Limits on the Power of Cryptographic Cheap Talk[⋆]

Pavel Hubáček[1,⋆⋆], Jesper Buus Nielsen[1], and Alon Rosen[2,⋆⋆⋆]

[1] Aarhus University
[2] IDC Herzliya

Abstract. We revisit the question of whether cryptographic protocols can replace correlated equilibria mediators in two-player strategic games. This problem was first addressed by Dodis, Halevi and Rabin (CRYPTO 2000), who suggested replacing the mediator with a secure protocol and proved that their solution is stable in the Nash equilibrium (NE) sense, provided that the players are computationally bounded.

We show that there exist two-player games for which no cryptographic protocol can implement the mediator in a sequentially rational way; that is, without introducing empty threats. This explains why all solutions so far were either sequentially unstable, or were restricted to a limited class of correlated equilibria (specifically, those that do not dominate any NE, and hence playing them does not offer a clear advantage over playing any NE).

In the context of computational NE, we classify necessary and sufficient cryptographic assumptions for implementing a mediator that allows to achieve a given utility profile of a correlated equilibrium. The picture that emerges is somewhat different than the one arising in semi-honest secure two-party computation. Specifically, while in the latter case every functionality is either "complete" (i.e., implies Oblivious Transfer) or "trivial" (i.e., can be securely computed unconditionally), in the former there exist some "intermediate" utility profiles whose implementation is equivalent to the existence of one-way functions.

1 Introduction

The field of game theory offers a variety of ways to reason about the behavior of rational players. One of the most famous analytic tools for that purpose is that of Nash equilibrium [14]. In the basic case of two-player games, a Nash equilibrium (NE) constitutes of two independent plans of action, one for each player, such that no player can unilaterally benefit by deviating from her own plan. The NE solution concept was subsequently generalized by Aumann [2], who allowed players to pick their actions in a correlated way. Correlated equilibria (CE) are in many cases preferable over NE, in part because they can potentially guarantee higher utility to the players. In order to be

[⋆] Supported by the European Research Commission Starting Grant 279447, ISF grant no. 334/08, and Danish National Research Foundation and The National Science Foundation of China (grant 61061130540) for the Sino-Danish Center for the Theory of Interactive Computation.

[⋆⋆] The author did part of this work while visiting IDC Herzliya supported by the ISF grant no. 334/08.

[⋆⋆⋆] Supported by the ISF grant no. 334/08.

able to act in a correlated manner, the players are assumed to have access to a mediator (sometimes referred to as correlation device), that provides them with private, correlated, recommendations on the action to be taken.

About a decade and a half ago, Dodis, Halevi and Rabin [7] pointed out the possibility of implementing the mediator without having to refer to any trusted party. To this end, they proposed the use of secure two-party computation, viewing the correlation device as a randomized functionality. Their approach, natural from the cryptographic perspective, gives rise to some game theoretical challenges that need to be addressed. Most notably, the cryptographic protocol preceding the actual play of the strategic game introduces new actions that are observable by the players. Since these actions take place sequentially, the model of the game needs to be adjusted to account for the strategic decisions that players need to take during the protocol execution. While these actions do not directly affect the utility in the underlying strategic game, they can nevertheless influence the players in their decision making. Such pre-play communication is referred to as *cheap talk* in the economic literature.

One crucial difference from the mediated setting, which is inherited from the sequential nature of protocols, is that one of the players may learn her recommendation before the other. If this player is not happy with the protocol's recommendation, she can simply decide to "abort," thus preventing the other player from learning his own recommendation. Another crucial difference is that player A (not necessarily the one who learns her recommendation first), can reveal extra information to player B, changing player B's knowledge and expectation on how player A is going to play.[3]

Given that such deviations can always be observed, it becomes necessary to specify what action players take in case deviation is detected. One could attempt to deter misbehavior by threatening with some punishment. However, it is not *a priori* clear what kind of punishment should a player invoke, assuming that the other player is rational. In the protocol of Dodis *et al.* [7], an "abort" action is punished by employing the min-max strategy (that is, the strategy that minimizes the maximal gain of the deviator). This approach suffers from the well known and often unavoidable shortcoming of being harmful to the punishing player. Consequently, the *threat* of playing the min-max strategy is *empty*, or in other words not credible. Punishing the other type of deviations, in which the deviating player reveals extra information, appears to be even more challenging, as a message reacting to such deviations might not even fall into the scope of the prescribed protocol (for instance, if the deviating player is the last to learn her recommendation, meaning that the protocol actually terminates at that point).

The issue of empty threats is classically handled by the requirement of subgame perfection (SPE), which requires strategies to be in equilibrium at any point during the protocol execution. This requirement insures that any threat is credible. One problem with subgame perfection, that is particularly acute when modeling behavior of computationally bounded players in a cryptographic protocol, is the requirement of optimality at any point in the protocol execution. This problem was first addressed by Gradwohl, Livne and Rosen [9], who by defining empty threats in an explicit manner, were able to reason

[3] For instance, the second player to learn his own recommendation could make his private view of the protocol public, thus revealing his recommendation to the first player and rendering the correlation device useless.

about sequential rationality in face of computationally bounded players. In addition to this modeling, their work proposed a simple cryptographic protocol for the class of convex hull Nash equilibria (i.e., correlated equilibria that can be expressed as a convex combination of the Nash equilibria of the game), assuming the existence of one-way functions. To avoid empty threats, their solution punishes the aborting player with her “worst" NE (i.e., the NE yielding the lowest payoff amongst all NE in the game). Indeed, since the punishment is a Nash equilibrium, a rational punishing player has no incentive to deviate from it, which renders the threat of playing this NE credible.

One significant shortcoming of the Gradwohl *et al.* [9] solution is that it only applies to convex combinations of Nash equilibria. Unfortunately, such equilibria are not very interesting since they do not enjoy the most beneficial feature of CE, namely the ability of dominating the payoffs achieved by any NE. This leaves open the question of whether there exists a sequentially rational cryptographic protocol for implementing the mediator in the cases where playing a CE is preferable over playing any NE.

1.1 Our Results

A necessary requirement for guaranteeing sequential rationality is the ability for a player to threaten credibly. For this to be possible the threat must consist of a rational plan of action. Otherwise, there is no guarantee that a rational player will follow through in case she is tested. We formalize this intuition by putting forward the notion of *Nash equilibrium punishable CE*. These are correlated equilibria for which the expected utility of any player given a recommendation by the mediator is never smaller than in her worst NE. This notion turns out to be crucial for implementing the mediator of a CE using a cryptographic protocol.

Theorem 2 (informal). *A correlated equilibrium can be implemented in a sequentially rational way using cryptographic cheap talk if and only if it is Nash equilibrium punishable.*

Given the above theorem, it is natural to ask whether every CE is NE-punishable. An affirmative answer would have implied that any player receiving an unsatisfactory recommendation from the cryptographic protocol can be threatened from aborting in a credible way.

Our answer to this question is negative. We show that there exist games with CE that are not NE-punishable. Moreover, these games have utility profiles that can be obtained only by those CE that are not NE-punishable (and so cannot be achieved by other NE-punishable equilibria). Additionally, both players prefer these utility profiles to utility profile of some other NE-punishable CE, thus both would be in favor of implementing such *preferable* CE.

Theorem 1 (informal). *There exist infinitely many strategic games with preferable CE that cannot be achieved by sequentially rational cryptographic cheap talk.*

The above theorem explains why all solutions so far were either sequentially unstable, or were restricted to a limited class of correlated equilibria.

In addition to the above results, we classify necessary and sufficient cryptographic assumptions for implementing a mediator that allows to achieve a given utility profile of a CE by a protocol that is in computational NE. We show that there are non-trivial CE in the convex hull of Nash equilibria[4] (CHNE) which can be implemented via cheap talk only if one-way functions exist.

Theorem 3 (informal). *If the payoff of all non-trivial convex hull Nash equilibria can be achieved via cryptographic cheap talk then one-way functions exist.*

As shown by Gradwohl *et al.* [9], if one-way functions exist then all non-trivial CE in the convex hull of NE can be implemented via computational (and moreover sequentially rational) cheap talk. Taken together these results fully characterize the assumptions under which all convex hull NE can be implemented. We also show that there exist CE outside CHNE which can only be cheap talk implemented if OT exists.

Theorem 4 (informal). *If the payoff of all correlated equilibria outside the convex hull of NE can be achieved via cryptographic cheap talk then there exists a protocol for oblivious transfer (OT).*

As shown by Dodis *et al.* [7], if there exists a protocol for OT then all correlated equilibria (including those outside the convex hull of NE) can be implemented via computational (but not necessarily sequentially rational) cheap talk. Taken together these results show that OT is complete for implementing all CE (regardless of the issue of sequential rationality). We conjecture that implementing *any* CE outside the CHNE and provide evidence to support the conjecture. We leave it as an open problem to prove or disprove the conjecture.

These are to our best knowledge the first results of this type. Previous work on rational cryptography has focused on sufficiency of cryptography for implementing equilibria. Our results suggest an intriguing connection between the distinction between CE and CHNE on one hand and the distinction between Cryptomania and Minicrypt on the other hand (see Impagliazzo [11]). The picture that emerges is somewhat different than the one arising in semi-honest secure two-party computation. While in the latter case every functionality is either “complete" (i.e. implies OT) or “trivial" (i.e. can be securely computed unconditionally), in the former there exist some “intermediate" utility profiles whose implementation is equivalent to the existence of one-way functions. The details are given in Sect. 6 and Sect. 7.

1.2 Related Work

Osborne and Rubinstein [15] provide a standard introduction to game theory. The notion of correlated equilibrium was introduced by Aumann [2]. A non-technical introduction motivating the notion of cheap talk is given in Farrell and Rabin [8]. Cheap talk implementation of a correlation device in game-theoretical framework was put forward by Bárány [4]. Aumann and Hart [3] show what equilibria payoffs can be achieved via cheap talk preceding games with imperfect information.

We already mentioned the works in [7, 9]. Teague [17], and subsequently Atallah *et al.* [1] gave a protocol for the general problem of correlated element selection achieving

[4] Note that NE, even though contained in the convex hull of NE, are trivial from our perspective, since there is no need for a mediator to play according to them.

better efficiency than [7], but preserving the original solution concept of computational NE. Using results from computational complexity to implement correlation devices was considered by Urbano and Vila [19], aiming for a similar result to Dodis *et al.* [7]. However, Teague [18] showed that their approach is flawed. An alternative solution concept for analyzing game theoretical properties of cryptographic protocols was suggested by Pass and shelat [16].

2 Preliminaries and Definitions

For $m \in \mathbb{N}$, we use $[m]$ to denote the set $\{1,\ldots,m\}$. For a finite set A, we use $\Delta(A)$ to denote the set of probability distributions over A.

Definition 1 (Two-player strategic game). *A* two-player strategic game Γ *is a triple* (A_1,A_2,u)*, where* A_i *is a set of actions of player* $i \in \{1,2\}$*, and* $u : A_1 \times A_2 \to \mathbb{R}^2$ *is a utility function assigning a utility profile to every action profile* $a \in A_1 \times A_2$*. We use* u_i *to denote the i'th output of u, i.e.,* $u(a) = (u_1(a),u_2(a))$*.*

In this work we only consider two-player games. Also, we talk about a $k \times k$ strategic game Γ if both players have k strategies in Γ, i.e., $|A_1| = |A_2| = k$. A classical example of strategic game is the game of Chicken as in Fig. 1a.

Definition 2 (Strategy profile). *A* strategy profile *for a strategic game* Γ *is a probability distribution* γ *on* $A_1 \times A_2$*, i.e.,* $\gamma \in \Delta(A_1 \times A_2)$*. We denote* $\gamma(a)$ *the probability assigned by* γ *to* $a \in A_1 \times A_2$*. The corresponding utility profile* $U(\gamma) \in \mathbb{R}^2$ *is given by* $U(\gamma) = (U_1(\gamma),U_2(\gamma))$*, where* $U_i(\gamma) = \sum_{(a_1,a_2)\in A_1\times A_2} \gamma(a_1,a_2)u_i(a_1,a_2)$ *for* $i \in \{1,2\}$*. If* $U(\gamma) = (v_1,v_2)$*, we say that* γ achieves the utility profile (v_1,v_2)*.*

Definition 3 (Correlated equilibrium). *A* correlated equilibrium (CE) *of a strategic game* (A_1,A_2,u) *is a strategy profile* $\gamma \in \Delta(A_1 \times A_2)$*, such that for every player* $i \in \{1,2\}$ *and every pair of strategies* $a_i, a'_i \in A_i$ *it holds that*

$$\sum_{a_{-i}\in A_{-i}} \gamma(a_i,a_{-i})u_i(a'_i,a_{-i}) \le \sum_{a_{-i}\in A_{-i}} \gamma(a_i,a_{-i})u_i(a_i,a_{-i}) \ .$$

We denote $U_i(\gamma|a_i)$ the expected utility of player i when given advice $a_i \in A_i$ and the other player also plays according to some advice sampled from γ, i.e., $U_i(\gamma|a_i) = \left(\sum_{a_{-i}\in A_{-i}} \gamma(a_i,a_{-i})\right)^{-1} \sum_{a_{-i}\in A_{-i}} \gamma(a_i,a_{-i})u_i(a_i,a_{-i})$.

Definition 4 ((Convex hull) Nash equilibrium). *A* Nash equilibrium (NE) *of a strategic game* $\Gamma = (A_1,A_2,u)$ *is a correlated equilibrium* σ *of* Γ*, such that* σ *is also a product distribution, i.e.,* $\sigma \in \Delta(A_1) \times \Delta(A_2)$*. A* convex hull Nash equilibrium (CHNE) *of a strategic game* Γ *is a correlated equilibrium of* Γ *that can be expressed as a convex combination of Nash equilibria of* Γ*.*

We denote $\mathrm{NE}(\Gamma), \mathrm{CHNE}(\Gamma)$, and $\mathrm{CE}(\Gamma)$ the set of Nash equilibria of Γ, the set of convex hull Nash equilibria of Γ, and the set of correlated equilibria of Γ respectively.[5]

[5] As a convention, we will use γ to denote a strategy profile that is a CE and σ to denote a strategy profile that is a NE.

We are interested in implementing correlated equilibria of two-player strategic games. Given such strategic game Γ one can visualize the utility profiles achievable by all its correlated equilibria in $\mathbb{R}^2$. Figure 1b depicts the polygon of utility profiles achievable by CE of the game of Chicken defined by the payoff matrix in Fig. 1a. The dark grey triangle corresponds to utility profiles achievable by the CHNE of Chicken, and its three corner points are exactly the payoffs of the three NE of the game of Chicken. One can see that the payoffs of CE of Chicken extend the region of CHNE payoffs in both directions, i.e., there are both CE that improve the CHNE payoffs (the white polygon) and those that are dominated by the CHNE payoffs (the light grey triangle).

	C	D
c	6,6	2,7
d	7,2	0,0

(a)

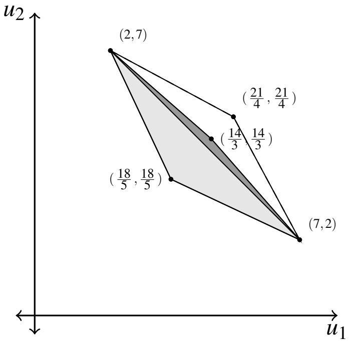

(b)

Fig. 1. (a) the game of Chicken (b) the utility profiles achievable by its CE.

There is a natural partial ordering on the utility profiles induced by the relation of Pareto dominance.

Definition 5 ((strict) Pareto dominance, weak Pareto optimality). *Let Γ be a strategic game, and $\gamma, \gamma' \in \mathrm{CE}(\Gamma)$. If $U_i(\gamma) > U_i(\gamma')$ for both $i \in \{1,2\}$, we say that γ* strictly Pareto dominates *γ'. We say that γ* Pareto dominates *γ' if for both $i \in \{1,2\}$ it holds that $U_i(\gamma) \geq U_i(\gamma')$, and there exist $i' \in \{1,2\}$ such that $U_{i'}(\gamma) > U_{i'}(\gamma')$. We say that a $\gamma^* \in \mathrm{CE}(\Gamma)$ is* weakly Pareto optimal *if there exists no $\gamma' \in \mathrm{CE}(\Gamma)$ that Pareto dominates γ^*.*

We sometimes abuse the notation and say that utility profile $v \in \mathbb{R}^2$ (strictly) Pareto dominates $v' \in \mathbb{R}^2$ if there exist $\gamma, \gamma' \in \mathrm{CE}(\Gamma)$, such that $v = U(\gamma), v' = U(\gamma')$ and γ (strictly) Pareto dominates γ'. Consider again the CE payoffs of Chicken in Fig. 1b. The two line segments between $(2,7)$ and $(\frac{14}{3},\frac{14}{3})$, and between $(\frac{14}{3},\frac{14}{3})$ and $(7,2)$ on the boundary of CHNE payoffs are exactly the *weakly Pareto optimal boundary* of the CHNE payoffs of Chicken.

3 Not all CE are NE-punishable

In this section, we show that there exists a barrier for using cryptography to implement any interesting correlated equilibrium without empty threats. Intuitively, for a correlated equilibrium to be implementable by a cryptographic protocol without empty threats, one must be able to effectively punish any deviating player by her worst NE.

Definition 6 (NE-punishable CE). *Let γ be a CE of a strategic game $\Gamma = (A_1, A_2, u)$. We say that γ is a* Nash equilibrium punishable correlated equilibrium *if for all $i \in \{1,2\}$ and every action $a_i \in A_i$ of player i played with non-zero probability in γ it holds that $U_i(\gamma|a_i) \geq U_i(\sigma_i)$, where σ_i is the worst Nash equilibrium for i in Γ.*

It is not at all obvious if there exists any strategic game with a CE that is not NE-punishable; it could also be the case that for any CE there exists a NE-punishable CE achieving the same utility profile. However, we show that none of the above is true. There are in fact many games with correlated equilibria that have some utility profile extending the polygon of CHNE payoffs, but no NE-punishable CE achieves such utility profile.

Theorem 1. *For any $k \in \mathbb{N}$. If $k > 3$, then there exists a $k \times k$ strategic game Γ with a correlated equilibrium $\gamma \in \mathrm{CE}(\Gamma) \setminus \mathrm{CHNE}(\Gamma)$, s.t. every $\gamma' \in \mathrm{CE}(\Gamma)$ with $U(\gamma') = U(\gamma)$ is not a NE-punishable CE of Γ.*

The proof is constructive. We start with a suitable $(k-1) \times (k-1)$ strategic game Λ and extend it into a $k \times k$ game Γ that exemplifies the theorem; the initial game Λ is characterised by some non-trivial properties (given by the criterion in Def. 7) that are exploited when we extend it.

Definition 7 (Extensibility Criterion). *A strategic game $\Lambda = (A_1, A_2, u)$ satisfies the* extensibility criterion *if there exists γ, a CE of Λ, with the following two properties:*

1. *γ strictly Pareto dominates any NE of Λ.*
2. *There exists $a \in A_i$ for some player $i \in \{1,2\}$, such that for every $\gamma' \in \mathrm{CE}(\Gamma)$ with $U(\gamma') = U(\gamma)$ it holds that $U_i(\gamma') > U_i(\gamma'|a)$.*

We use the fact that any strategic game Λ satisfying the extensibility criterion has a CE γ preferable for both players to any NE of Λ. The CE γ is preserved as a correlated equilibrium in the extended game Γ. We are able to carefully devise the payoffs of Γ such that its unique NE is strictly Pareto dominated by γ, however for at least one of the players there exists a recommendation in γ that is inferior to the unique NE.

Lemma 1. *For any $k \in \mathbb{N}^+$, if there exists a $(k-1) \times (k-1)$ strategic game Λ_{k-1} that satisfies the extensibility criterion, then there exists a $k \times k$ strategic game Γ with a correlated equilibrium $\gamma \in \mathrm{CE}(\Gamma) \setminus \mathrm{CHNE}(\Gamma)$, s.t. every $\gamma' \in \mathrm{CE}(\Gamma)$ with $U(\gamma') = U(\gamma)$ is not a NE-punishable CE of Γ.*

Proof. We show how to extend $\Lambda_{k-1} = (A, B, u)$ with one additional action for each player to define Γ. Let a_0 be the new action of player A and b_0 be the new action of player B, thus $\Gamma = (A \cup \{a_0\}, B \cup \{b_0\}, u')$. The utility function u' of Γ corresponds to the utility function of Λ_{k-1} for every action profile in $A \times B$. For some $s, t \in \mathbb{R}$, u' is defined on the remaining action profiles as: $u'(a_0, b_0) = (t, t)$, and $u'(a_0, b) = u'(a, b_0) = (s, s)$ for every $b \in B$ and every $a \in A$.

We show that it is possible to select s and t such that the claim holds. Recall that Λ_{k-1} satisfies the extensibility criterion, so there exists a CE γ satisfying the two conditions from Def. 7. Let i be the player and $a \in A_i$ be the advice from the second condition of the extensibility criterion. Denote v the expectation of player i in γ given recommendation a, i.e., $v = U_i(\gamma|a)$. We can assume without loss of generality that γ is the CE with maximal v. Let v' be the maximal utility obtained in Λ_{k-1} by any of the players in some NE, i.e., $v' = \max(U_A(\sigma_A^*), U_B(\sigma_A^*))$, where σ_i^* is the best NE for player i. Set s such that $\max(v, v') < s < U_i(\gamma)$, and let $t = (s + U_i(\gamma))/2$.

If s and t are selected as above, then no Nash equilibrium of Λ_{k-1} is a Nash equilibrium of Γ. Moreover, the action profile (a_0, b_0) is a unique NE of Γ achieving the utility profile (t,t). However, γ is still a correlated equilibrium in Γ, and the expectation of player i when given a as a recommendation is strictly smaller than the utility obtained by player i in the unique NE (a_0, b_0) of Γ. Thus γ is not a NE-punishable CE.

Consider any other CE γ' of Γ that achieves the same utility profile as γ. Both t and s are smaller than $U_i(\gamma)$, thus any new correlated equilibrium achieving $U(\gamma)$ satisfies the second condition from the extensibility criterion. Since $U_i(\gamma|a) \geq U_i(\gamma'|a)$, any such γ' is also not NE-punishable. □

It remains to show that games satisfying the extensibility criterion exist for any $k > 2$.

Lemma 2. *For every $k \in \mathbb{N}$ with $k > 2$, there exists a $k \times k$ strategic game Λ_k that satisfies the extensibility criterion.*

Proof. Let $c, d, e, f, g \in \mathbb{R}$ be real numbers such that $c < d < e < f < g$, where $g - f < e - c$, and $3f < (e - c)$.[6] Consider the $k \times k$ game $\Lambda_k = (A = \{a_1, \ldots, a_k\}, B = \{b_1, \ldots, b_k\}, u)$ with the utility function $u : A \times B \to \mathbb{R}^2$ defined as follows:

- $u(a_j, b_j) = (f, g)$ for every $j \in [k-1]$,
- $u(a_k, b_k) = (d, e)$,
- $u(a_j, b_{j+1}) = (g, f)$ for every $j \in [k-2]$,
- $u(a_{k-1}, b_k) = (e, f)$,
- $u(a_k, b_1) = (g, d)$, and
- $u(a, b) = (c, c)$ otherwise.

To illustrate the corresponding payoff matrix, we give the payoff matrix of Λ_4 in Fig. 2.

	b_1	b_2	b_3	b_4
a_1	f,g	g,f	c,c	c,c
a_2	c,c	f,g	g,f	c,c
a_3	c,c	c,c	f,g	e,f
a_4	g,d	c,c	c,c	d,e

Fig. 2. The payoff matrix of Λ_4.

Due to the restrictions on the entries in the payoff matrix, there is no pure Nash equilibrium in Λ_k. Indeed, for every action profile $(a, b) \in A \times B$ there exists either an

[6] The two conditions $g - f < e - c$ and $3f < (e - c)$ are required for ease of exposition when describing the candidate CE. In fact, Λ_k defined without this conditions would also satisfy the claim of Lemma 2.

action a' of player A or an action b' of player B, such that A prefers (a',b) to (a,b) or B prefers (a,b') to (a,b). Following the same reasoning, Λ_k can only have fully mixed Nash equilibria. Notice that any of such NE assigns non-zero probability to the action profiles with utility profile (c,c).

We describe a candidate CE for the claim of Lemma 2. Let γ_k be any probability distribution on $A \times B$ satisfying these conditions.

1. $\gamma_k(a_k,b_1) = \gamma_k(a_{k-1},b_k) = \gamma_k(a_k,b_k) = \frac{g-f}{3(g-f)+(2k-3)(e-c)}$,
2. $\gamma_k(a,b) = \frac{e-c}{3(g-f)+(2k-3)(e-c)}$ for every $(a,b) \notin \{(a_k,b_1),(a_{k-1},b_k),(a_k,b_k)\}$ such that $u(a,b) \neq (c,c)$, and
3. $\gamma_k(a,b) = 0$ otherwise.

A proof of the following claim is given in the full version.

Claim. Any such probability distribution $\gamma_k \in \Delta(A \times B)$ is a correlated equilibrium of Λ_k.

Moreover, γ_k has in its support only the action profiles that do not yield the utility profile (c,c). Therefore, any such CE strictly Pareto dominates any completely mixed NE of Λ_k.

The expectation $U_A(\gamma_k)$ of player A is

$$((k-1)f+(k-2)g)\,\frac{e-c}{3(g-f)+(2k-3)(e-c)} + \frac{(d+e+g)(g-cf)}{3(g-f)+(2k-3)(e-c)} \, ,$$

and this is strictly larger than f when $3f < (e-c)$. On the other hand, any correlated equilibrium γ_k' of Λ_k that achieves the same utility profile as γ_k must assign non-zero probability to every action profile with utility profile different from (c,c). Since the highest utility of player A obtained from any action profile in which A plays action a_{k-1} is f, the expectation of A in any such correlated equilibrium γ_k' when given recommendation a_{k-1} is at most f. Therefore, Λ_k satisfies the extensibility criterion. □

4 Computational Cheap Talk Simultaneous Move Games

In this section we present an overview of our game theoretical model and solution concepts. Full details are given in the full version.

Our core object of study is so-called computational cheap talk, simultaneous move (CTSM) games. A CTSM game without types is fully *specified* by a strategic game (A_1,A_2,u). The game itself is an extensive game with imperfect information modeling an interactive protocol, where the agents take turn in exchanging messages, with agent 1 arbitrarily being chosen to send the first message. At some point each agent must additionally pick an action $a_i \in A_i$ for (A_1,A_2,u). The utility of a play is $u(a_1,a_2)$, i.e., the utility does not depend on the communication, only the actions. We assume that the agents do not get any information on what the action of the other party is, and hence consider the choice of actions for (A_1,A_2,u) as simultaneous moves. The strategy σ_i of agent i specifies which messages to send in response to the messages sent by the other agent, and which action to pick for (A_1, A_2, u) at the end of the cheap talk. We require that σ_i is poly-time, to allow using cryptography. Any mixed strategy should also by poly-time computable. To conveniently model this, we technically only allow pure strategies, and then we give each such strategy an extra input r_i, which is a uniformly random bit-string not observed by the other agent. Any mixing must be implemented by $\sigma_i(r_i)$ in poly-time.

As described above, for each strategic game (A_1, A_2, u), we have a CTSM game. Correspondingly, for each CTSM game, we have a strategic game, which is just the game (A_1, A_2, u) used to specify it. We say that a CE for a strategic game can be *cheap talk implemented* if there exists a strategy $\sigma = (\sigma_1, \sigma_2)$ for the corresponding CTSM game which obtains the same utility profile as the CE and which is a *computational* NE, which is just an ε-NE for a negligible ε. We say that a CE for a strategic game can be *ETF cheap talk implemented* if it can be cheap talk implemented by some σ which is additionally *empty-threat free*. We define empty-threat freeness along the lines of [9], specialize their general definition to the setting of CTSM games and generalizing to handle imperfect information. The details are in the full version. Here we sketch and motivate the definition.

An *empty threat* posed by me in a CTSM game is a part of my future strategy which I do not currently play and which I would not play should you call my bluff by deviating in a way making the threatening strategy active. You would *demonstrate* the existence of such a future empty threat posed by me by specifying a deviation by you which would make me deviate from playing the supposedly empty threat. We adopt this constructive definition, an advantage being that we can insist that the demonstration be poly-time. Note, however, that using an empty threat to force me to deviate from a threat does not convincingly demonstrate that my threat was empty. We therefore require that your demonstrator itself is empty threat free in future play. Formally we require that the deviation meant to demonstrate the existence of a future empty threat occurs in response to some event D, for *deviate*, and require that the demonstration be empty-threat free in the sub-game defined by D occurring.

Another qualification is that a deviation which makes me abstain from my threat, but which does not at the same time result in you receiving a larger expected utility does not demonstrate that I posed an empty threat. Yes, your deviation made me not execute the threat, but the threat did not serve to prevent you from this particular deviation, as you have no incentive for your deviation in the first place. All in all, a credible demonstration that I am posing an empty threat on you would therefore be an event D observable by you, and a deviation, which you only make when D occurs, which has the property that it leads to an empty-threat free future play, in the sub-game defined by D occurring, in which you have higher utility.

Formalizing the above definition and making it work well with the computational issue, is highly non-trivial, but none of the details really matter for the intuition of the results we describe later. For details, see the full version. Here we only mention and motivate the two main technical choices.

Since our definition of ETF is recursive, we need a last round to start from. Yet, our strategies are allowed any polynomial number of rounds, and the nature of most settings naturally modeled by CTSM games does not make it seem reasonable to postulate some exogenous fixed last round of communication, so we do not want to build a fixed last round into our model. Also, it is by far always given that a party can commit to an external action, like a bid in a real-life auction, until long after the cheap talk protocol was run, so we cannot guarantee that no more communication can take place after the protocol was run. I.e., the natural strategy space contains the possibility of more communication than needed exactly by the protocol in question, so our model should capture this. We essentially handle this by considering CTSM games families of games, $\varGamma = \{\varGamma_R\}_{R\in\mathbb{N}}$, where all $\varGamma_R$ have the same corresponding strategic game, and where $\varGamma_R$ has a fixed last round in round R. This allows to easily define ETF for each $\varGamma_R$, and we then say that σ is ETF if there exists R_0 such that it is ETF for all $\varGamma_R$ for $R \geq R_0$. I.e., the stability of a protocol is in particular not jeopardized by leaving some empty rounds after the execution of the strategy, i.e., rounds in which communication could have taken place. Robustness to the presence of such possible communication seems crucial for stability in real world networks.

We have chosen to use a similar mechanism to model poly-time. For a fixed strategic game (A_1,A_2,u) and $T \in \mathbb{N}$, let $\varGamma^T$ be the CTSM game corresponding to (A_1,A_2,u), where the messages and the action must be computable in time exactly T. For a polynomial p we consider a family of games $\varGamma_p = \{\varGamma^{(\kappa)} = \varGamma^{p(\kappa)}\}_{\kappa\in\mathbb{N}}$. A strategy $\sigma = \{\sigma^{(\kappa)}\}_{\kappa\in\mathbb{N}}$ for $\varGamma_p$ is one where $\sigma^{(\kappa)}$ is a strategy for $\varGamma^{(\kappa)}$. A strategy σ for $\varGamma$ is clearly poly-time. We say that σ is a computational NE for $\varGamma_p$ if there exists negligible ε such that $\sigma^{(\kappa)}$ is an $\varepsilon(\kappa)$-NE for $\varGamma^{(\kappa)}$. We call it a computational CTSM for (A_1,A_2,u) if there exists a polynomial p_0 such that it is a computational NE for $\varGamma_p$ for all $p \geq p_0$. Using the same flavor of definition to handle the computational issue and the no-last-round issue, allows to give one natural definition handling both issues.

Note that the above two design choices force proposed protocols to run in some fixed polynomial number of rounds and some fixed poly-time, whereas deviations are allowed to deviate to larger polynomials. This seems natural and strong.

To play a NE of any strategic game it is sufficient for the players to randomize independently, and there is no need for any cheap talk. The players need some publicly observable lottery to play according to a CHNE, that can be implemented using the protocol of Gradwohl *et al.* [9]. However, a CE outside the convex hull of NE needs some non-trivially correlated randomness. Motivated by our results from Sect. 6 and Sect. 7, we categorize correlated equilibria payoffs using the terminology of Impagliazzo [11].

Definition 8 (Trivial, Minicrypt, and Cryptomania utility profiles). *Let $\varGamma$ be a strategic game, and $v \in \mathbb{R}^2$ be a utility profile achieved by some $\gamma \in \mathrm{CE}(\varGamma)$.*

- *We call v a* trivial utility profile *if there exists $\sigma \in \mathrm{NE}(\varGamma)$ achieving v.*
- *We call v a* Minicrypt utility profile *if γ is a CHNE and there is no NE achieving v.*
- *We call v a* Cryptomania utility profile *if γ is not a CHNE.*

5 NE-punishable CE versus Empty-threat free NE

We can now formally relate NE-punishable CE and empty-threat free computational NE.

Theorem 2. *Let $\Gamma = (A_1, A_2, u)$ be a strategic game and let $\tilde{\Gamma}$ be the corresponding CTSM game. If there exists a strategy profile σ, a computational ETFE of $\tilde{\Gamma}$ with utility profile (v_1, v_2), then there exists a NE-punishable CE γ for Γ achieving the same utility profile (v_1, v_2).*

The theorem is proven in the full version. Here we provide a sketch of the proof. Consider any computational ETFE σ of $\tilde{\Gamma}$. Remember that σ is a family of strategies, and the utility profile of the members of the family need not converge to a fixed utility profile. However, we assume in the premise of the theorem that it does converge, to some (v_1, v_2). In the same vain, the action profiles of the members need not converge. However, the distribution of the action profile of all the strategies, i.e., the probability distribution over which actions $(a_1, a_2) \in A_1 \times A_2$ they make the players play, belong to a fixed compact space as we consider finite games Γ. Hence we can pick an infinite sub-sequence which converges to some probability distribution γ on $A_1 \times A_2$. It is possible to show that γ is a CE. Namely, in the games of the convergent sub-sequence, the incentive to deviate given any particular action is converging to 0, as σ in particular is an ε-NE for a negligible ε. This means that the incentive to deviate in the limit point γ is 0, by compactness. For the same reason γ has utility profile (v_1, v_2). We now assume that γ is not NE-punishable, and use this to show that σ is not empty threat free, which proves the theorem by contradiction.

If γ is not NE-punishable, then there exist $i \in \{1,2\}$ and an action $a_i \in A_i$ such that a_i occurs with non-zero probability and such that $U_i(\gamma|a_i) < U_i(\sigma_i^*)$, where σ_i^* is the worst NE for player i and $U_i(\gamma|a_i)$ is the expected utility of player i when playing γ given that the recommendation is a_i.

To prove that σ is not a computational ETFE we must pick a strategy space with enough rounds to run σ, or more rounds, and show that σ is not an ε-ETFE in this strategy space for any negligible ε. This in turn means that we must give an event D observable by P_2 (assume w.l.o.g. that $i = 2$) and a deviation for P_2 in the face of D for which he gets noticeably better expected utility in all ETF plays in the sub-game defined by D occurring.

As for the strategy space, pick the one which after the run of σ leaves at least one extra round of communication and where it is player 2 who sends a message in the last round of the strategy space. As the event D, pick the event that the output of running σ_2 is the bad action a_i for which $U_i(\gamma|a_i) < U_i(\sigma_i^*)$ and that κ is among the values in the infinite sub-sequence which converges to γ. As for the deviation, let player 2 play exactly as in σ_2, except that if D occurs, then player 2 does not play a_i. Instead, it waits until the last communication round where it sends its entire view of the protocol to player 1. Then player 2 picks an action a_2^* according to σ_2^*, and plays a_2^*. To show that σ is not a computational ETFE, it is now sufficient to show that in all ETF continuations after the last communication round, in the sub-game defined by D occurring, player 2 gets noticeably better expected utility than by playing σ. If this is not the case, then there exists an ETF continuation $\tilde{\sigma}$ after the last communication round, in the sub-game defined by D occurring, such that player 2 gets utility close to what he gets by playing σ when D occurs, which in turn is lower than what he gets by playing the worst NE. It follows that the utility profile of $\tilde{\sigma}$ is not the utility profile of a CHNE. Namely, a CHNE has a utility profile which is a convex combination of utility profiles for NE, so no player can get less than in his worst NE.

To conclude the proof by contradiction it is now sufficient to prove that $\tilde{\sigma}$ is a CHNE. Recall that $\tilde{\sigma}$ is played in the sub-game with a common prior C corresponding to the view of the parties after D occurred. Since player 2 sends his entire view to player 1 when D occurs, in the common prior C, player 1 can efficiently compute the signal of player 2. Denote the signal of player i by s_i. We use that $s_2 = s(s_1)$ for a fixed poly-time function s. If we give unbounded computing time to player 1 and only give it the signal s_2, then it can re-sample a random $(s'_1, s'_2) \leftarrow C$ with $s(s'_1) = s_2$ and play according to $\sigma_1(s'_1)$. This will lead to exactly the same strategy, and the unbounded computing power of player 1 does not allow it better deviations: since player 1 can efficiently compute $s_2 = s(s_1)$ from s_1 and since it knows the code σ_2 of player 2, it can use random runs of $\sigma_2(s_2)$ to sample the strategy profile of player 2 up to exponentially good precision in poly-time and and then in poly-time compute an optimal response to this fixed and now known strategy. Hence the unbounded computing power can at most give inverse exponentially more utility, which does not disturb the ε-NE. But then we have an ε-NE where the players have a common signal s_2. It is possible to use compactness of the strategy space to show that a sub-sequence of an ε-NE converges to a CHNE. The details are in the full version.

It is instructive to see how the above *reveal your view* deviation defeats some of the obvious attempts at circumventing the impossibility result.

Consider first a relaxed version of NE-punishable, which we could call *one-sided punishable*, where we only require that there *exists* $i \in \{1,2\}$ such that for every action $a_i \in A_i$ of player i played with non-zero probability in γ it holds that $U_i(\gamma|a_i) \geq U_i(\sigma_i)$, where σ_i is the worst Nash equilibrium for i in Γ. Say $i = 1$ without loss of generality. Consider the protocol which runs an unfair, active secure two-party computation where first player 1 learns a_1 and then in the following round player 2 learns a_2 or learns that player 1 aborted. If player 1 aborts, then player 2 punishes by playing the worst NE for player 1. It seems this should work as player 1 now has no incentive to deviate and player 2 cannot deviate as he learns his recommendation a_2 last. However, this does not work! What player 2 will do if he receives a bad recommendation a_2, i.e., one where $U_2(\gamma|a_2) < U_2(\sigma_2)$, where σ_2 is the worst Nash equilibrium for i in Γ, is to send his entire view, including a_2 to player 1, just before actions are to be played. Now that player 1 has no uncertainty on the view of player 1, all stable ways for the two players to pick their actions in the face of this deviation will give player 2 a payoff which is at least as good as in σ_2.

Consider then the attempt to use gradual release to give a_1 and a_2 to the players, the hope being that we can release a_1 and a_2 in a way such that when learning a_i it is too late to prevent the other party from learning a_{-i}. Again, this is in vain, as the *reveal your view* deviation is played *after* both a_1 and a_2 are fully revealed. For the same reason techniques for fair computation between rational players will fail too, like the protocol in Groce and Katz [10].

We consider it very interesting future work to consider variations of empty-threat freeness which prevent the *reveal your view* deviation, more specifically, can we give realistic models of empty-threat freeness allowing to implement larger classes of CE?

6 All Minicrypt Payoffs iff One-Way Functions Exist

Recall that we denote Minicrypt utility profiles to be the utility profiles achieved by some non-trivial CHNE. In this section we justify the name by showing that there exists a Minicrypt utility profile which requires one-way functions to be computational cheap talk implemented. This complements the result by Gradwohl *et al.* [9] that one-way functions are sufficient to implement any Minicrypt utility profile.

6.1 Implementing All Minicrypt Payoffs Implies One-Way Functions

In this section we show how to use a computational cheap talk implementation of some CHNE achieving a Minicrypt payoff to construct a protocol for weak coin-flip.

Given a two-party protocol $\pi = (\pi_1, \pi_2)$ with no inputs, and outputs which are in $\{0,1\}$. Let $y_i(\pi) \in \{0,1\}$ denote the output of π_i after running π. Note that $y_i(\pi)$ is a random variable, with the universe being the randomness used by P_1 and P_2 in the run of the protocol. A weak coin-flip protocol is such a protocol, where the following holds:

1. If both players are honest, then they output the same value, i.e., $y_1(\pi_1, \pi_2) = y_2(\pi_1, \pi_2)$. Moreover, $\Pr[y_1(\pi_1, \pi_2) = 0] = \Pr[y_1(\pi_1, \pi_2) = 1] = \frac{1}{2}$.
2. For any efficient strategy π_1^* of P_1 it holds that $\Pr[y_2(\pi_1^*, \pi_2) = 0] \leq \frac{1}{2} + \varepsilon$ for a negligible ε.
3. For any efficient strategy π_2^* of P_2 it holds that $\Pr[y_1(\pi_1, \pi_2^*) = 1] \leq \frac{1}{2} + \varepsilon$ for a negligible ε.

It follows from the seminal work of Impagliazzo and Luby [12] that weak coin-flip implies one-way functions.[7]

Consider the CTSM game specified by $\Gamma = (A_1, A_2)$, where $A_1 = \{\mathtt{c}, \mathtt{d}\}$, $A_2 = \{\mathtt{C}, \mathtt{D}\}$, and the utility function u is given by the payoff matrix:

	C	D
c	1,1	0,4
d	4,0	0,0

The probability distribution selecting $(\mathtt{c}, \mathtt{D})$ and $(\mathtt{d}, \mathtt{C})$ with equal probability is a convex hull NE achieving the utility profile $(2,2)$. We show that if it is possible to implement such CHNE using cryptographic cheap talk, then one-way functions exist.

Theorem 3. *If there exists in the CTSM game corresponding to Γ a computational NE σ achieving utility profile $(2,2)$, then one-way functions exist.*

Proof. Consider the two-party protocol π given in Fig. 3.

The following statements are logically equivalent.

[7] The notion is defined slightly different in [12], but by letting a party P_i who outputs "REJECT" output i instead, the notions become equivalent. Note also that opposed to what is common in contemporary definitions, see e.g. [13], we do not require that the winner can be determined from the communication of the protocol. This is in line with the original definition in [12], so we can still use the implication of one-way functions.

1. For $i \in \{1,2\}$, party P_i runs the cheap talk phase of strategy σ_i of P_i in the strategy profile σ, using uniformly random randomizers. All the messages are forwarded to party P_{-i}, and the round function is computed on the messages forwarded from P_{-i}.
2. If in round m the strategy σ_i plays d or C, then P_i outputs $y_i = 0$. If σ_i plays c or D, then P_i outputs $y_i = 1$.

Fig. 3. Protocol for weak coin-flip given a cheap talk implementation of a specific CHNE.

1. There exists an efficient π_1^* such that P_2 outputs 0 in (π_1^*, π_2) with probability $p_0 > \frac{1}{2}$.
2. There exists an efficient σ_1^* such that P_2 plays C in (σ_1^*, σ_2) with probability $p_0 > \frac{1}{2}$.
3. There exists an efficient σ_1^* such that P_1 has utility $u_0 > 2$ in (σ_1^*, σ_2).
4. There exists an efficient σ_1^* such that P_1 has utility $u_0 > 2$ in (σ_1^*, σ_2) and such that P_1 never plays c.

By construction statement 1 implies statement 2. If statement 2 is true, then the strategy $\sigma_1^\dagger$ which plays like σ_1^* and then plays d has expected utility $4p_0 > 2$. Statement 3 implies statement 4 because d is weakly dominating for P_1, i.e, P_1 never gets less utility by playing d instead of c. If statement 4 is true, then $4\alpha + 0(1-\alpha) > 2$, where α is the probability that P_2 plays c in (σ_1^*, σ_2). This implies that $\alpha > \frac{1}{2}$. By letting π_1^* be the strategy playing like σ_1^*, this implies statement 1.

If both parties follow the protocol in Fig. 3 then they both output the same bit b, and it is 0 or 1 with equal probability. Since σ is a computational equilibrium of $(\Gamma, C_\emptyset)$, any player can increase her utility by at most negligible amount. Thus, any player can bias the output of the protocol by at most negligible amount towards her preferred outcome, and the protocol is a weak coin-flip protocol. □

7 All Cryptomania Payoffs iff OT Exists

In this section we show that there exist Cryptomania profiles which imply OT. Implementing any Cryptomania profile given OT follows from [7]. We will also conjecture that implementing any Cryptomania profile implies OT and give supporting evidence.

We recall the notion of *random Rabin OT*. It is a secure two-party computation specified by a randomized function $f(x_1,x_2) = (y_1,y_2)$. The outputs do not depend on the inputs (x_1,x_2). The output y_1 is a bit $y_1 \in \{0,1\}$. The output y_2 is a trit $y_2 \in \{0,1,\perp\}$. The bit y_1 is uniformly random. The probability that $y_2 = \perp$ is $\frac{1}{2}$, independent of y_1. And, if $y_2 \neq \perp$, then $y_2 = y_1$. Note that this implies that party 1 gets no information on whether $y_2 = y_1$ or $y_2 = \perp$ and that if $y_2 = \perp$, then party 2 has no information on y_1. We call a protocol a *semi-honest random Rabin OT* if it implements random Rabin OT against parties guaranteed to follow the protocol in the model [5]. Semi-honest random Rabin OT is interesting as it is known to be complete for two-party computation, even for active secure two-party computation which can tolerate that the parties deviate from the protocol.

Given semi-honest random Rabin OT one can empty-threat free implement any NE-punishable CE. One uses an active-secure two-party computation to sample the CE and punishes a deviating party by playing the worst NE for that party. The proof that this is empty-threat free follows the proof of Gradwohl *et al.* [9]. We now show that OT is needed for having an implementation of all Cryptomania profiles.

7.1 Playing Chicken well implies OT

In this section we show that there exists a version of Chicken which has a CE with a weakly Pareto optimal utility profile which cannot be obtained using a computational NE in the corresponding cheap talk game, unless OT exists. The game has two actions per player, which shows that even in the simplest non-trivial game setting, one can only harvest the maximal utility if OT exists.

Consider the CTSM game specified by $\Gamma_{\mathtt{chicken}} = (A_1, A_2, u)$, where $A_1 = \{\mathtt{c}, \mathtt{d}\}$, $A_2 = \{\mathtt{C}, \mathtt{D}\}$ and the utility function u is given by the payoff matrix:

	C	D
c	15,15	6,21
d	21,6	0,0

Theorem 4. *If there exists a computational NE σ for the CTSM game corresponding to $\Gamma_{\mathtt{chicken}}$ achieving utility profile* $(14,14)$*, then there exists a protocol for semi-honest random Rabin OT.*

Proof. Let σ be as in the premise. We assume that $u(\sigma) = (14,14)$—extending the proof to handling the case where the payoff of each player i is $14 - \varepsilon_i$ for a negligible ε_i is standard. In the following we use $\mathrm{view}_i(\sigma) = \mathrm{view}_i(\Gamma, \sigma, C)$ to denote the view of player i when the parties play according to σ.

Consider the following two-party protocol π:

1. Party P_i runs the cheap talk phase of strategy σ_i of P_i in the strategy profile σ, using uniformly random randomizers.
2. If in the last round the strategy σ_i plays $\mathtt{c}$ or $\mathtt{C}$, then P_i outputs $b_i = 1$. If σ_i plays $\mathtt{d}$ or $\mathtt{D}$, then P_i outputs $b_i = 0$.

Let view_i denote the view of party P_i in a run of this protocol. We are going to analyze the distribution of the output of the parties and the distribution of their views, and then conclude that they imply OT.

Since the expected utility $(14,14)$ is symmetric, we know that σ plays $(\mathtt{d},\mathtt{C})$ as much as it plays $(\mathtt{c},\mathtt{D})$; call the probability of playing each of these α. Let β denote the probability that σ plays $(\mathtt{c},\mathtt{C})$. We clearly have that $2\alpha \leq 1-\beta$. The expected utility is therefore $\alpha(21,6)+\alpha(6,21)+\beta(15,15) \leq 2\alpha(13.5,13.5)+(1-2\alpha)(15,15)$. From $14 \leq 2\alpha 13.5+(1-2\alpha)15$, it follows that $\alpha \leq \frac{1}{3}$. This means that the expected utility is at most $\frac{1}{3}21+\frac{1}{3}6+\beta 15$. From $\frac{1}{3}21+\frac{1}{3}6+\beta 15 \geq 14$, we get that $\beta \geq \frac{1}{3}$. The expected utility of P_2 when σ_2 plays $\mathtt{C}$ is $\frac{\beta}{\alpha+\beta}15+\frac{\alpha}{\alpha+\beta}6$. If P_2 would switch to $\mathtt{D}$ when σ_1 says to play $\mathtt{C}$, then the expected utility of P_2 would become $\frac{\beta}{\alpha+\beta}21+\frac{\alpha}{\alpha+\beta}0$. It follows from the fact that σ is a computational NE that $\beta 21 \leq \beta 15+\alpha 6-\varepsilon$ for some negligible ε. We will assume that $\varepsilon = 0$—handling the negligible ε is standard. From $\beta 21 \leq \beta 15+\alpha 6$ we get that $\beta \leq \alpha$. From $\alpha \leq \frac{1}{3}$, $\beta \geq \frac{1}{3}$ and $\beta \leq \alpha$ we get that $\alpha = \beta = \frac{1}{3}$. This means that the joint output of $(\mathsf{P}_1, \mathsf{P}_2)$ in π is uniform on $\{(0,1),(1,0),(1,1)\}$. One can show that an expected constant number of samples from this distribution is sufficient to implement random Rabin OT, see the full version for the details. This, however, is not sufficient to conclude the proof, as the transcript of π might leak information. To finish the proof we therefore have to show that the parties have no extra information to their outputs, i.e., show that

$$[\mathrm{view}_1 \,|b_1 = 1 \wedge b_2 = 1] \approx [\mathrm{view}_1 \,|b_1 = 1 \wedge b_2 = 0]$$
$$[\mathrm{view}_2 \,|b_1 = 1 \wedge b_2 = 1] \approx [\mathrm{view}_2 \,|b_1 = 0 \wedge b_2 = 1] \;,$$

where $\approx$ denotes computational indistinguishability. We show the first relation. The second follows using a symmetric argument.

Assume that there exists an efficient distinguisher D which can distinguish $[\mathrm{view}_1 \,|b_1 = 1 \wedge b_2 = 1]$ and $[\mathrm{view}_1 \,|b_1 = 1 \wedge b_2 = 0]$ with non-negligible probability, i.e., $|\Pr[D([\mathrm{view}_1 \,|b_1 = 1 \wedge b_2 = 1]) = 1] - \Pr[D([\mathrm{view}_1 \,|b_1 = 1 \wedge b_2 = 0]) = 1]|$ is non-negligible. Since we work with non-uniform complexity, we can assume that it is always the case that $\Pr[D([\mathrm{view}_1 \,|b_1 = 1 \wedge b_2 = 1] = 1)] \geq \Pr[D([\mathrm{view}_1 \,|b_1 = 1 \wedge b_2 = 0] = 1)]$. Now consider the following strategy σ_1^*. It plays like σ_1, except that if σ_1 recommends to play $\mathtt{c}$, then σ_1^* switches to $\mathtt{d}$ when $D(\mathrm{view}_1) = 1$, where view_1 is the view of P_1. Note that σ_1 recommending to play $\mathtt{c}$ is logically equivalent to $b_1 = 1$. I.e., $\mathrm{view}_1 \in \{[\mathrm{view}_1 \,|b_1 = 1 \wedge b_2 = 1], [\mathrm{view}_1 \,|b_1 = 1 \wedge b_2 = 0]\}$. Furthermore, since $\alpha = \beta$, we have that b_2 is uniformly random. We use this to compute the utility of switching. We look at the cases that the joint play of σ is $(\mathtt{c},\mathtt{C})$ and $(\mathtt{c},\mathtt{D})$ separately. If the joint play is $(\mathtt{c},\mathtt{C})$, then we switch with probability $\Pr[D([\mathrm{view}_1 \,|b_1 = 1 \wedge b_2 = 1]) = 1]$, for a gain of $\Pr[D([\mathrm{view}_1 \,|b_1 = 1 \wedge b_2 = 1]) = 1](21 - 15)$. If the joint play is $(\mathtt{c},\mathtt{D})$, then we switch with probability $\Pr[D([\mathrm{view}_1 \,|b_1 = 1 \wedge b_2 = 0]) = 1]$, for a gain of $\Pr[D([\mathrm{view}_1 \,|b_1 = 1 \wedge b_2 = 0]) = 1](0-6)$. This gives a total gain of $6(\Pr[D([\mathrm{view}_1 \,|b_1 = 1 \wedge b_2 = 1]) = 1] - \Pr[D([\mathrm{view}_1 \,|b_1 = 1 \wedge b_2 = 0]) = 1])$. This means that the gain is six times the advantage of D, which is non-negligible. This is a contradiction to σ being a computational NE. □

7.2 Perfectly Implementing any CE outside CHNE Implies Unconditional OT

We now justify the conjecture that cheap talk implementing any Cryptomania profile implies OT. In particular, we show that if the implementation had been perfect, in the sense that it only leaks the recommendations, then one can always implement OT. We leave it as an open problem to investigate whether the additional protocol transcript of a cheap talk implementation of the correlation device in general leaks sufficiently little information that the result also holds for computational cheap talk implementations.

Theorem 5. *Let γ be a Cryptomania correlation device for a game Γ, i.e., it outputs recommendations which are not in the CHNE of Γ. Then given a polynomial number of samples of γ, two parties can implement unconditionally secure OT against semi-honest adversaries in the model [5].*

We use the result of Crépeau, Morozov and Wolf [6] that any non-trivial Discrete Memoryless Channel implies OT. Thus, it suffices to show that there are some correlation devices that can be used to simulate a non-trivial DMC; the existence of any such correlation device would consequently imply the existence of OT.

Definition 9 (Discrete Memoryless Channel). *A* discrete memoryless channel *is characterized by an input alphabet $\mathcal{A}_X$, an output alphabet $\mathcal{A}_Y$, and a set of conditional probability distributions $P_{y|x}$ for each $x \in \mathcal{A}_X$.*

Note that the binary symmetric channel with probability of error $p \in [0,1]$ is a special case of DMC with $\mathcal{A}_X = \mathcal{A}_Y = \{0,1\}$, and the conditional probabilities $P_{1|0} = P_{0|1} = p$, and $P_{0|0} = P_{1|1} = 1-p$.

Wolf and Wullschleger [20] considered the problem of two parties with access to correlated random variables X, and Y trying to simulate a DMC characterized by the conditional probabilities $P_{Y|X}$. A correlated equilibrium γ of a strategic two player game corresponds to an identical situation. The two players have access to two correlated random variables that are defined by the randomized advice about what action each one of them should take in the game. Given access to the correlation device, the players can simulate a discrete memoryless channel as described in Fig. 4.

To send bit $d \in \{0,1\}$ from party A to party B:

1. Both players get advice according to γ, and use rejection sampling to make sure that the pair of advice they get is an element $(a,b) \in \{a_0, a_1\} \times \{b_0, b_1\}$ for some actions a_0, a_1 of player A and b_0, b_1 of player B. They use the correlation device for γ multiple times, until both a_0 and a_1 appear in the list of advice received by player A.
2. Party A erases some advice from her list to make a_0 and a_1 equiprobable, and sends to B the index i of the first occurrence of a_d in her list.
3. Party B outputs d', such that $b_{d'}$ is the i-th advice in the list of player B.

Fig. 4. Simulating a DMC when given access to some correlation device for a CE γ.

This procedure simulates a DMC defined by the conditional probabilities $P_{y|x}$ corresponding to the CE restricted by the rejection sampling to $\{a_0, a_1\} \times \{b_0, b_1\}$; for example the probability of receiving 0 after sending 1 is $P_{0|1} = \gamma(a_1, b_0)/(\gamma(a_1, b_0) + \gamma(a_1, b_1))$. Note that this procedure in general does not simulate the binary symmetric channel.[8] However, we show that for non-trivial CE the properties of the associated DMC are good enough to imply OT.

We are interested in DMCs that are non-trivial in the following sense.

Definition 10 (Crépeau *et al.*[6]). *We call a channel $P_{Y|X}$* trivial *if there exist, after removal of all redundant input symbols, partitions of the (remaining) ranges $\mathcal{X}$ of X and $\mathcal{Y}$ of Y , $\mathcal{X} = \mathcal{X}_1 \cup \ldots \cup \mathcal{X}_n, \mathcal{Y} = \mathcal{Y}_1 \cup \ldots \cup \mathcal{Y}_n$, and channels $P_{Y_i|X_i}$, where the ranges of X_i and Y_i are $\mathcal{X}_i$ and $\mathcal{Y}_i$, respectively, such that*

$$P_{Y|X=x}(y) = \begin{cases} P_{Y_i|X_i=x}(y) & \textit{if } x \in \mathcal{X}_i, y \in \mathcal{Y}_i, \\ \\ 0 & \textit{if } x \in \mathcal{X}_i, y \in \mathcal{Y}_j, i \neq j \end{cases}$$

holds and such that the capacity of the channel $P_{Y_i|X_i}$ is 0 *for all i.*

[8] Some non-trivial CE indeed give rise to well-known channels. For example the CE from previous section corresponds to the Z-channel.

The following lemma justifies the use of correlated equilibria outside the convex-hull of NE to simulate non-trivial DMCs.

Lemma 3. *Let $\varGamma$ be a strategic game, and γ some correlated equilibrium of $\varGamma$. If γ is a CE of $\varGamma$ outside the convex hull of NE, then there exist a pair of actions $a_i \neq a_j$ of player A and a pair of actions $b_k \neq b_l$ of player B, such that the restriction of γ to $\{a_i, a_j\} \times \{b_k, b_l\}$ allows to simulate a non-trivial DMC.*

Proof. Recall that $P_{b|a} = \gamma(a,b)/(\gamma(a,b_k)+\gamma(a,b_l))$ for any $(a,b) \in \{a_i,a_j\} \times \{b_k,b_l\}$. Since γ is not a CHNE of $\varGamma$, there must exist actions $a_i \neq a_j$ of player A and $b_k \neq b_l$ of player B, such that

$$P_{b_k|a_i} \neq P_{b_k|a_j}, \text{ or } P_{b_l|a_i} \neq P_{b_l|a_j} \tag{7.1}$$

(or else γ is a completely mixed NE of $\varGamma$). We want to show that the conditional probabilities $P_{b|a}$ characterize a channel with non-zero capacity. Condition (7.1) ensures that it is never the case that $P_{b_k|a_i} = P_{b_l|a_i} = P_{b_k|a_j} = P_{b_l|a_j} = 1/2$. Thus, the resulting DMC does not have entropy 1 (i.e. it has non-zero capacity). On the other hand, we need to show that the resulting DMC has enough entropy to be non-trivial, i.e., that it is not a perfect channel or a channel outputting always the same symbol. It suffices to show that among the tuples of actions consistent with the condition (7.1) we can in fact select the actions a_i, a_j and b_k, b_l so that at most one of the conditional probabilities $P_{b|a}$ is zero. Equivalently, we instead show that it is possible to select the actions where at most one of $\gamma(a,b)$ is equal to zero. We call a candidate *bad* if it has more than one 0. Note that no bad candidate has $\gamma(a_i,b_k) = \gamma(a_j,b_k) = 0$ or $\gamma(a_i,b_l) = \gamma(a_j,b_l) = 0$, since then $P_{b_k|a_i} = P_{b_k|a_j} = 0$ and $P_{b_l|a_i} = P_{b_l|a_j} = 1$, respectively $P_{b_l|a_i} = P_{b_l|a_j} = 0$ and $P_{b_k|a_i} = P_{b_k|a_j} = 1$. So, bad candidates are either of the row type, $\gamma(a_i,b_k) = \gamma(a_i,b_l) = 0$ or $\gamma(a_j,b_k) = \gamma(a_j,b_l) = 0$, or the diagonal type, $\gamma(a_i,b_k) = \gamma(a_j,b_l) = 0$ or $\gamma(a_i,b_l) = \gamma(a_j,b_k) = 0$. One can use this to show that it holds for any two actions a_i and a_j that the residual distribution given a_i and a_j is either identical or disjoint. I.e., either $\gamma(a_i,b_k)/\gamma(a_i) = \gamma(a_j,b_k)/\gamma(a_j)$ for all b_k or $\gamma(a_i,b_k) = 0 \vee \gamma(a_j,b_k) = 0$ for all b_k. This shows that the distribution is a sum of product distributions, each a NE, i.e., a CHNE. There are more details in the full version. □

The following theorem characterizes DMCs with respect to the possibility of their use to create unconditional OT:

Theorem 6 (Crépeau *et al.*[6]). *Let two players A and B be connected by a non-trivial channel $P_{Y|X}$. Then, for any $\alpha > 0$, there exists a protocol for unconditionally secure OT from A to B with failure probability at most α, where the number of uses of the channel is of order $O(\log(1/\alpha)^{2+\varepsilon})$ for any $\varepsilon > 0$. Trivial channels, on the other hand, do not allow for realizing OT in an unconditional way.*

Lemma 3 together with the above result of Crépeau *et al.* [6] give the sought proof of Theorem 5.

Acknowledgements

The authors would like to thank for discussions and useful comments to Ronen Gradwohl, Jonathan Katz, Noam Livne, Peter Bro Miltersen, and Margarita Vald.

References

[1] Mikhail J. Atallah, Marina Blanton, Keith B. Frikken, and Jiangtao Li, *Efficient correlated action selection*, Financial Cryptography (Giovanni Di Crescenzo and Aviel D. Rubin, eds.), Lecture Notes in Computer Science, vol. 4107, Springer, 2006, pp. 296–310.

[2] Robert J. Aumann, *Subjectivity and correlation in randomized strategies*, Journal of Mathematical Economics **1** (1974), no. 1, 67–96.

[3] Robert J. Aumann and Sergiu Hart, *Long cheap talk*, Econometrica **71** (2003), no. 6, 1619–1660.

[4] Imre Bárány, *Fair distribution protocols or how the players replace fortune*, Mathematics of Operations Research **17** (1992), no. 2, 327–340.

[5] Ran Canetti, *Universally composable security: A new paradigm for cryptographic protocols*, IACR Cryptology ePrint Archive **2000** (2000), 67.

[6] Claude Crépeau, Kirill Morozov, and Stefan Wolf, *Efficient unconditional oblivious transfer from almost any noisy channel*, SCN (Carlo Blundo and Stelvio Cimato, eds.), Lecture Notes in Computer Science, vol. 3352, Springer, 2004, pp. 47–59.

[7] Yevgeniy Dodis, Shai Halevi, and Tal Rabin, *A cryptographic solution to a game theoretic problem*, CRYPTO (Mihir Bellare, ed.), Lecture Notes in Computer Science, vol. 1880, Springer, 2000, pp. 112–130.

[8] Joseph Farrell and Matthew Rabin, *Cheap talk*, The Journal of Economic Perspectives **10** (1996), no. 3, 103–118.

[9] Ronen Gradwohl, Noam Livne, and Alon Rosen, *Sequential rationality in cryptographic protocols*, FOCS, IEEE Computer Society, 2010, pp. 623–632.

[10] Adam Groce and Jonathan Katz, *Fair computation with rational players*, EUROCRYPT (David Pointcheval and Thomas Johansson, eds.), Lecture Notes in Computer Science, vol. 7237, Springer, 2012, pp. 81–98.

[11] Russell Impagliazzo, *A personal view of average-case complexity*, Structure in Complexity Theory Conference, 1995, pp. 134–147.

[12] Russell Impagliazzo and Michael Luby, *One-way functions are essential for complexity based cryptography (extended abstract)*, FOCS, IEEE Computer Society, 1989, pp. 230–235.

[13] Hemanta K. Maji, Manoj Prabhakaran, and Amit Sahai, *On the computational complexity of coin flipping*, FOCS, IEEE Computer Society, 2010, pp. 613–622.

[14] John Nash, *Non-cooperative games*, Annals of mathematics **54** (1951), no. 2, 286–295.

[15] Martin J. Osborne and Ariel Rubinstein, *A course in game theory*, MIT Press, 1994.

[16] Rafael Pass and Abhi Shelat, *Renegotiation-safe protocols*, ICS (Bernard Chazelle, ed.), Tsinghua University Press, 2011, pp. 61–78.

[17] Vanessa Teague, *Selecting correlated random actions*, Financial Cryptography (Ari Juels, ed.), Lecture Notes in Computer Science, vol. 3110, Springer, 2004, pp. 181–195.

[18] ______, *Problems With Coordination in Two-Player Games: Comment on "Computational Complexity and Communication"*, Econometrica **76** (2008), no. 6, 1559–1564.

[19] Amparo Urbano and Jose E. Vila, *Computational complexity and communication: Coordination in two–player games*, Econometrica **70** (2002), no. 5, 1893–1927.

[20] Stefan Wolf and Jürg Wullschleger, *Zero-error information and applications in cryptography*, Information Theory Workshop, 2004. IEEE, October 2004, pp. 1 – 6.