normal-form coarse-correlated equilibrium (NFCCE)* (Moulin and Vial 1978) is the coarse equivalent of NFCE. Before the game starts, players decide whether to commit to playing according to the normal-form plan that was sampled by the mediator (from some correlated distribution known to all players), without observing such a plan first. Players who decide to commit will privately receive the plan that was sampled for them; players that decide to not commit will not receive any recommended plan, and are free to play according to any strategy they desire. Since players know less at the time of commitment than either in NFCE or EFCE, a social-welfare-maximizing NFCCE is always guaranteed to achieve equal or higher social welfare than any NFCE or EFCE. No coarse equivalent of EFCE is known in the literature.

In this paper, we introduce the coarse equivalent of EFCE, which we coin *extensive-form coarse-correlated equilibrium*

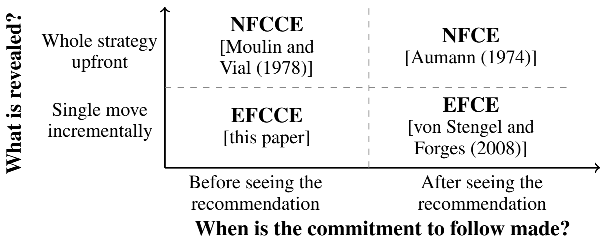

Figure 1: Correlated and coarse-correlated solution concepts.

(EFCCE). It is an intermediate solution concept between EFCE and NFCCE. Specifically, EFCCE is akin to EFCE in that each recommended move is only revealed when the players reach the decision point for which the recommendation is relevant. However, unlike EFCE, the acting player must choose whether or not to commit to the recommended move *before* the move is revealed to them, instead of after. Figure 1 shows how EFCCE fits within the family of correlated and coarse-correlated solution concepts.

We prove that EFCCEs are a subset of NFCCEs and a superset of EFCEs, and give an example of a game in which the three solution concepts lead to distinct solution sets. So—because a social-welfare-maximizing EFCCE guarantees a higher social welfare than any EFCE—our EFCCE solution concept is more appealing than EFCE in applications where the mediator has enough contractual power to enforce that agents that commit to follow the recommended move actually do play the recommended move. This can be the case, for example, if players are able to enter into binding contracts with the mediator or the mediator has enough extraneous power over the players.

We also show that the problem of computing a social-welfare-maximizing EFCCE can be represented as a bilinear saddle-point problem, which can be solved in polynomial time in two-player extensive-form games with no chance moves but not in games with more than two players or two-player games with chance moves. Finally, we note that in two-player games with no chance moves, EFCCE leads to a linear program whose size is smaller than EFCE; because of this, EFCCE can also be used as a computationally lighter relaxation of EFCE (for example, as a routine in the algorithm by Čermák et al.; Bosanský et al. (2016; 2017) for computing a strong Stackelberg equilibrium).

Finally, we show that the problem of computing a social-welfare-maximizing NFCCE can be expressed as a bilinear saddle-point problem, which can be solved in polynomial time in two-player extensive-form games with no chance moves (the problem is known to be NP-hard in games with more than two players and/or chance moves (von Stengel and Forges 2008)). This formulation is significant, as it enables several new classes of algorithms to be employed to find a social-welfare-maximizing NFCCE. In particular, we show that it enables a linear programming formulation that in our experiments is two to four orders of magnitude faster than the prior state of the art.

Preliminaries

Extensive-form games are played on a game tree, and can capture both sequential and simultaneous moves, as well as private information. Each node v in the game tree belongs to exactly one player $i \in \{1, \dots, n\} \cup \{\mathtt{C}\}$ whose turn it is to move. Player $\mathtt{C}$ is a special player called the *chance player*; it is used to denote random events that happen in the game, such as drawing a card from a deck or tossing a coin. The edges leaving v represent actions that the player can take at that node; we denote the set of actions available at v by A_v. In order to capture private information, the set of nodes that belong to each player $i \in \{1, \dots, n\}$ are partitioned into a collection $\mathcal{I}_i$ of nonempty sets. Each $I \in \mathcal{I}_i$ is called an *information set* of Player i, and is a set of nodes that Player i cannot distinguish among, given what the player has observed so far. In this paper, we assume *perfect recall*, that is, no player forgets what he or she knew earlier. Necessarily, for any $I \in \mathcal{I}_i$ and $u, v \in I$, it must be $A_u = A_v$, or otherwise Player i would be able to distinguish between u and v. For this reason, we will often write A_I to mean the set of available actions at any node in I. Finally, two information sets I_i, I_j for Player i and j, respectively, are *connected*, denoted by $I_i \rightleftharpoons I_j$, if there exist $u \in I_i, v \in I_j$ such that the path from the root to u passes through v or vice versa.

Any node v for which A_v is empty is called a *leaf*, and denotes an end state of the game. We denote the set of leaves of the game by Z. Each $z \in Z$ is associated with a tuple of n payoffs (one for each non-chance player); we denote by $u_i(z)$ the payoff for Player $i \in \{1, \dots, n\}$ in leaf z.

Sequences (Σ)

The set of *sequences* of Player i, denoted by Σ_i, is defined as the set $\Sigma_i := \{(I, a) : I \in \mathcal{I}_i, a \in A_I\} \cup \{\varnothing_i\}$, where the special sequence $\varnothing_i$ is called the *empty sequence* of Player i. Given a node v that belongs to Player i, the *parent sequence* of v, denoted by $\sigma_i(v)$, is the last sequence $(I, a) \in \Sigma_i$ encountered on the path from the root of the game tree to v; if no such sequence exists (i.e., Player i never acts before v), we let $\sigma_i(v) = \varnothing_i$. The *parent sequence* $\sigma(I)$ of an information set $I \in \mathcal{I}_i$ is defined as $\sigma(I) := \sigma(v)$ where v is any node in I (all choices produce the same parent sequence, since the game is assumed to have perfect recall).

A pair of sequences is *relevant* if their information sets are on the same branch. Formally, given two sequences σ_i and σ_j for two distinct Players i and j, respectively, we say that the pair (σ_i, σ_j) is relevant if either one sequence is the empty sequence, or $\sigma_i = (I_i, a_i), \sigma_j = (I_j, a_j)$ and $I_i \rightleftharpoons I_j$.

Reduced-Normal-Form Plans (Π)

A *normal-form plan* for Player i defines a choice of action $a_I \in A_I$ for *every* information set $I \in \mathcal{I}_i$ of the player. However, this representation contains irrelevant information, as some information sets may become unreachable after the player makes certain decisions higher up the tree. A *reduced-normal-form plan* π is a normal-form plans where this irrelevant information is removed: it defines a choice of action for *every* information set $I \in \mathcal{I}_i$ that is still reachable as a result of the other choices in π itself. We denote the set of reduced-normal-form plans of Player i by Π_i.

We now define certain subsets of Π_i. The reader is encouraged to refer to Figure 2 while reading the definitions to see what these subsets are in a small example. Given an

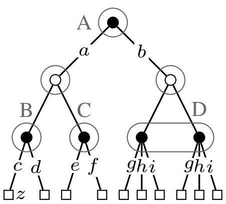

	ABCD
π_1	$a\,c\,e-$
π_2	$a\,c\,f-$
π_3	$a\,d\,e-$
π_4	$a\,d\,f-$
π_5	$b\,{-}{-}\,g$
π_6	$b\,{-}{-}\,h$
π_7	$b\,{-}{-}\,i$

$$\Pi_1(\text{A}) = \{\pi_1, \dots, \pi_7\}$$
$$\Pi_1(\text{B}) = \{\pi_1, \dots, \pi_4\}$$
$$\Pi_1(\text{D}) = \{\pi_5, \pi_6, \pi_7\}$$
$$\Pi_1((\text{C}, f)) = \{\pi_2, \pi_4\}$$
$$\Pi_1((\text{D}, g)) = \{\pi_5\}$$
$$\Pi_1(z) = \{\pi_1, \pi_2\}$$

Figure 2: (Left) Sample game tree. Black round nodes belong to Player 1, white round nodes belong to Player 2, and white square nodes are leaves. Rounded, gray lines denote information sets. (Center) Set Π_1 of reduced-normal-form plans for Player 1. Each plan identifies an action at each information set; information sets that are not reachable given the actions that were chosen higher up in the tree are denoted by '–'. (Right) Examples of certain subsets of Π_1 defined in the Preliminaries section.

information set I of Player i, we denote by $\Pi_i(I)$ the subset of reduced-normal-form plans π where Player i plays so as to reach I whenever possible (the possibility depends on the opponent's actions up to that point), and can play any other actions at points of the game where reaching I is not possible anymore. Given a sequence $\sigma = (I, a) \in \Sigma_i$, $\Pi_i(\sigma)$ further curtails the set $\Pi_i(I)$ by requiring that at I, Player i plays action a. We let $\Pi_i(\varnothing_i) := \Pi_i$. We will also make frequent use of the shorthand $\Pi_i(z) := \Pi_i(\sigma_i(z))$ to denote the set of reduced-normal-form plans in which Player i tries to reach leaf by $z \in Z$. Finally, a *reduced-normal-form strategy* for Player i is a probability distribution over Π_i.

Polytope of Sequence-Form Strategies (Q)

The sequence-form representation is a more compact way of representing normal-form strategies of a player in a perfect-recall extensive-form game (Romanovskii 1962; Koller, Megiddo, and von Stengel 1996; von Stengel 1996). Formally, fix a player $i \in \{1, \dots, n\}$, and let μ be a probability distribution over Π_i. The *sequence-form strategy induced by* μ is the real vector $\boldsymbol{y}$, indexed over $\sigma \in \Sigma_i$, defined as

$$y(\sigma) := \sum_{\pi \in \Pi_i(\sigma)} \mu(\pi). \tag{1}$$

The set of sequence-form strategies that can be induced as μ varies over the set of all possible probability distributions over Π_i is denoted by Q_i. Koller, Megiddo, and von Stengel (1996) prove that it is a convex polytope (called the *sequence-form polytope*) $Q_i = \{\boldsymbol{y} \in \mathbb{R}^{|\Sigma_i|} : \boldsymbol{F}_i\boldsymbol{y} = \boldsymbol{f}_i,\ \boldsymbol{y} \geq \boldsymbol{0}\}$, where $\boldsymbol{F}_i$ is a sparse $|\mathcal{I}_i| \times |\Sigma_i|$ matrix with entries in $\{0, 1, -1\}$, and $\boldsymbol{f}_i$ is a vector with entries in $\{0, 1\}$.

Polytope of Extensive-Form Correlation Plans (Ξ)

Given any probability distribution μ over $\bigtimes_{i=1}^n \Pi_i$ in an extensive-form game, the *correlation plan* $\boldsymbol{\xi}$ induced by μ is defined as the real vector, indexed over tuples $(\sigma_1, \dots, \sigma_n) \in \bigtimes_{i=1}^n \Sigma_i$ of pairwise-relevant sequences, where each entry is

$$\xi(\sigma_1, \dots, \sigma_n) := \sum_{\substack{\pi_1 \in \Pi_1(\sigma_1) \\ \cdots \\ \pi_n \in \Pi_n(\sigma_n)}} \mu(\pi_1, \dots, \pi_n). \tag{2}$$

The set of correlation plans $\boldsymbol{\xi}$ that can be induced as μ varies over the set of all possible probability distributions is denoted by Ξ and called the *polytope of extensive-form correlation plans*. It is always a polytope in a space of dimension polynomial in the input game description. Furthermore, in two-player games without chance moves, Ξ can be described as the intersection of a polynomial number (in the size of the game tree) linear constraints, as shown by von Stengel and Forges (2008). They also proved that this property does not always hold if the game has chance moves or more than two players. Finally, for any $i \in \{1, \dots, n\}$, $\sigma \in \Sigma_i$, and $z \in Z$, we introduce the following notation that we will use frequently: $\xi_i(\sigma; z) := \xi(\sigma_1(z), \dots, \sigma_{i-1}(z), \sigma, \sigma_{i+1}(z), \dots, \sigma_n(z))$.

Saddle-Point Formulation of NFCCE

In this section, we show that the problem of finding an NFCCE in an n-player extensive-form game with perfect recall can be expressed as a bilinear saddle-point problem, that is, an optimization problem of the form

$$\operatorname*{argmin}_{\boldsymbol{x} \in \mathcal{X}} \max_{\boldsymbol{w} \in \mathcal{W}} \boldsymbol{x}^\top \boldsymbol{A} \boldsymbol{w},$$

where $\mathcal{X}$ and $\mathcal{W}$ are convex and compact sets. In our specific case, $\mathcal{X}$ and $\mathcal{W}$ will be convex polytopes in low-dimensional spaces (in particular, $\mathcal{X} = \Xi$). As we will show later, this formulation immediately implies that in two-player games with no chance moves, a social-welfare-maximizing NFCCE can be computed in polynomial time as the solution of a linear program.

We now go through the steps that enable us to formulate the problem of finding an NFCCE as a bilinear saddle-point problem. The general structure of the argument is similar to that of Farina et al. (2019) in the context of EFCE, and we will use it again later when dealing with EFCCE.

By definition, a correlated distribution μ over $\bigtimes_{i=1}^n \Pi_i$ is an NFCCE if no player has an incentive to unilaterally deviate from the recommended plan assuming that nobody else does. More formally, let i be any player, and let $\hat{\mu}_i$ be any probability distribution over Π_i, independent of μ. Playing according to $\hat{\mu}_i$ must give Player i expected utility $\hat{u}_i$ at most equal to the expected utility u_i of committing to the mediator's recommendation. In order to express $\hat{u}_i$ and u_i as a function of μ and $\hat{\mu}_i$, it is necessary to quantify the probability of the game ending in any leaf $z \in Z$. When Player i deviates and plays according to $\hat{\mu}_i$, the probability that the game ends in z is equal to the probability that the mediator samples from μ plans $\pi_j \in \Pi_j(z)$ for any Player j other than i, and that Player i samples from $\hat{\mu}_i$ a plan $\hat{\pi}_i \in \Pi_i(z)$. Correspondingly, using the independence of μ and $\hat{\mu}_i$, we can write

$$\hat{u}_i = \sum_{z \in Z} \left[u_i(z) \left(\sum_{\substack{\pi_i \in \Pi_i \\ \pi_j \in \Pi_j(z)\ \forall j \neq i}} \mu(\pi_1, \dots, \pi_n) \right) \left(\sum_{\hat{\pi}_i \in \Pi_i(z)} \hat{\mu}_i(\hat{\pi}_i) \right) \right].$$

On the other hand, the probability that leaf z is reached when all players commit to the mediator's recommendation is equal to the probability that the mediator samples from μ plans $\pi_j \in \Pi_j(z)$ for all players $j \in \{1, \dots, n\}$:

$$u_i = \sum_{z \in Z} \left[u_i(z) \left(\sum_{\pi_j \in \Pi_j(z)\ \forall j} \mu(\pi_1, \dots, \pi_n) \right) \right]. \tag{3}$$

Using the definition of extensive-form correlation plan (2) and sequence-form strategy (1) we can convert the requirement that $\hat{u}_i \le u_i$ for all choices of i and $\hat{\mu}_i$ into the following equivalent condition:

Proposition 1. *An extensive-form correlation plan $\boldsymbol{\xi} \in \Xi$ is an NFCCE if and only if, for any player $i \in \{1, \dots, n\}$ and sequence-form strategy $\boldsymbol{y}_i \in Q_i$,*

$$\sum_{z \in Z} u_i(z)\xi_i(\varnothing_i; z) y_i(\sigma_i(z)) \le \sum_{z \in Z} u_i(z)\xi_i(\sigma_i(z); z). \quad (4)$$

Inequality (4) is of the form $\boldsymbol{\xi}^\top \boldsymbol{A}_i \boldsymbol{y}_i - \boldsymbol{b}_i^\top \boldsymbol{\xi} \le 0$, where $\boldsymbol{A}_i$ and $\boldsymbol{b}_i$ are suitable sparse matrices/vectors that only depend on i. With this new notation, we can rewrite the condition in Proposition 1 as follows: $\boldsymbol{\xi} \in \Xi$ is an NFCCE if and only if

$$\begin{aligned} & \max_{i \in \{1,\dots,n\}} \max_{\boldsymbol{y}_i \in Q_i} \left\{ \boldsymbol{\xi}^\top \boldsymbol{A}_i \boldsymbol{y}_i - \boldsymbol{b}_i^\top \boldsymbol{\xi} \right\} \le 0 \\ \iff & \max_{\substack{\boldsymbol{\lambda} \in \Delta^n \\ \boldsymbol{y}_i \in Q_i \ \forall i}} \left\{ \sum_{i=1}^n \lambda_i \left(\boldsymbol{\xi}^\top \boldsymbol{A}_i \boldsymbol{y}_i - \boldsymbol{b}_i^\top \boldsymbol{\xi} \right) \right\} \le 0 \\ \iff & \max_{\substack{\boldsymbol{\lambda} \in \Delta^n \\ \tilde{\boldsymbol{y}}_i \in \lambda_i Q_i \ \forall i}} \left\{ \sum_{i=1}^n \boldsymbol{\xi}^\top \boldsymbol{A}_i \tilde{\boldsymbol{y}}_i - \lambda_i \boldsymbol{b}_i^\top \boldsymbol{\xi} \right\} \le 0, \end{aligned} \quad (5)$$

where in the last transformation we operated a change of variable $\tilde{\boldsymbol{y}}_i := \lambda_i \boldsymbol{y}_i$; it is a simple exercise to prove that this change of variable is legitimate and that the domain of the maximization is a convex polytope. An NFCCE always exists. In particular, any

$$\boldsymbol{\xi} \in \operatorname*{argmin}_{\boldsymbol{\xi} \in \Xi} \max_{\substack{\boldsymbol{\lambda} \in \Delta^n \\ \tilde{\boldsymbol{y}}_i \in \lambda_i Q_i \ \forall i}} \left\{ \sum_{i=1}^n \boldsymbol{\xi}^\top \boldsymbol{A}_i \tilde{\boldsymbol{y}}_i - \lambda_i \boldsymbol{b}_i^\top \boldsymbol{\xi} \right\} \quad (6)$$

is an NFCCE because it satisfies (5). Since the domains of the minimization and maximization problems are both convex polytopes, and the objective function is bilinear, the optimization problem in (6) is a bilinear saddle-point problem.

Enforcing a Lower Bound on Social Welfare

Given an NFCCE μ, social welfare is defined as $\mathrm{SW} := \sum_{i=1}^n u_i$, where u_i is as in Equation (3). Hence, social welfare is a linear function of the correlation plan $\boldsymbol{\xi}$, which can be expressed as $\mathrm{SW} : \Xi \ni \boldsymbol{\xi} \mapsto \boldsymbol{c}^\top \boldsymbol{\xi}$ where $\boldsymbol{c} := \sum_{i=1}^n \boldsymbol{b}_i$. Consequently, an NFCCE that guarantees a given lower bound τ on the social welfare can be expressed as in (6) where the domain of the minimization is changed from $\boldsymbol{\xi} \in \Xi$ to $\boldsymbol{\xi} \in \Xi \cap \{\boldsymbol{\xi} : \boldsymbol{c}^\top \boldsymbol{\xi} \ge \tau\}$. This preserves the polyhedral nature of the optimization domain. Finally, the same construction can be used verbatim if social welfare is replaced with any linear function of $\boldsymbol{\xi}$.

Connection to Linear Programming

The saddle-point formulation in (6) can be mechanically translated into a linear program (LP) by taking the dual of the internal maximization problem, that is, of (5). Specifically, the dual problem is the linear program

$$(7): \begin{cases} \min & u \\ \text{s.t.} & u - \boldsymbol{v}_i^\top \boldsymbol{f}_i + \boldsymbol{b}_i^\top \boldsymbol{\xi} \ge 0 \quad \forall i \in \{1, \dots, n\} \\ & \boldsymbol{F}_i^\top \boldsymbol{v}_i - \boldsymbol{A}_i^\top \boldsymbol{\xi} \ge \boldsymbol{0} \quad \forall i \in \{1, \dots, n\} \\ & u \in \mathbb{R}, \boldsymbol{v}_i \in \mathbb{R}^{|\mathcal{I}_i|} \quad \forall i \in \{1, \dots, n\}. \end{cases}$$

(See the Preliminaries section for the meaning of $\boldsymbol{F}_i$ and $\boldsymbol{f}_i$). By strong duality, the value of (7) is the same as the value of the primal problem, that is, the maximum 'deviation benefit' $\hat{u}_i - u_i$ across all players $i \in \{1, \dots, n\}$ and probability distributions $\hat{\mu}_i$ over Π_i. Hence, we can find an NFCCE $\boldsymbol{\xi}$ that maximizes any given objective $\boldsymbol{c}^\top \boldsymbol{\xi}$ by adding the constraint $u \le 0$ and solving the modified LP

$$(8): \begin{cases} \max & \boldsymbol{c}^\top \boldsymbol{\xi} \\ \text{s.t.} & u - \boldsymbol{v}_i^\top \boldsymbol{f}_i + \boldsymbol{b}_i^\top \boldsymbol{\xi} \ge 0 \quad \forall i \in \{1, \dots, n\} \\ & \boldsymbol{F}_i^\top \boldsymbol{v}_i - \boldsymbol{A}_i^\top \boldsymbol{\xi} \ge \boldsymbol{0} \quad \forall i \in \{1, \dots, n\} \\ & \boldsymbol{\xi} \in \Xi \\ & u \le 0, \boldsymbol{v}_i \in \mathbb{R}^{|\mathcal{I}_i|} \quad \forall i \in \{1, \dots, n\}. \end{cases}$$

The LP above always has a polynomial number of variables, but potentially an exponential number of constraints because of the condition $\boldsymbol{\xi} \in \Xi$. However, in two-player extensive-form games with no chance moves, (8) is *guaranteed* to have a polynomial number of constraints as Ξ can be described compactly (von Stengel and Forges 2008). Hence, in those games a social-welfare-maximizing NFCCE can be computed in polynomial time.

EFCCE: An Intermediate Solution Concept

In this section, we introduce a new solution concept which we coin *extensive-form coarse-correlated equilibrium (EFCCE)*. It combines the idea of *coarse correlation*—that is, players must decide whether they want to commit to following the recommendations issued by the correlation device before observing such recommendations—with the idea of *extensive-form correlation*—that is, recommendations are revealed incrementally as the players progress down the game tree. EFCCE is akin to EFCE in that each recommended move is only revealed when the players reach the decision point for which the recommendation is relevant. However, unlike EFCE, the acting player must choose whether or not to commit to the recommended move *before* such a move is revealed to them, instead of after. Each choice is binding only with respect to the decision point for which the choice is made, and players can make different choices at different decision points. Just like EFCE, defections (that is, deciding to *not* commit to following the correlation device's recommended move) come at the cost of future recommendations, as the correlation device will stop issuing recommendations to the defecting player. As with all correlated equilibria, the correlated distribution from which the recommendations are sampled must be such no player has incentives to unilaterally defect when no other player does.

Saddle-Point Formulation

In this section, we show that also an EFCCE can be expressed as the solution to a bilinear saddle-point problem. To do so, we use the idea of *trigger agents* (Gordon, Greenwald, and Marks 2008; Dudik and Gordon 2009):

Definition 1. *Let $i \in \{1, \dots, n\}$ be a player, let $\hat{I} \in \mathcal{I}_i$ be an information set for Player i, and let $\hat{\mu}_i$ be a probability distribution over $\Pi_i(\hat{I})$. An $(\hat{I}, \hat{\mu}_i)$-*trigger agent for Player i *is a player that commits to and follows all recommendations issued by the mediator until they reach a node $v \in \hat{I}$ (if any). When any node $v \in \hat{I}$ is reached, the player 'gets triggered', stops committing to the recommendations and instead plays* *according to a reduced-normal-form plan sampled from $\hat{\mu}_i$ until the game ends.*

By definition, a correlated distribution μ over $\times_{i=1}^{n} \Pi_i$ is an EFCCE when, for all $i \in \{1, \ldots, n\}$, the value u_i that Player i obtains by following the recommendations is at least as large as the expected utility $\hat{u}_{\hat{I}}$ attained by any $(\hat{I}, \hat{\mu}_i)$-trigger agent for that player (assuming nobody else deviates). The expected utility for Player i when everybody commits to following the mediator's recommendations is as in Equation (3). In order to express the expected utility of the $(\hat{I}, \hat{\mu}_i)$-trigger agent, we start by computing the probability of the game ending in each possible leaf $z \in Z$. Let $(\pi_1, \ldots, \pi_n)$ be the tuple of reduced-normal-form plans that was sampled by the mediator. Two cases must be distinguished:

- The path from the root of the game tree to z passes through a node $v \in \hat{I}$. We denote the set of such leaves by $Z_{\hat{I}}$. In this case, the trigger agent commits to following all recommendations until just before $\hat{I}$, and then plays according to a plan $\hat{\pi}_i \in \Pi_i(\hat{I})$ sampled from the distribution $\hat{\mu}_i$ from $\hat{I}$ onwards. The following conditions are necessary and sufficient for the game to terminate at z: $\pi_j \in \Pi_j(z)$ for all $j \in \{1, \ldots, n\} \setminus \{i\}$, $\pi_i \in \Pi_i(\hat{I})$, and $\hat{\pi}_i \in \Pi_i(z)$. Hence, the probability that the game ends at $z \in Z_{\hat{I}}$ is
$$p_z := \left(\sum_{\substack{\pi_i \in \Pi_i(\hat{I}) \\ \pi_j \in \Pi_j(z)\ \forall j \neq i}} \mu(\pi_1, \ldots, \pi_n) \right) \left(\sum_{\hat{\pi}_i \in \Pi_i(z)} \hat{\mu}_i(\hat{\pi}_i) \right). \quad (9)$$
- Otherwise, the trigger agent never gets triggered, and instead commits to following all recommended moves until the end of the game. The probability that the game ends at $z \in Z \setminus Z_{\hat{I}}$ is therefore
$$q_z := \sum_{\pi_j \in \Pi_j(z)\ \forall j} \mu(\pi_1, \ldots, \pi_n). \quad (10)$$

With this information, the expected utility of the $(\hat{I}, \hat{\mu}_i)$-trigger agent is $\hat{u}_{\hat{I}} = \sum_{z \in Z_{\hat{I}}} u_i(z)\, p_z + \sum_{z \in Z \setminus Z_{\hat{I}}} u_i(z)\, q_z$.

Using to the definition of extensive-form correlation plan (2) and sequence-form strategy (1), we can rewrite the condition $u_i \geq \hat{u}_{\hat{I}}$ (which must hold for all choices of i, $\hat{I} \in \mathcal{I}_i$, and probability distribution $\hat{\mu}_i$ over $\Pi_i(\hat{I})$) compactly. In particular, denoting by $\boldsymbol{y}_{i,\hat{I}}$ the sequence-form strategy for the $(\hat{I}, \hat{\mu}_i)$-trigger agent, we can rewrite (9) as $p_z = \xi_i(\sigma_i(\hat{I}); z) y_{i,\hat{I}}(\sigma_i(z))$ and (10) as $q_z = \xi_i(\sigma_i(z); z)$. Furthermore, the fact that $\hat{\mu}_i$ has support $\Pi_i(\hat{I})$ translates into the constraints $y_{i,\hat{I}}(\sigma(\hat{I})) = 1$. For this reason, it makes sense to introduce the symbol
$$Q_{i,\hat{I}} := Q_i \cap \{\boldsymbol{y}_{i,\hat{I}} : y_{i,\hat{I}}(\sigma(\hat{I})) = 1\}.$$
Correspondingly, we let $\boldsymbol{F}_{i,\hat{I}}$ and $\boldsymbol{f}_{i,\hat{I}}$ be the constraint matrix and vector such that $Q_{i,\hat{I}} = \{\boldsymbol{y}_{i,\hat{I}} \geq \boldsymbol{0} : \boldsymbol{F}_{i,\hat{I}} \boldsymbol{y}_{i,\hat{I}} = \boldsymbol{f}_{i,\hat{I}}\}$. Putting everything together, we have the following:

Proposition 2. *An extensive-form correlation plan $\boldsymbol{\xi} \in \Xi$ is an EFCCE iff, for any player $i \in \{1, \ldots, n\}$, information set $\hat{I} \in \mathcal{I}_i$, and sequence-form strategy $\boldsymbol{y}_{i,\hat{I}} \in Q_{i,\hat{I}}$,*
$$\sum_{z \in Z_{\hat{I}}} u_i(z) \xi_i(\sigma_i(\hat{I}); z) y_{i,\hat{I}}(\sigma_i(z)) \leq \sum_{z \in Z_{\hat{I}}} u_i(z) \xi_i(\sigma_i(z); z). \quad (11)$$

Inequality (11) is in the form $\boldsymbol{\xi}^\top \boldsymbol{A}_{i,\hat{I}} \boldsymbol{y}_{i,\hat{I}} - \boldsymbol{b}_{i,\hat{I}}^\top \boldsymbol{\xi} \leq 0$ where $\boldsymbol{A}_{i,\hat{I}}$ and $\boldsymbol{b}_{i,\hat{I}}$ are suitable matrices/vectors that only depend on the trigger information set $\hat{I}$ of Player i. From here, one can follow the same steps that we already took in the case of NFCCE and obtain a bilinear saddle-point formulation and an LP for EFCCE. For space reasons, we only state the LP, which we also implemented (see Experiments section):
$$(12): \begin{cases} \max & \boldsymbol{c}^\top \boldsymbol{\xi} \\ \text{s.t.} & u - \boldsymbol{v}_{i,\hat{I}}^\top \boldsymbol{f}_{i,\hat{I}} + \boldsymbol{b}_{i,\hat{I}}^\top \boldsymbol{\xi} \geq 0 \quad \forall i, \hat{I} \in \mathcal{I}_i \\ & \boldsymbol{F}_{i,\hat{I}}^\top \boldsymbol{v}_{i,\hat{I}} - \boldsymbol{A}_{i,\hat{I}}^\top \boldsymbol{\xi} \geq \boldsymbol{0} \quad \forall i, \hat{I} \in \mathcal{I}_i \\ & \boldsymbol{\xi} \in \Xi \\ & u \leq 0, \boldsymbol{v}_{i,\hat{I}} \in \mathbb{R}^{|\mathcal{I}_i|} \quad \forall i, \hat{I} \in \mathcal{I}_i. \end{cases}$$

The LP (12) has a polynomial number of variables, and in two-player games with no chance moves it also has a polynomial number of constraints due to the polynomial description of Ξ (von Stengel and Forges 2008). In particular, in two-player games with no chance moves, a social-welfare-maximizing EFCCE can be computed in polynomial time by setting the objective function $\boldsymbol{c}^\top \boldsymbol{\xi}$ to be the social welfare
$$\boldsymbol{c}^\top \boldsymbol{\xi} := \sum_{z \in Z} \left[\left(\sum_{i=1}^{n} u_i(z) \right) \xi(\sigma_1(z), \ldots, \sigma_n(z)) \right].$$

Finally, the EFCCE linear program (12) has more constraints and variables than that for NFCCE (see (8)), but fewer than that for EFCE (see Supplemental Material). Empirically, this results in intermediate run times compared to NFCCE and EFCE, as confirmed by our experiments presented later.

Complexity Results

As we have already pointed out, in the case of two-player games without chance moves, the LP (12) has a polynomial number of constraints and variables, and can therefore be solved in polynomial time using standard LP technology. As in NFCCE and EFCE, the same does not hold for two-players games with chance moves nor for games with three players or more players, with or without chance moves. In particular, the following results can be easily obtained by using the same reduction employed by von Stengel and Forges (2008) (all proofs are in the Supplemental Material):

Definition 2 ($\text{SW}_{\text{EFCCE}}(\kappa)$)**.** *Given an extensive-form game Γ and a real number κ, $SW_{EFCCE}(\kappa)$ denotes the problem of deciding whether or not Γ admits an EFCCE with social welfare at least κ.*

Proposition 3. *$SW_{EFCCE}(\kappa)$ is NP-Hard in two-player games with chance moves, as well as in three-player games, with or without chance moves.*

Relationships among Solution Concepts

In this section, we analyze the relationship between EFCE, EFCCE, and NFCCE. We start with an inclusion lemma, which shows that the solution concept that we just introduced, EFCCE, is a superset of EFCE and a subset of NFCCE (all proofs are available in the Supplemental Material):

Proposition 4. *Let Γ be a perfect-recall extensive-form game. Then we have the following inclusion of equilibria*
$$\text{EFCE} \subseteq \text{EFCCE} \subseteq \text{NFCCE}.$$

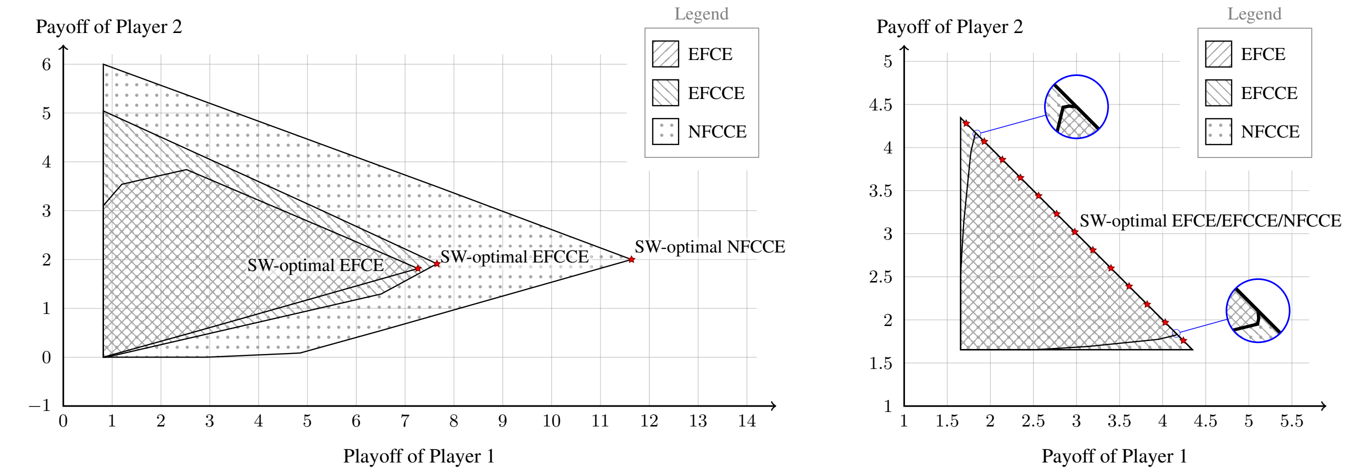

Figure 3: Space of payoff vectors that can be induced by EFCE, EFCCE, and NFCCE in an instance of the Sheriff game (left) and of 3-card Goofspiel (right). The 'SW-optimal' symbols indicate payoffs corresponding to social-welfare-maximizing equilibria.

Proposition 4 applies to games with more than two players and/or chance moves as well. Let $U_{\mathrm{EFCE}}, U_{\mathrm{EFCCE}}, U_{\mathrm{NFCCE}}$ denote the set of expected payoff vectors that can be induced by EFCE, EFCCE, and NFCCE, respectively. Then, Proposition 4 is an important ingredient for our next proposition:

Proposition 5. *The sets $U_{EFCE}, U_{EFCCE}, U_{NFCCE}$ are convex polytopes. Furthermore, $U_{EFCE} \subseteq U_{EFCCE} \subseteq U_{NFCCE}$.*

Proposition 4 also implies the following relationship between the maximum social welfare that can be obtained by EFCE, EFCCE, and NFCCE:

Corollary 1. *Let $SW^*_{EFCE}, SW^*_{EFCCE}, SW^*_{NFCCE}$ denote the maximum social welfare that can be reached by EFCE, EFCCE, and NFCCE. Then, $SW^*_{EFCE} \le SW^*_{EFCCE} \le SW^*_{NFCCE}$.*

Figure 3 shows the set of payoff vectors that can be induced by EFCE, EFCCE, and NFCCE in an instance of the Sheriff game (Farina et al. 2019) (left) and an instance of a 3-card Goofspiel game (Ross 1971) (right).[1] In the Sheriff game instance, we have that both inclusions in Proposition 5 are strict, while in the Goofspiel game the inclusion $U_{\mathrm{EFCE}} \subsetneq U_{\mathrm{EFCCE}}$ is strict and $U_{\mathrm{EFCCE}} = U_{\mathrm{NFCCE}}$. The appendix contains an instance of a Battleship game (Farina et al. 2019) in which the inclusion $U_{\mathrm{EFCCE}} \subsetneq U_{\mathrm{NFCCE}}$ is strict while $U_{\mathrm{EFCE}} = U_{\mathrm{EFCCE}}$.

Experiments

We experimentally compared NFCCE, EFCCE, and EFCE in terms of maximum social welfare and run time. In our experiments, we used instances from three different two-player games with no chance moves: Sheriff (Farina et al. 2019), Battleship (Farina et al. 2019), and Goofspiel (Ross 1971). Sheriff is a bargaining game, in which two players—the Smuggler and the Sheriff—must settle on an appropriate bribe so as to avoid the Sheriff inspecting the Smuggler's cargo, which might or might not contain illegal items. The correlation device recommends to the Smuggler how many illegal items to include in the cargo and what size bribes to offer in the bargaining rounds, and to the Sheriff what feedback to give in the bargaining rounds and whether to inspect the cargo. Battleship is a parametric version of the classic board game, where two competing fleets take turns at shooting at each other. The correlation device recommends to each player where to place the ship and where to shoot in every round. Finally, Goofspiel is a card game in which two players repeatedly bid to win a common public card, which will be discarded in case of bidding ties. The correlation device recommends the bids. The three games were chosen so as to illustrate three different applications in which an intermediate form of centralized control (the correlation device) is beneficial: bargaining in Sheriff, conflict resolution in Battleship, and bidding in Goofspiel. See the Supplemental Material for a detailed description of the games.

We used Gurobi 8.1.1 (Gurobi Optimization 2019) to solve the linear programs for NFCCE (8), EFCCE (12), and EFCE (16).[2] We used the barrier algorithm without crossover, and we let Gurobi automatically determine the recommended number of threads for execution. All experiments were run on a 64-core machine with 512 GB of RAM.

Our experimental results are in Table 1 (top) for Battleship, Table 1 (bottom) for Sheriff, and Table 2 in the Supplemental Material for Goofspiel. Each table is split into three parts. **Part (a)** contains information about the parameters that were used to generate the game instances (see the Supplemental Material for a description of their effects). It also shows the size of the instances in terms of number of sequences pairs, defined as the product $|\Sigma_1| \times |\Sigma_2|$ of the number of sequences of the players, and number of *relevant* pairs of sequences (see Preliminaries section). **Part (b)** compares the run times of our algorithm (column 'LP'). In the case of NFCCE, we

[1] The polytopes of reachable payoff vectors were computed using Polymake, a tool for computational polyhedral geometry (Gawrilow and Joswig 2000; Assarf et al. 2017).

[2] This third LP, and its derivation, are given in the Supplemental Material. This third LP was already known (Farina et al. 2019), but we present it following the notation and structural derivation that we use above for the two new LPs.

Battleship: Grid Size	Rounds	Seq. Pairs $\lvert\Sigma_1\rvert \times \lvert\Sigma_2\rvert$	Rel. Seq. Pairs	NFCCE LP	NFCCE CG-LP	NFCCE CG-MILP	EFCCE LP	EFCE LP	NFCCE	EFCCE	EFCE
$(2,2)$	2	6.795×10^5	3.524×10^4	358ms	30m 08s	24.78s	512ms	692ms	0	-2.100	-2.100
$(3,2)$	2	3.571×10^6	2.645×10^5	6.68s	> 24h	3h 17m	2.67s	3.18s	0	-1.270	-1.270
$(3,2)$	3	7.209×10^8	3.893×10^6	1m 55s	> 24h	> 24h	3m 53s	12m 43s	0	-1.500	-1.500
$(3,2)$	4	4.438×10^{10}	2.644×10^7	25m 49s	oom	oom	1h 03m	2h 57m	0	-1.955	-1.955

Sheriff: $n_{\max}$	$b_{\max}$	r	Seq. Pairs $\lvert\Sigma_1\rvert \times \lvert\Sigma_2\rvert$	Rel. Seq. Pairs	NFCCE LP	NFCCE CG-LP	NFCCE CG-MILP	EFCCE LP	EFCE LP	NFCCE	EFCCE	EFCE
10	2	2	1.045×10^4	3.717×10^3	105ms	6m 17s	42.63s	250ms	311ms	13.64	9.565	9.078
10	2	3	3.706×10^5	3.363×10^4	1.03s	12h 52m	> 24h	1.38s	2.14s	13.64	10	10
10	3	3	1.886×10^6	9.577×10^4	2.07s	> 24h	> 24h	2.44s	9.62s	18.18	15	15
20	5	4	5.373×10^9	1.326×10^7	13m 20s	oom	oom	1h 34m	6h 26m	28.57	25	25
20	6	4	1.797×10^{10}	2.774×10^7	2h 16m	oom	oom	5h 06m	21h 40m	33.33	30	30
20	7	4	5.131×10^{10}	5.277×10^7	6h 32m	oom	oom	11h 59m	oom	38.10	35	—

(a) Instance parameters and dimensions (b) Run times of algorithms (c) Maximum social welfare

Table 1: Experimental results on instances of the Battleship game (top) and Sheriff game (bottom). 'oom': Out of memory.

also compare against the only known prior polynomial-time algorithms to compute social-welfare-maximizing NFCCE in extensive-form games, which are both based on the column generation technique, and were introduced by Celli, Coniglio, and Gatti (2019). We implemented both algorithms proposed by Celli, Coniglio, and Gatti (2019): 'CG-LP', based on a linear programming oracle, and 'CG-MILP', based on a mixed integer linear programming oracle. CG-LP is guaranteed to compute a social-welfare-maximizing NFCCE in polynomial time, whereas CG-MILP requires exponential time in the worst case but was faster in practice in that prior paper. These algorithms were implemented in AMPL and rely on the Gurobi backend. Finally, **Part (c)** reports the value of the maximum social welfare that can be attained by NFCCE, EFCCE, and EFCE.

Comparison of Run Time

As expected, increasing the coarseness of the equilibrium from EFCE to EFCCE to NFCCE reduces the linear program size and results in a shorter run time. Empirically, the NFCCE LP is up to four times faster than the EFCCE LP, and the EFCCE LP is in turn between two to four times faster than the EFCE LP. Furthermore, our results indicate that the NFCCE LP that we developed (8) is two to four orders of magnitude faster than CG-LP and CG-MILP, and that it is able to scale to game instances up to five orders of magnitude larger than CG-LP and CG-MILP can in 24 hours. This difference in performance is likely due, at least in part, to the fact that the algorithms by Celli, Coniglio, and Gatti (2019) have a number of constraints that scales with the total number $|\Sigma_1| \times |\Sigma_2|$ of sequence pairs in the game, whereas our LP formulation has a number of constraints and variables that grows with the number of *relevant* sequence pairs, which is a tiny fraction of the total number of sequence pairs in practice.

Comparison of Maximum Social Welfare

Our results experimentally confirm Corollary 1: as the coarseness of the equilibrium increases from EFCE to EFCCE to NFCCE, so does the value of the maximum social welfare that the mediator can induce. The maximum social welfare attained by NFCCE is strictly larger than EFCCE and EFCE in Battleship and Sheriff (Table 1), while it is the same in Goofspiel (Table 2 in the Supplemental Material).

Experimentally, the maximum social welfare that can be obtained through EFCCE is often equal to the maximum social welfare that can be obtained through EFCE. While this does not imply that the set of reachable payoffs is the same (see Figure 3), it is an indication that EFCCE is a fairly tight relaxation of EFCE. That, combined with the fact that it can be solved up to four times faster than EFCE in practice, suggest that this new solution concept is worthwhile.

Conclusions

In this paper we studied two instantiations of the idea of coarse correlation in extensive-form games: normal-form coarse-correlated equilibrium and extensive-form coarse-correlated equilibrium. For both solution concepts, we gave saddle-point problem formulations and linear programs.

We proved that EFCCE, which we introduced for the first time, is an intermediate solution concept between NFCCE and the extensive-form correlated equilibrium introduced by von Stengel and Forges (2008). In particular, the set of payoffs that can be reached by EFCCE is always a superset of those that can be reached by EFCE, and a subset of those that can be reached by NFCCE. Empirically, EFCCE is a fairly tight relaxation of EFCE, and a social-welfare-maximizing EFCCE can be computed up to four times faster. This suggests that EFCCE is a worthy solution concept. Also, it can be a suitable and faster alternative in algorithms that rely on EFCE as a subroutine, such as the algorithm by Černỳ, Boŷanskỳ, and Kiekintveld (2018) for computing a strong Stackelberg equilibrium.

Finally, we compared the run time of our algorithm for computing a social-welfare-maximizing NFCCE, and showed that it is two to four orders of magnitude faster than the only previously known algorithms for that problem. Our algorithm can also scale to game instances up to five orders of magnitude larger than the prior state of the art, thus enabling the computation of coarse-correlated solution concepts in reasonably-sized extensive-form games for the first time.

References

Ashlagi, I.; Monderer, D.; and Tennenholtz, M. 2008. On the value of correlation. *Journal of Artificial Intelligence Research* 33:575–613.

Assarf, B.; Gawrilow, E.; Herr, K.; Joswig, M.; Lorenz, B.; Paffenholz, A.; and Rehn, T. 2017. Computing convex hulls and counting integer points with polymake. *Mathematical Programming Computation* 9(1):1–38.

Aumann, R. 1974. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics* 1:67–96.

Bosanský, B.; Brânzei, S.; Hansen, K. A.; Lund, T. B.; and Miltersen, P. B. 2017. Computation of stackelberg equilibria of finite sequential games. *ACM Transaction on Economics and Computation (TEAC)* 5(4):23:1–23:24.

Celli, A.; Coniglio, S.; and Gatti, N. 2019. Computing optimal ex ante correlated equilibria in two-player sequential games. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, 909–917. International Foundation for Autonomous Agents and Multiagent Systems.

Čermák, J.; Bošanskỳ, B.; Durkota, K.; Lisý, V.; and Kiekintveld, C. 2016. Using correlated strategies for computing Stackelberg equilibria in extensive-form games. In *AAAI*.

Černỳ, J.; Boỳanskỳ, B.; and Kiekintveld, C. 2018. Incremental strategy generation for stackelberg equilibria in extensive-form games. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, 151–168. ACM.

Dudik, M., and Gordon, G. J. 2009. A sampling-based approach to computing equilibria in succinct extensive-form games. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, 151–160. AUAI Press.

Farina, G.; Ling, C. K.; Fang, F.; and Sandholm, T. 2019. Correlation in extensive-form games: Saddle-point formulation and benchmarks. ArXiV preprint.

Gawrilow, E., and Joswig, M. 2000. *Polymake: a Framework for Analyzing Convex Polytopes*. Basel: Birkhäuser Basel. 43–73.

Gilboa, I., and Zemel, E. 1989. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior* 1:80–93.

Gordon, G. J.; Greenwald, A.; and Marks, C. 2008. No-regret learning in convex games. In *Proceedings of the 25th international conference on Machine learning*, 360–367. ACM.

Gurobi Optimization, L. 2019. Gurobi optimizer reference manual.

Koller, D.; Megiddo, N.; and von Stengel, B. 1996. Efficient computation of equilibria for extensive two-person games. *Games and Economic Behavior* 14(2).

Moulin, H., and Vial, J.-P. 1978. Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory* 7(3-4):201–221.

Moulin, H.; Ray, I.; and Gupta, S. S. 2014. Coarse correlated equilibria in an abatement game. Technical report, Cardiff Economics Working Papers.

Romanovskii, I. 1962. Reduction of a game with complete memory to a matrix game. *Soviet Mathematics* 3.

Ross, S. M. 1971. Goofspiel–the game of pure strategy. *Journal of Applied Probability* 8(3):621–625.

von Stengel, B., and Forges, F. 2008. Extensive-form correlated equilibrium: Definition and computational complexity. *Mathematics of Operations Research* 33(4):1002–1022.

von Stengel, B. 1996. Efficient computation of behavior strategies. *Games and Economic Behavior* 14(2):220–246.

Supplemental Material

Formulation of EFCE

In this section, we show that an EFCE can also be expressed as the solution to a bilinear saddle-point problem. To do so, we resort again to the idea of *trigger agents* (Gordon, Greenwald, and Marks 2008; Dudik and Gordon 2009), slightly modifying the definition of trigger agent that we have given in the section about EFCCE to allow for deviations to happen *after* the recommendations have been received:

Definition 3. *Let $i \in \{1, \ldots, n\}$ be a player, let $\hat\sigma = (\hat I, \hat a) \in \Sigma_i$ be a sequence of Player i, and let $\hat\mu_i$ be a probability distribution over $\Pi_i(\hat I)$. A $(\hat\sigma, \hat\mu_i)$*-trigger agent for Player *i is a player that commits to and follows all recommendations issued by the mediator unless they get recommended to play action $\hat a$ at information set $\hat I$; if this happens, the player 'gets triggered', stops following the recommendations and instead plays according to a reduced-normal-form plan sampled from $\hat\mu_i$ until the game ends.*

By definition, a correlated distribution μ over $\bigtimes_{i=1}^n \Pi_i$ is an EFCE when, for all $i \in \{1, \ldots, n\}$, the value u_i that Player i obtains by following the recommendations is at least as large as the expected utility $\hat u_{\hat\sigma}$ attained by any $(\hat\sigma, \hat\mu_i)$-trigger agent for that player (assuming nobody else deviates). The expected utility for Player i when everybody follows the mediator's recommendations is as in Equation (3). On the other hand, the expected utility for the $(\hat\sigma, \hat\mu_i)$-trigger agent can be computed similarly to the one for the $(\hat I, \hat\mu_i)$-trigger agent that we have computed in the section about EFCCE. We thus compute the probability of the game ending in any terminal node $z \in Z$, distinguishing three cases:

- The path from the root of the game tree to z includes playing action $\hat a$ at information set $\hat I$. We denote the set of such leaves by $Z_{\hat\sigma}$. In this case, the trigger agent follows all recommendations until $\hat a$ gets recommended, and then plays according to a reduced-normal-form plan $\hat\pi \in \Pi_i(\hat I)$ sampled from the distribution $\hat\mu_i$ from $\hat I$ onwards. To reach leaf z, however, we need to have that the reduced-normal-form plan $\hat\pi$ includes playing $\hat a$ at $\hat I$. Hence, the following conditions are necessary and sufficient for the game to terminate at z: $\pi_j \in \Pi_j(z)$ for all $j \in \{1, \ldots, n\} \setminus \{i\}$, $\pi_i \in \Pi_i(\hat\sigma)$, and $\hat\pi \in \Pi_i(z)$. Correspondingly, the probability that the game ends at $z \in Z_{\hat\sigma}$ is
$$p_z := \left(\sum_{\substack{\pi_i \in \Pi_i(\hat\sigma) \\ \pi_j \in \Pi_j(z)\ \forall j \neq i}} \mu(\pi_1, \ldots, \pi_n) \right) \left(\sum_{\pi_i \in \Pi_i(z)} \hat\mu_i(\pi_i) \right). \quad (13)$$
- The path from the root of the game tree to z passes through information set $\hat I$ but does not include playing action $\hat a$ at $\hat I$; that is, $z \in Z_{\hat\sigma} \setminus Z_{\hat I}$. The game can end in this state both if the mediator recommends all the players to play in order to reach z, and thus the trigger agent never gets triggered, or if the mediator recommends to play $\hat a$ at $\hat I$ but then the trigger agent deviates and plays according to the reduced-normal-form plan $\hat\pi \in \Pi_i(z)$ sampled from the distribution $\hat\mu_i$. Hence, the probability that the game ends at $z \in Z_{\hat I} \setminus Z_{\hat\sigma}$ is the sum of two terms:
$$q_z := \sum_{\pi_j \in \Pi_j(z)\ \forall j} \mu(\pi_1, \ldots, \pi_n) + \left(\sum_{\substack{\pi_i \in \Pi_i(\hat\sigma) \\ \pi_j \in \Pi_j(z)\ \forall j \neq i}} \mu(\pi_1, \ldots, \pi_n) \right) \left(\sum_{\pi_i \in \Pi_i(z)} \hat\mu_i(\pi_i) \right).$$
- Otherwise, the trigger agent never gets triggered, and follows all recommended moves until the end of the game. The probability that the game ends at $z \in Z \setminus Z_{\hat I}$ is therefore
$$r_z := \sum_{\pi_j \in \Pi_j(z)\ \forall j} \mu(\pi_1, \ldots, \pi_n). \quad (14)$$

With this information, the expected utility of the $(\hat\sigma, \hat\mu_i)$-trigger agent is computed as
$$\hat u_{\hat\sigma} = \sum_{z \in Z_{\hat\sigma}} u_i(z)\, p_z + \sum_{z \in Z_{\hat I} \setminus Z_{\hat\sigma}} u_i(z)\, q_z + \sum_{z \in Z \setminus Z_{\hat I}} u_i(z)\, r_z.$$

By following the same steps taken for NFCCE and EFCCE in the body of the paper, we can now write the constraint that defines an EFCE in the following, compact way:
$$\sum_{z \in Z_{\hat I}} u_i(z) \xi_i(\hat\sigma; z) y_{i,\hat\sigma}(\sigma_i(z)) \le \sum_{z \in Z_{\hat\sigma}} u_i(z) \xi_i(\sigma_i(z); z), \quad (15)$$

which needs to hold for all players $i \in \{1, \ldots, n\}$, sequences $\hat\sigma \in \Sigma_i$, and deviation strategies
$$\begin{aligned} \boldsymbol{y}_{i,\hat\sigma} \in Q_{i,\hat I} &:= \{\boldsymbol{y}_{i,\hat I} \in Q_i : y_{i,\hat I}(\sigma_i(\hat I)) = 1\} \\ &=: \{\boldsymbol{y}_{i,\hat I} : \boldsymbol{F}_{i,\hat I} \boldsymbol{y}_{i,\hat I} = \boldsymbol{f}_{i,\hat I},\ \boldsymbol{y}_{i,\hat I} \ge \boldsymbol{0}\}. \end{aligned}$$

Inequality (15) is in the form $\boldsymbol{\xi}^\top \boldsymbol{A}_{i,\hat\sigma} \boldsymbol{y}_{i,\hat\sigma} - \boldsymbol{b}_{i,\hat\sigma}^\top \boldsymbol{\xi} \le 0$ where $\boldsymbol{A}_{i,\hat\sigma}$ and $\boldsymbol{b}_{i,\hat\sigma}$ are suitable matrices/vectors that only depend on the trigger sequence $\hat\sigma$ of Player i. From this formulation, we can obtain a linear program for finding an EFCE that optimizes over any linear function of ξ, following the steps that we have outlined in the section about linear programming for NFCCE. Here we give only the final LP.

$$(16): \begin{cases} \max & \boldsymbol{c}^\top \boldsymbol{\xi} \\ \text{s.t.} & u - \boldsymbol{v}_{i,\hat\sigma}^\top \boldsymbol{f}_{i,\hat I} + \boldsymbol{b}_{i,\hat\sigma}^\top \boldsymbol{\xi} \ge 0 \quad \forall i, \hat\sigma \in \Sigma_i \\ & \boldsymbol{F}_{i,\hat I}^\top \boldsymbol{v}_{i,\hat\sigma} - \boldsymbol{A}_{i,\hat\sigma}^\top \boldsymbol{\xi} \ge \boldsymbol{0} \quad \forall i,\ \hat\sigma \in \Sigma_i \\ & \boldsymbol{\xi} \in \Xi \\ & u \le 0, \boldsymbol{v}_{i,\,\hat\sigma} \in \mathbb{R}^{|\Sigma_i|} \quad \forall i,\ \hat\sigma \in \Sigma_i. \end{cases}$$

Proofs

Proposition 4. *Let Γ be a perfect-recall extensive-form game. Then we have the following inclusion of equilibria*

$$\text{EFCE} \subseteq \text{EFCCE} \subseteq \text{NFCCE}.$$

Proof. We break the proof into two parts, which can be read independently. We will make use of the notation $\sigma \in I$ to mean that $\sigma = (I, a)$ for some $a \in A_I$.

EFCE $\subseteq$ EFCCE Let $\boldsymbol{\xi} \in \Xi$ be an EFCE. We need to show that, given any player i, decision point $\hat{I} \in \mathcal{I}_i$, and extensive-form strategy $\boldsymbol{y}_{\hat{I}} \in Q_{i,\hat{I}}$, Inequality (11) is satisfied. The crucial ingredient in the proof is the fact that for any $I \in \mathcal{I}_i$ and $z \in Z$,

$$\xi_i(\sigma_i(I); z) = \sum_{\sigma_i \in I} \xi_i(\sigma_i; z)$$

which is a trivial consequence of the definition of Ξ (2) and the fact that for all $I \in \mathcal{I}_i$,

$$\Pi_i(I) = \dot{\bigcup}_{\sigma \in I} \Pi_i(\sigma),$$

where $\dot{\cup}$ denotes disjoint union, that is, union operation that requires that the sets to be unioned are disjoint, which is the case in our setting. Hence,

$$\begin{aligned}
&\sum_{z \in Z_{\hat{I}}} u_i(z) \xi_(\sigma_i(\hat{I}), z) y_{\hat{I}}(\sigma_i(z)) \\
&= \sum_{z \in Z_{\hat{I}}} \sum_{\hat{\sigma} \in \hat{I}} u_i(z) \xi_i(\hat{\sigma}; z) y_{\hat{I}}(\sigma_i(z)) \\
&= \sum_{\hat{\sigma} \in \hat{I}} \left[\sum_{z \in Z_{\hat{I}}} u_i(z) \xi_i(\hat{\sigma}; z) y_{\hat{I}}(\sigma_i(z)) \right]. \qquad (\star)
\end{aligned}$$

Since $\boldsymbol{\xi}$ is an EFCE and $y_{\hat{I}}(\sigma_i(\hat{I})) = 1$ by definition of $Q_{i,\hat{I}}$, the quantity in square brackets in $(\star)$ is upper bounded as in Inequality (15). Hence we can write

$$\begin{aligned}
(\star) &\leq \sum_{\hat{\sigma} \in \hat{I}} \sum_{z \in Z_{\hat{\sigma}}} u_i(z) \xi_i(\sigma_i(z); z) \\
&= \sum_{z \in Z_{\hat{I}}} u_i(z) \xi_i(\sigma_i(z); z),
\end{aligned}$$

where the second equality follows from the fact that the collection of sets $\{Z_{\hat{\sigma}}\}_{\hat{\sigma} \in \hat{I}}$ forms a partition of $Z_{\hat{I}}$. This shows that Inequality (11) holds for any applicable choice of player i, decision point $\hat{I}$ and deviation strategy $\boldsymbol{y}_{\hat{I}}$, and therefore $\boldsymbol{\xi}$ is an EFCCE.

EFCCE $\subseteq$ NFCCE Let $\boldsymbol{\xi} \in \Xi$ be an EFCCE. We need to show that, for any player i and extensive-form deviation strategy $y_i \in Q_i$, Inequality (4) is satisfied. To this end, let $\mathcal{I}_i^*$ be the set of *initial* decision points for the given player, defined as all information sets that has an empty parent sequence; in symbols, $\mathcal{I}_i^* := \{I \in \mathcal{I}_i : \sigma_i(I) = \varnothing_i\}$. The collection of sets $\{Z_I\}_{I \in \mathcal{I}_i^*}$ is a partition of the set of all leaves Z. Hence, using the fact that $\sigma_i(I) = \varnothing_i$ for all $I \in \mathcal{I}_I^*$:

$$\begin{aligned}
&\sum_{z \in Z} u_i(z) \xi_i(\varnothing_i; \sigma_i(z)) y_i(\sigma_i(z)) \\
&= \sum_{I \in \mathcal{I}_i^*} \sum_{z \in Z_I} u_i(z) \xi_i(\varnothing_i; z) y_i(\sigma_i(z)) \\
&= \sum_{I \in \mathcal{I}_i^*} \sum_{z \in Z_I} u_i(z) \xi_i(\sigma_i(I); z) y_i(\sigma_i(z)) \\
&\leq \sum_{I \in \mathcal{I}_i^*} \sum_{z \in Z_I} u_i(z) \xi_i(\sigma_i(z); z) \\
&= \sum_{z \in Z} u_i(z) \xi_i(\sigma_i(z); z),
\end{aligned}$$

where the inequality follows from Inequality (11), which is applicable since $\boldsymbol{\xi}$ is an EFCCE by hypothesis, and $y(\sigma_i(I)) = y(\varnothing_i) = 1$ by definition of Q_i. This shows that Inequality (4) holds for any player i, and therefore $\boldsymbol{\xi}$ is an NFCCE. $\square$

Proposition 5. *The sets $U_{EFCE}, U_{EFCCE}, U_{NFCCE}$ are convex polytopes. Furthermore, $U_{EFCE} \subseteq U_{EFCCE} \subseteq U_{NFCCE}$.*

Proof. First, observe that the set of NFCCEs is a convex polytope, since it is the intersection of Ξ with the linear inequalities 4, one for each player i. (Equivalent statements hold for EFCCE and EFCE.) Second, the function that maps a $\boldsymbol{\xi} \in \Xi$ to the tuple of expected payoffs (one for each player) under the assumptions that players do not deviate from the recommendation strategy encoded by $\boldsymbol{\xi}$, namely

$$\boldsymbol{\xi} \mapsto \begin{pmatrix} \vdots \\ \sum_{z \in Z} u_i(z) \xi_i(\sigma_i(z); z) \\ \vdots \end{pmatrix}_{i \in \{1,\dots,n\}},$$

is linear. Since the image of a convex polytope with respect to a linear function is a convex polytope, the first part of the statement follows. The second part of the statement follows trivially from Proposition 4. $\square$

Proposition 6. *$SW_{EFCCE}(\kappa)$ is NP-Hard in two-player games with chance moves.*

Proof. Given a SAT instance (C, V) in disjunctive normal form, where C is the set of clauses and V the set of variables, we can build an extensive-form game as follows (see Figure 4 for an example of such a game tree):

- At the start of the game, chance selects one action in the set $\{a_\phi \mid \phi \in C\}$ uniformly at random, that is, it picks non-deterministically a clause in the SAT formula.
- Each action a_ϕ leads to the single node in information set I_ϕ of Player 1, where he can choose an action to play in the set $\{a_{\phi,l} \mid l \in \phi\}$, that is, it selects a literal in the clause ϕ.
- All actions $a_{\phi,l}$ with $l = v$ or $l = \bar{v}$ for some $v \in V$ lead to a node in information set I_v of Player 2, where he can choose an action to play in the set $\{a_v, a_{\bar{v}}\}$, that is, it selects a truth assignment for variable v.
- Actions $\{a_v, a_{\bar{v}}\}$ lead to terminal nodes, where the players receive utility $(0, 0)$ if $l = v$ and action $a_{\bar{v}}$ was played or if $l = \bar{v}$ and action a_v was played, and they receive utility $(1, 1)$ otherwise.

If the SAT formula is satisfiable, this game admits a pure Nash equilibrium in which Player 2 plays the truth assignment that satisfies it and Player 1 selects a satisfiable literal

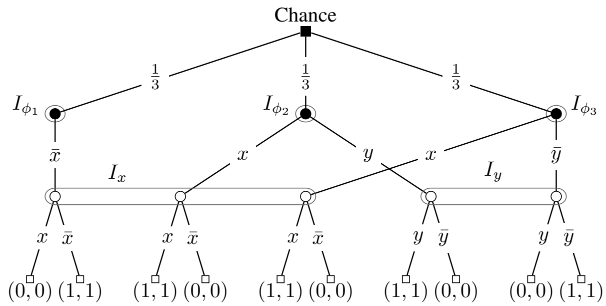

Figure 4: Two-player extensive-form game with chance moves for the SAT instance $C = \{\bar{x}, x \vee y, x \vee \bar{y}\}$.

for each clause; in this case, the expected utility for both player is 1 and thus the social welfare is 2. If the SAT formula is not satisfiable, then for any truth assignment there exist at least one clause for which no literal can evaluate to true, which means that any strategy profile will lead to at least one $(0, 0)$ outcome and thus the social welfare will be strictly smaller than 2. Since any Nash equilibrium is trivially also an EFCCE, and since 2 is the optimal outcome of the game, the pure Nash equilibrium described above is also an optimal EFCCE. Hence, any polynomial-time algorithm for deciding whether it exists an EFCCE with a social-welfare of at least 2 could be employed to decide whether a SAT formula is satisfiable or not. □

Proposition 7. *$SW_{EFCCE}(\kappa)$ is NP-Hard in three-player games, with or without chance moves.*

Proof. The proof is similar to the one for two players with chance, with the only difference that we now have to employ the third player to simulate the random chance move. This requires only a small modification in the extensive-form game. In fact, it is sufficient to replace the player at the root node with Player 3, that can choose one action in the set $\{a_\phi \mid \phi \in C\}$. The game then proceeds unchanged, with Player 3 receiving as payoff $1 - u_1$ where u_1 is the payoff received in the same terminal node by Player 1.

If the SAT formula is satisfiable, then the game admits at least one Nash equilibrium with a social welfare of 2. Regardless of the strategy employed by Player 3, Player 1 can in fact always select one literal for each clause such that they all evaluate to true under the truth assignment played by Player 2. If the SAT formula is not satisfyable, then every Nash equilibrium has a social welfare strictly smaller than 2, as Player 3 is incentivized to play in a way as to reach clauses that are not satisfiable under the truth assignment played by Player 2, and there is always at least one such clause for each strategy of Player 2. Thus, the same argument of the two player with chance proof still holds, hence the SAT problem can be reduced to the $\text{SW}_{\text{EFCCE}}(\kappa)$ one. □

Proposition 3. *$SW_{EFCCE}(\kappa)$ is NP-Hard in two-player games with chance moves, as well as in three-player games, with or without chance moves.*

Proof. See Proposition 6 and Proposition 7 □

Additional Figure

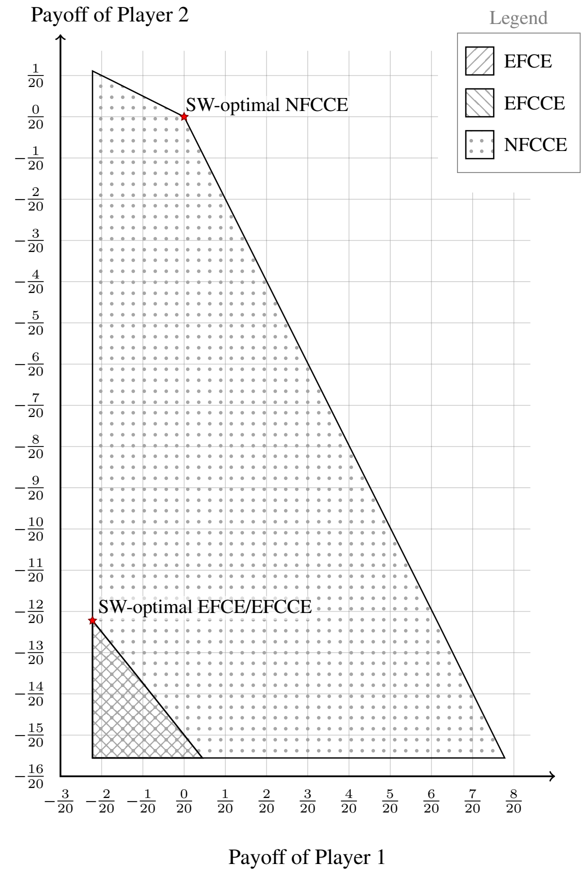

Figure 5: Payoff-space representation of a battleship game with a 2×1 board, one ship per player, and two rounds. In this case, $U_{\text{EFCE}} = U_{\text{EFCCE}} \subsetneq U_{\text{NFCCE}}$.

Game Instances Used in Our Experiments

We give a detailed description of the three games used in the Experiments section.

Sheriff

Sheriff (Farina et al. 2019) is a bargaining game inspired by the board game *Sheriff of Nottingham*. The game is played by two players: the *Smuggler* and the *Sheriff*. At the beginning of the game, the Smuggler must decide the number n of illegal items to load in his or her cargo (from 0 up to $n_{\max}$). Each illegal item has the fixed value v. After that, the two players engage in r rounds of bargaining. At each round, the Smuggler selects a bribe ranging from 0 to $b_{\max}$, and the Sheriff gives feedback as to whether or not the Sheriff would inspect the Smuggler's cargo given that bribe amount. The Sheriff's feedback is binding only in the last round of bargaining. If during the last round of bargaining the Sheriff accepts the bribe, the game ends, the Smuggler receives a utility of $v \cdot n$, minus the bribe amount b that was offered in the last bargaining round, and the Sheriff receives a utility

Num Ranks r	Seq. Pairs $\lvert\Sigma_1\rvert \times \lvert\Sigma_2\rvert$	Rel. Seq. Pairs
2	4.900×10^1	3.700×10^1
3	3.364×10^3	6.640×10^2
4	8.409×10^5	1.793×10^4
5	5.247×10^8	6.772×10^5
6	6.799×10^{11}	6.974×10^6

(a) Instance parameters and dimensions

NFCCE			EFCCE	EFCE
LP	CG-LP	CG-MILP	LP	LP
27ms	86ms	167ms	18ms	17ms
26ms	6.83s	355ms	27ms	29ms
222ms	6m 06s	1m 42s	378ms	444ms
9.99s	> 24h	> 24h	14.50s	20.27s
14m 01s	oom	oom	42m 26s	1h 11m

(b) Run times of algorithms

NFCCE	EFCCE	EFCE
3	3	3
6	6	6
10	10	10
15	15	15
21	21	21

(c) Maximum social welfare

Table 2: Experimental results on several instances of the Goofspiel game. 'oom': Out of memory.

equal to b. If the Sheriff does not accept the bribe in last bargaining round (that is, the Sheriff decides to inspect the cargo), the game can end in two possible ways:

- if the cargo contains no illegal items (that is, $n = 0$), the utility of the Smuggler is set to a fixed penalty amount s, while the utility of the Sheriff is $-s$;
- if the cargo contains some illegal items (that is, $n > 0$), the utility of the Smuggler is set to $-p \cdot n$, where p is a parameter of the game called the Smuggler's penalty, and the utility of the Sheriff is $p \cdot n$.

The Sheriff instances that we use in our experiments are parametric over the maximum number n_{max} of illegal items that the Smuggler can load in her cargo, the maximum bribe b_{max} that can be offered to the Sheriff, and the number r of bargaining rounds between the two players. Increasing any of this parameters affects the size of the resulting game instance. The other parameters of the game are set to the fixed values $v = 5, p = 1, s = 1$ in all of our game instances.

Battleship

Battleship (Farina et al. 2019) is a parametric version of the classic board game, where two competing fleets take turns at shooting at each other. At the beginning of the game, the players take turns at secretly placing a set of ships on separate grids (one for each player) of size $w \times h$. Each ship has a size (measured in terms of contiguous grid cells) and a value, and must be placed so that all the cells that make up the ship are fully contained within each player's grids and do not overlap with any other ship that the player has already positioned on the grid. After all ships have been placed. the players take turns at firing at their opponent. Ships that have been hit at all their cells are considered sunk. The game continues until either one player has sunk all of the opponent's ships, or each player has completed r shots. At the end of the game, each player's payoff is calculated as the sum of the values of the opponent's ships that were sunk, minus γ (a constant known as *loss multiplier*) times the sum of the ships which that player has lost.

Our instances of the Battleship game are parametric on the grid size (w, h), and the maximum number of rounds r that players have. The loss multiplier γ was set to the fixed value 2 in all of our instances. Furthermore, each player has always a single ship of length 1 and value 1.

Goofspiel

The variant of Goofspiel (Ross 1971) that we use in our experiments is a two-player card game, employing three identical decks of r cards each. At the beginning of the game, each player receives one of the decks to use it as its own hand, while the last deck is put face down between the players, with cards in increasing order of rank from top to bottom. Cards from this deck will be the prizes of the game. In each round, the players privately select a card from their hand as a bet to win the topmost card in the prize deck. The selected cards are simultaneously revealed, and the highest one wins the prize card. In case of a tie, the prize card is discarded. Each prize card's value is equal to its face value, and at the end of the game the players' score are computed as the sum of the values of the prize cards they have won.

Additional Experimental Results

Table 2 shows some experimental results for Goofspiel. The Experiments section in the body of the paper describes how the table is structured.