

Limits on the Power of Cryptographic Cheap Talk

Pavel Hubáček*

Jesper Buus Nielsen[†]

Alon Rosen[‡]

August 14, 2013

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 [16]. 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 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.

*Department of Computer Science, Aarhus University, Denmark. Partially supported by the European Research Commission Starting Grant 279447. Part of this work was supported by ISF grant no. 334/08. The author acknowledges support from the Danish National Research Foundation and The National Science Foundation of China (under the grant 61061130540) for the Sino-Danish Center for the Theory of Interactive Computation, within which part of this work was performed. Email: hubacek@cs.au.dk.

[†]Department of Computer Science, Aarhus University, Denmark. Supported by the European Research Commission Starting Grant 279447. Supported by the Danish Council for Independent Research via DFF Starting Grant 10-081612. Email: jbn@cs.au.dk.

[‡]School of Computer Science, Herzliya IDC, 46150, Israel. Work supported by ISF grant no. 334/08. E-mail: alon.rosen@idc.ac.il.

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.¹

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 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.

¹For 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.

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 5.1 (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 3.2 (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 equilibria² (CHNE) which can be implemented via cheap talk only if one-way functions exist.

Theorem 6.1 (informal). *If the payoff of all non-trivial convex hull Nash equilibrium 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 7.1 (informal). *If the payoff of all correlated equilibrium 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.

²Note 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.

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

For introductory text on game theory see Osborne and Rubinstein [17]. 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 Barany [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 [19], and subsequently Atallah *et al.* [1] gave a protocol for the general problem of correlated element selection achieving 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 [21], aiming for a similar result to Dodis *et al.* [7]. However, Teague [20] showed that their approach is flawed. An alternative solution concept for analyzing game theoretical properties of cryptographic protocols was suggested by Pass and Shelat [18].

2 Preliminaries and Definitions

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

Definition 2.1 (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 $u : A_1 \times A_2 \rightarrow \mathbb{R}^2$ 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 2.2 (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 2.3 (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

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

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., $U_i(\gamma|a_i) = \left(\sum_{a_{-i} \in A_{-i}} \gamma(a_i, a_{-i}) \right)^{-1} \sum_{a_{-i} \in A_{-i}} \gamma(a_i, a_{-i}) u_i(a_i, a_{-i})$.

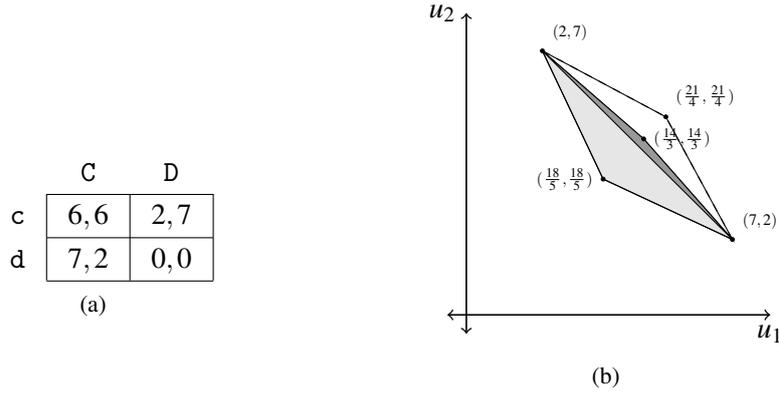


Figure 1: (a) the game of Chicken (b) the utility profiles achievable by its CE.

Definition 2.4 ((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 $\text{NE}(\Gamma)$, $\text{CHNE}(\Gamma)$, and $\text{CE}(\Gamma)$ 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.³

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).

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

Definition 2.5 ((strict) Pareto dominance, weak Pareto optimality). Let Γ be a strategic game, and $\gamma, \gamma' \in \text{CE}(\Gamma)$. 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 $\gamma^* \in \text{CE}(\Gamma)$ is *weakly Pareto optimal* if there exists no $\gamma' \in \text{CE}(\Gamma)$ 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 $\gamma, \gamma' \in \text{CE}(\Gamma)$, 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.

³As 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.

Definition 3.1 (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 3.2. *For any $k \in \mathbb{N}$. If $k > 3$, then there exists a $k \times k$ strategic game Γ with a correlated equilibrium $\gamma \in \text{CE}(\Gamma) \setminus \text{CHNE}(\Gamma)$, s.t. every $\gamma' \in \text{CE}(\Gamma)$ 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. 3.3) that are exploited when we extend it.

Definition 3.3 (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 $\gamma' \in \text{CE}(\Gamma)$ 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 3.4. *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 $\gamma \in \text{CE}(\Gamma) \setminus \text{CHNE}(\Gamma)$, s.t. every $\gamma' \in \text{CE}(\Gamma)$ 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. 3.3. 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., $v' = \max(U_A(\sigma_A^*), U_B(\sigma_B^*))$, 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. \square

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

Lemma 3.5. *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 $3f < (e - c)$.⁴ Consider the $k \times k$ game $\Lambda_k = (A = \{a_1, \dots, a_k\}, B = \{b_1, \dots, b_k\}, u)$ with the utility function $u : A \times B \rightarrow \mathbb{R}^2$ 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

Figure 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 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 3.5. 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 Appendix D.

Claim 3.6. *Any such probability distribution $\gamma_k \in \Delta(A \times B)$ is a correlated equilibrium of Λ_k .*

⁴The 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 3.5.

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

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

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. \square

We also justify that our counterexample is minimal in the sense that there is no 2×2 strategic game that could be used in the context of Lemma 3.4 (for proof see Appendix D).

Lemma 3.7. *There is no 2×2 strategic game that satisfies the extensibility criterion.*

Since among the 2×2 strategic games only the games with two pure Nash equilibria can have a correlated equilibrium improving the utility profiles achieved by CHNE, we get the following corollary.

Corollary 3.8. *If Γ is a 2×2 strategic game, then every $\gamma \in \text{CE}(\Gamma)$ is NE-punishable.*

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 Appendix A.

Our core object of study is so-called computational cheap talk, simultaneous move (CTSM) games. A CTSM game without types⁵ 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 be 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*.

⁵In this work we focus on finite games without types. This was the setting of the founding paper in [7], and giving a full characterization of this setting turns out to be plentiful technically involved. We leave it as future work to study cheap talk games with types. However, our model, and some of our results, apply to the more general setting of infinite games with types. Technically, in terms of the formal model in Appendix A, we always consider the CTSM game (T_1, T_2, A_1, A_2, b) , where $T_1 = T_2 = \{\top\}$ and $b = 1$, and we always analyze games assuming the empty common prior C_0 for Γ , which always output $(\top, \varepsilon, \top, \varepsilon)$.

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 Appendix A. 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. The details and their motivation has therefore been deferred to Appendix A. 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 don't 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, $\Gamma = \{\Gamma_R\}_{R \in \mathbb{N}}$, where all Γ_R have the same corresponding strategic game, and where Γ_R has a fixed last round in round R . This allows to easily define ETF for each Γ_R , and we then say that σ is ETF if there exists R_0 such that it is ETF for all Γ_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 Γ^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 $\Gamma_p = \{\Gamma^{(k)} = \Gamma^{p(k)}\}_{k \in \mathbb{N}}$. A strategy $\sigma = \{\sigma^{(k)}\}_{k \in \mathbb{N}}$ for Γ_p is one where $\sigma^{(k)}$ is a strategy for $\Gamma^{(k)}$. A strategy σ for Γ is clearly poly-time. We say that σ is a computational NE for Γ_p if there exists negligible ε such that $\sigma^{(k)}$ is an $\varepsilon(k)$ -NE for $\Gamma^{(k)}$. 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 Γ_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 4.1 (Trivial, Minicrypt, and Cryptomania utility profiles). Let Γ be a strategic game, and $v \in \mathbb{R}^2$ be a utility profile achieved by some $\gamma \in \text{CE}(\Gamma)$.

- We call v a *trivial utility profile* if there exists $\sigma \in \text{NE}(\Gamma)$ 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 5.1. *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 Appendix B.1. 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 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 given in Appendix B.1.

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 (see Appendix E).

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.⁶

Consider the CTSM game specified by $\Gamma = (A_1, A_2)$, where $A_1 = \{c, d\}$, $A_2 = \{C, D\}$, and the utility function u is given in Fig. 3. The probability distribution selecting (c, D) and (d, C) with equal probability

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

Figure 3: A game of Chicken.

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 6.1. *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. 4.

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 $y_i = 1$.

Figure 4: Protocol for weak coin-flip given a cheap talk implementation of a specific CHNE.

The following statements are logically equivalent.

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) .

⁶The 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.

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 σ_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. 4 then they both output the same bit b , and it is 0 or 1 with equal probability. Since σ is a computational equilibrium of (Γ, C_θ) , 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. \square

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]. See Appendix E.2 for the details. 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 $\Gamma_{\text{chicken}} = (A_1, A_2, u)$, where $A_1 = \{c, d\}$, $A_2 = \{C, D\}$ and the utility function u is given by:

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

Theorem 7.1. *If there exists a computational NE σ for the CTSM game corresponding to Γ_{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 $\text{view}_i(\sigma) = \text{view}_i(\Gamma, \sigma, C)$ 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 c or C, then P_i outputs $b_i = 1$. If σ_i plays d or 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 (d, C) as much as it plays (c, D); call the probability of playing each of these α . Let β denote the probability that σ plays (c, 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 C is $\frac{\beta}{\alpha+\beta}15 + \frac{\alpha}{\alpha+\beta}6$. If P_2 would switch to D when σ_1 says to play 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 (P_1, P_2) in π is distributed as follows:

prob.	P_1	P_2
$\frac{1}{3}$	0	1
$\frac{1}{3}$	1	0
$\frac{1}{3}$	1	1

One can show that an expected constant number of samples from this distribution is sufficient to implement random Rabin OT, see Appendix C 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

$$[\text{view}_1 | b_1 = 1 \wedge b_2 = 1] \approx [\text{view}_1 | b_1 = 1 \wedge b_2 = 0]$$

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

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 $[\text{view}_1 | b_1 = 1 \wedge b_2 = 1]$ and $[\text{view}_1 | b_1 = 1 \wedge b_2 = 0]$ with non-negligible probability, i.e., $|\Pr[D([\text{view}_1 | b_1 = 1 \wedge b_2 = 1]) = 1] - \Pr[D([\text{view}_1 | b_1 = 1 \wedge b_2 = 0]) = 1]|$ is non-negligible. Since we work with non-uniform complexity, we can assume that it is always the case that $\Pr[D([\text{view}_1 | b_1 = 1 \wedge b_2 = 1]) = 1] \geq \Pr[D([\text{view}_1 | b_1 = 1 \wedge b_2 = 0]) = 1]$. Now consider the following strategy σ_1^* . It plays like σ_1 , except that if σ_1 recommends to play c, then σ_1^* switches to d when $D(\text{view}_1) = 1$, where view_1 is the view of P_1 . Note that σ_1 recommending to play c is logically equivalent to $b_1 = 1$. I.e., $\text{view}_1 \in \{[\text{view}_1 | b_1 = 1 \wedge b_2 = 1], [\text{view}_1 | b_1 = 1 \wedge b_2 = 0]\}$. 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 (c, C) and (c, D) separately. If the joint play is (c, C), then we switch with probability $\Pr[D([\text{view}_1 | b_1 = 1 \wedge b_2 = 1]) = 1]$, for a gain of $\Pr[D([\text{view}_1 | b_1 = 1 \wedge b_2 = 1]) = 1](21 - 15)$. If the joint play is (c, D), then we switch with probability $\Pr[D([\text{view}_1 | b_1 = 1 \wedge b_2 = 0]) = 1]$, for a gain of $\Pr[D([\text{view}_1 | b_1 = 1 \wedge b_2 = 0]) = 1](0 - 6)$. This gives a total gain of $6(\Pr[D([\text{view}_1 | b_1 = 1 \wedge b_2 = 1]) = 1] - \Pr[D([\text{view}_1 | b_1 = 1 \wedge b_2 = 0]) = 1])$. 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. \square

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 7.2. *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 7.3 (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 Wullschlegler [22] 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. 5.

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 .

Figure 5: 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.⁷ 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 7.4 (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 , $\mathcal{X} = \mathcal{X}_1 \cup \dots \cup \mathcal{X}_n$, $\mathcal{Y} =$

⁷Some non-trivial CE indeed give rise to well-known channels. For example the CE from previous section corresponds to the Z-channel.

$\mathcal{Y}_1 \cup \dots \cup \mathcal{Y}_n$, 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

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

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

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

Lemma 7.5. *Let Γ be a strategic game, and γ some correlated equilibrium of Γ . If γ is a CE of Γ 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 Γ , there must exist actions $a_i \neq a_j$ of player A and $b_k \neq b_l$ of player B, such that

$$P_{b_k|a_i} \neq P_{b_k|a_j}, \text{ or } P_{b_l|a_i} \neq P_{b_l|a_j} \quad (7.1)$$

(or else γ is a completely mixed NE of Γ). 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 for it to be a non-trivial DMC, 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.

Assume that it is not possible to select the actions such that at most one of $\gamma(a, b)$ is equal to zero. Then all the candidate tuples (a_i, a_j, b_k, b_l) consistent with condition (7.1) fall into one of the following types:

1. $\gamma(a_i, b_k) = \gamma(a_i, b_l) = 0$ or $\gamma(a_j, b_k) = \gamma(a_j, b_l) = 0$,
2. $\gamma(a_i, b_k) = \gamma(a_j, b_l) = 0$ or $\gamma(a_i, b_l) = \gamma(a_j, b_k) = 0$,
3. there is exactly one $(a, b) \in \{a_i, a_j\} \times \{b_k, b_l\}$ s.t. $\gamma(a, b)$ is non-zero.

Note that the tuples such that $\gamma(a_i, b_k) = \gamma(a_j, b_k) = 0$ or $\gamma(a_i, b_l) = \gamma(a_j, b_l) = 0$ cannot be consistent with (7.1), 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 $P_{b_k|a_i} = P_{b_k|a_j} = 1$.

We give an algorithm that allows to decompose γ into a convex combination of NE of Γ . We call actions a_i and a_j of player A *disjoint* if there is no action b_k of player B such that $\gamma(a_i, b_k)$ and $\gamma(a_j, b_k)$ are simultaneously non-zero (so, every action a_i which is played with non-zero probability is not disjoint with itself). First, we prove the following claim about actions that are not disjoint.

Claim 7.6. *Let $a_i \neq a_j$ be two actions of player A that are not disjoint. Then the conditional distribution $\gamma(a_i)$ of γ restricted to the row of a_i and the conditional distribution $\gamma(a_j)$ of γ restricted to the row of a_j are identical.*

Proof. Assume for a contradiction that the two conditional distributions are not identical. Then there exists a tuple (a_i, a_j, b_l, b_m) such that $P_{b_l|a_i} \neq P_{b_l|a_j}$, or else $\gamma(a_i)$ and $\gamma(a_j)$ are identical as shown in Claim 7.7. If $b_k = b_l$ or $b_k = b_m$ (where b_k is the action of player B such that $\gamma(a_i, b_k)$ and $\gamma(a_j, b_k)$ are

simultaneously non-zero), then we are done since (a_i, a_j, b_l, b_m) is a tuple consistent with (7.1), and it is neither one of the above three possible types (since $\gamma(a_i, b_k)$ and $\gamma(a_j, b_k)$ are simultaneously non-zero). Otherwise, if (a_i, a_j, b_l, b_m) is one of the above three types, let without loss of generality b_l be such that $\gamma(a_i, b_l) \neq 0$ or $\gamma(a_j, b_l) \neq 0$. Then (a_i, a_j, b_k, b_l) is also a candidate tuple not of one of the three possible types. \square

Claim 7.7. *Let a_i, a_j be actions of player A such that for every pair of actions b_k, b_l of player B it holds that $P_{b_k|a_i} = P_{b_k|a_j}$. Then the two conditional distributions of $\gamma(a_i)$ and $\gamma(a_j)$ are identical.*

Proof. If the assumption of the claim holds, then $\gamma(a_i, b_k)$ and $\gamma(a_j, b_k)$ are either simultaneously non-zero or simultaneously zero for all actions b_k of player B. We can therefore restrict ourselves to actions b_k such that $\gamma(a_i, b_k)$ is non-zero (let these be w.l.o.g. the first actions of player B). We show the claim by induction.

The base case is given by the assumption of the claim, since for b_1, b_2 we have that $P_{b_1|a_i} = P_{b_1|a_j}$ and $P_{b_2|a_i} = P_{b_2|a_j}$. Assume now that it holds for every $b_m \in \{b_1 \dots, b_{n-1}\}$ that

$$\frac{\gamma(a_i, b_m)}{\gamma(a_i, b_1) + \dots + \gamma(a_i, b_{n-1})} = \frac{\gamma(a_j, b_m)}{\gamma(a_j, b_1) + \dots + \gamma(a_j, b_{n-1})}.$$

Since $\gamma(a_i, b_m)$ and $\gamma(a_j, b_m)$ are both non-zero, this is equivalent to

$$\frac{\gamma(a_i, b_1) + \dots + \gamma(a_i, b_{n-1})}{\gamma(a_i, b_m)} = \frac{\gamma(a_j, b_1) + \dots + \gamma(a_j, b_{n-1})}{\gamma(a_j, b_m)}.$$

For b_m and b_n it also holds that $P_{b_m|a_i} = P_{b_m|a_j}$, so we get

$$\frac{\gamma(a_i, b_1) + \dots + \gamma(a_i, b_{n-1})}{\gamma(a_i, b_m)} + \frac{\gamma(a_i, b_m) + \gamma(a_i, b_n)}{\gamma(a_i, b_m)} = \frac{\gamma(a_j, b_1) + \gamma(a_j, b_{n-1})}{\gamma(a_j, b_m)} + \frac{\gamma(a_j, b_m) + \gamma(a_j, b_n)}{\gamma(a_j, b_m)}.$$

Therefore

$$\frac{\gamma(a_i, b_1) + \dots + \gamma(a_i, b_n)}{\gamma(a_i, b_m)} = \frac{\gamma(a_j, b_1) + \gamma(a_j, b_n)}{\gamma(a_j, b_m)},$$

that shows the inductive step. \square

The algorithm to decompose γ into a convex combination of NE goes as follows. If there is any action of player A left which is played with non-zero probability, take one such action a_i . Let $B(a_i)$ be the set of the actions b_k of player B such that $\gamma(a_i, b_k) > 0$, and let $A(a_i)$ be all the actions of player A not disjoint with a_i . As shown in Claim 7.6, for all $a_j \in A(a_i)$ the conditional distribution of γ restricted to a_j is identical to the conditional distribution of γ restricted to a_i . Thus, $\gamma(a_j, b_m) = 0$ for all $a_j \in A(a_i)$ and $b_m \notin B(a_i)$. Moreover, $\gamma(a_h, b_k) = 0$ for all $a_h \notin A(a_i)$ and $b_k \in B(a_i)$, as $\gamma(a_i, b_k) \neq 0$ and a_i and a_h are disjoint. If we restrict γ to $A(a_i) \times B(a_i)$ and normalize, then γ is a (possibly mixed) NE of the restricted game. Remove all the actions in $A(a_i)$ and $B(a_i)$ from the action space and repeat this procedure again.

The above algorithm terminates after finitely many steps since Γ is a finite game, and it decomposes γ into a convex combination of NE, with the weights being the inverse of the normalization factors. This is a contradiction with γ being a CE outside the convex-hull of NE. One can thus always find some actions in the support of γ that allow simulating a non-trivial DMC. \square

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

Theorem 7.8 (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 $O(\log(1/\alpha)^{2+\epsilon})$ for any $\epsilon > 0$. Trivial channels, on the other hand, do not allow for realizing OT in an unconditional way.*

Lemma 7.5 together with the above result of Crépeau *et al.* [6] give the sought proof of Theorem 7.2.

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] Peter Bro Miltersen, Jesper Buus Nielsen, and Nikos Triandopoulos, *Privacy-enhancing auctions using rational cryptography*, CRYPTO (Shai Halevi, ed.), Lecture Notes in Computer Science, vol. 5677, Springer, 2009, pp. 541–558.
- [15] Hervé Moulin and J-P Vial, *Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon*, International Journal of Game Theory **7** (1978), no. 3, 201–221.
- [16] John Nash, *Non-cooperative games*, Annals of mathematics **54** (1951), no. 2, 286–295.
- [17] Martin J. Osborne and Ariel Rubinstein, *A course in game theory*, MIT Press, 1994.

- [18] Rafael Pass and Abhi Shelat, *Renegotiation-safe protocols*, ICS (Bernard Chazelle, ed.), Tsinghua University Press, 2011, pp. 61–78.
- [19] Vanessa Teague, *Selecting correlated random actions*, Financial Cryptography (Ari Juels, ed.), Lecture Notes in Computer Science, vol. 3110, Springer, 2004, pp. 181–195.
- [20] ———, *Problems With Coordination in Two-Player Games: Comment on "Computational Complexity and Communication"*, *Econometrica* **76** (2008), no. 6, 1559–1564.
- [21] Amparo Urbano and Jose E. Vila, *Computational complexity and communication: Coordination in two-player games*, *Econometrica* **70** (2002), no. 5, 1893–1927.
- [22] Stefan Wolf and Jürg Wullschleger, *Zero-error information and applications in cryptography*, Information Theory Workshop, 2004. IEEE, October 2004, pp. 1 – 6.

A Computational Cheap Talk Simultaneous-Move Games

In this appendix we give full details for our game theoretical model and solution concepts. We intend to follow [9] as closely as possible, but where [9] start with a clean purely game theoretic notion and then dirty it up to handle computational issues and the fact that communication protocols are considered, we instead start with a clean definition of what a communication protocol is and then define the game theoretic notions around this skeleton. We get a less general definition, but also, we fell, a more precisely specified and workable definition.

A.1 Discussion of Basic Model Choices

Our goal is to analyze games which use cryptography in the cheap talk phase. We will therefore have to restrict the set of strategies to the set of efficient strategies or include the price of computation into the utility function. We consider the inclusion of the price of computation into the utility function as the purest solution and probably the one with best predictive power in general. However, including the price of computation also has the potential to considerably complicate analysis, possibly taking focus away from the more interesting issues. We have therefore instead chosen to restrict the strategy space to the efficient ones. A consequence of this design choice is the by now well-known one that we need to include a negligible slack parameter ε into the solution concepts. For instance, instead of NE we need to consider an ε -NE for a negligible ε . This is so because an efficient strategy will always have some small probability of breaking the applied cryptography, e.g., by just making a guess at the keys of the other agents.⁸

We will model an efficient strategy as a strategy which can be implemented in strict polynomial time. This deviates slightly from the usual approach in cryptography, which uses expected polynomial time. However, for every expected polynomial time strategy with utility u there exists another strategy which is strict polynomial time and which gets utility $u' = u - \varepsilon$ for a negligible ε . Since we already committed to having a negligible slack parameter in our model, little is therefore lost. It seems, however, that it buys us a lot in simplicity of definition. Namely, if we went for expected polynomial time, we would have to formalize what it means for a strategy to be expected polynomial time, with the problems this give: Is a strategy expected polynomial time, if terminates in expected polynomial time given that the strategy of the other parties are fixed, or should it guarantee to terminate in expected polynomial time no matter the strategy of the other parties? The first choice is clearly too liberal, as the other parties might strategically deviate if it could make some other player become inefficient, which in practice would mean it would never terminate the computation. The second choice is too restrictive, as *any* strategy would

⁸A solution concept including the price of computation could handle this by having the expected utility, ε , be too small compared to the price of the extra computation needed to make the guess.

include inefficient adversarially chosen strategies which makes the strategy in consideration inefficient only because it breaks the applied cryptography. Consider, e.g., your strategy in a joint strategy which runs a secure coin-flipping protocol in each round and terminates if the coin comes out 0. If the other agent has unbounded computing power, it might break the coin-flip protocol to make it always output 1 and hence make the game run forever. Yet, we would like to be able to analyze exactly such protocols. Should we then require that an efficient strategy guarantees to terminate in expected polynomial time no matter the strategies of the other parties, as long as they are efficient? or as long as they are efficient and rational? These would appear to lead to recursive definitions! One can, however, resolve this and give a satisfactory definition, but we do not know of a definition simple enough for definition and analysis that it is worth the complication, in particular as very little seems to be gained by picking expected polynomial time over polynomial time. So, we go for simplicity. Also, many of the motivating settings we want to analyze have a cheap talk phased followed by an exogenous deadline for a forced move in the game which determines the utility, say a bid in an auction or turning a truck off collision course or not. These settings do not pair well with an expected running time of the cheap talk phase, but impose a notion of worst case running time. Therefore choosing expected running time might in fact lead to a loss of generality.

Inspired by the model in [14] we use a non-uniform notion of efficiency as opposed to the definition in [7]. The motivation is the non-uniform definition allows to model the computational setting using a sequence of games with *finite* action spaces. The approach in [7] gives infinite action spaces, the set of Turing machines. Having a finite strategy space is sometimes convenient in formalizing solution concepts and in analysis, so we prefer this choice.

A final design choice, which is usual in computational games, is that we make mixing of strategies explicit. I.e., we do not let the strategy space be the set of all probability distributions on the actions in the action space, as is usual in game theory, we only allow those probability distributions which can be efficiently implemented, as motivated above. We model this by making a strategy by a fixed algorithm which takes as input a uniformly random randomizer, which is used to mixed the chosen actions.

A.2 Cheap Talk Simultaneous-Move Games

An m -round cheap talk simultaneous-move (CTSM) game is a tuple

$$\begin{aligned} \Gamma = (T_1, T_2, A_1, A_2, u : T_1 \times T_2 \times A_1 \times A_2 \rightarrow \mathbb{R}^2, \\ S_1, S_2, m \in \mathbb{N}, b \in \{1, 2\}, \\ \{R^{(j)}, S^{(j)}, \Sigma^{(j)}\}_{j \in [m-1]}, R_1^{(m)}, \Sigma_1^{(m)}, R_2^{(m)}, \Sigma_2^{(m)}). \end{aligned}$$

Before the game player i is given a type $t_i \in T_i$. At the same time player i is possibly given some signal $s_i \in S_i$ about the type of the other party. We consider a Bayesian game, where the types and messages are drawn using some known distribution C , the so-called common prior. Then, there are $m - 1$ rounds of cheap talk, as specified below, followed by a simultaneous move game, where player 1 plays $a_1 \in A_1$ and player 2 plays $a_2 \in A_2$. The utility is given by $(u_1, u_2) = u(t_1, t_2, a_1, a_2)$, where u_i is the utility of player i , i.e., the cheap talk does not explicitly affect the utility.

For $i = 1, 2$ we use P_{-i} to denote the other player than P_i , i.e. P_{3-i} . We use the same notation when indexing strategies. If $\sigma = (\sigma_1, \sigma_2)$ we use $(\sigma_{-i}, \sigma_i^*)$ to denote (σ_1, σ_2^*) when $i = 2$ and (σ_1^*, σ_2) when $i = 1$.

A.3 Structure of a game

The structure of a game is as follows.

- In rounds $j \in [m - 1]$ the player takes turn making a move. In round $j = 1$ it is P_b who moves. Let $P : [m - 1] \rightarrow \{P_1, P_2\}$ denote the corresponding player function, where $P(j) = P_b$ iff j is odd.

- For rounds $j \in [m-1]$ the action space of $P(j)$ is $S^{(j)}$. We call $S^{(j)}$ the message space.
- In round $j = m$ the action space of each P_i is A_i .
- We call $R^{(j)}$ the randomizer space of round j .
- We define spaces $S_i^{(<j)}$, where $S_i^{(<j)}$ is the recall space of P_i about rounds earlier than round j .
- If $j = 1$, then $S_i^{(<j)} = S_i$.
- If $1 < j \leq m$ and $P(j-1) = P_i$, then $S_i^{(<j)} = S_i^{(<j-1)} \times R^{(j-1)}$ and $S_{-i}^{(<j)} = S_{-i}^{(<j-1)} \times S^{(j-1)}$.
- If $1 \leq j < m$ and $P(j) = P_i$, then $\Sigma^{(j)}$ is the set of possible strategies of player i in that round. It is a subset of the functions with functionality $T_i \times S_i^{(<j)} \times R^{(j)} \rightarrow S^{(j)}$. For notational convenience, we let $\Sigma_i^{(j)} = \Sigma^{(j)}$ and $\Sigma_{-i}^{(j)} = \{\top\}$ when $P(j) = P_i$ and $j < m$.
- If $j = m$, then $\Sigma_i^{(m)}$ is the set of possible strategies of player i in the last round. It is a subset of the functions with functionality $T_i \times S_i^{(<m)} \times R_i^{(m)} \rightarrow A_i$.

A.4 Structure of a strategy

We let $\Sigma_i = \times_{j=1}^m \Sigma_i^{(j)}$ and $\Sigma = \Sigma_1 \times \Sigma_2$. We let $\Sigma^{(j)} = \Sigma_1^{(j)} \times \Sigma_2^{(j)}$. Note that for $j < m$ and $P_i = P(j)$, we have that $\Sigma^{(j)}$ and $\Sigma_i^{(j)}$ essentially are identical. A strategy profile for a game is an element $(\sigma_1, \sigma_2) \in \Sigma_1 \times \Sigma_2$, i.e., $\sigma_i = (\sigma_i^{(1)}, \dots, \sigma_i^{(m-1)}, \sigma_i^{(m)}) \in \times_{j=1}^m \Sigma_i^{(j)}$. The outcome of the game is defined via letting each player follow its strategy. In more detail:

- A round $j \in [m-1]$ where $P_i = P(j)$ and a strategy $\sigma^{(j)} \in \Sigma^{(j)}$ defines a randomized *round function* $\text{Rnd}_{\Gamma, \sigma^{(j)}}^{(j)} : T_1 \times S_1^{(<j)} \times T_2 \times S_2^{(<j)} \rightarrow T_1 \times S_1^{(<j+1)} \times T_2 \times S_2^{(<j+1)}$, as follows: Let the input be $(t_1, s_1^{(<j)}, t_2, s_2^{(<j)})$. It first computes $s^{(j)} \leftarrow \sigma_i^{(j)}(t_i, s_i^{(<j)}, r_i^{(j)})$ for a uniformly random $r_i^{(j)} \in R_i^{(j)}$. Then it lets $s_{-i}^{(<j+1)} = (s_{-i}^{(<j)}, s^{(j)})$ and $s_i^{(<j+1)} = (s_i^{(<j)}, r_i^{(j)})$, and outputs $(t_1, s_1^{(<j+1)}, t_2, s_2^{(<j+1)})$.
- A round m and a strategy profile $\sigma^{(m)} = (\sigma_1^{(m)}, \sigma_2^{(m)}) \in \Sigma_1^{(m)} \times \Sigma_2^{(m)}$ defines a randomized *round function* $\text{Rnd}_{\Gamma, \sigma^{(m)}}^{(m)} : T_1 \times S_1^{(<m)} \times T_2 \times S_2^{(<m)} \rightarrow A_1 \times A_2$, as follows: Let the input be $(t_1, s_1^{(<m)}, t_2, s_2^{(<m)})$. It computes $a_1 \leftarrow \sigma_1^{(m)}(s_1^{(<m)}, r_2^{(m)})$ and $a_2 \leftarrow \sigma_2^{(m)}(s_2^{(<m)}, r_1^{(m)})$ for uniformly random $r_i^{(m)} \in R_i^{(m)}$, and outputs (a_1, a_2) .
- For a strategy profile $(\sigma_1, \sigma_2) \in \Sigma_1 \times \Sigma_2$, we define a randomized function $\text{Play}_{\Gamma, \sigma} : T_1 \times S_1 \times T_2 \times S_2 \rightarrow A_1 \times A_2$ given by

$$\text{Play}_{\Gamma, \sigma}(\cdot) = \text{Rnd}_{\Gamma, \sigma^{(m)}}^{(m)} \circ \dots \circ \text{Rnd}_{\Gamma, \sigma^{(2)}}^{(2)} \circ \text{Rnd}_{\Gamma, \sigma^{(1)}}^{(1)} .$$

A.5 Playing a Game

A common prior C for a game is a distribution on $T_1 \times S_1 \times T_2 \times S_2$. We use $D[T_1 \times S_1 \times T_2 \times S_2]$ to denote the set of such distributions. To have a fully specified play of a game we need to specify the common prior. We call the actions (a_1, a_2) played in the last round the *outcome* of the game, and we define further properties of the game via the expected utility of the players given the outcome.

Definition A.1 (Expected Utility). Let Γ be a CTSM game, let σ be a strategy profile for Γ , and let C be a common prior for Γ . We use $\text{Play}_{\Gamma, \sigma}(C)$ to denote the random variable described as follows: sample $(t_1, s_1, t_2, s_2) \leftarrow C$, sample $(a_1, a_2) \leftarrow \text{Play}_{\Gamma, \sigma}(t_1, s_1, t_2, s_2)$ and output (a_1, a_2) . For fixed Γ and C

and for $i = 1, 2$ we define a utility function $u_i : \Sigma \rightarrow \mathbb{R}$, by letting $u_i(\sigma)$ be the expected value of u_i in $(u_1, u_2) = u(t_1, a_1, t_2, a_2)$ when $(a_1, a_2) \leftarrow \text{Play}_{\Gamma, \sigma}(t_1, s_1, t_2, s_2)$ and $(t_1, s_1, t_2, s_2) \leftarrow C$. When we need to make the game and the common prior explicit we write $u_i(\sigma, \Gamma, C)$.

For later use we will need a notion of conditioned utility, i.e., expected utility given that some event happens. For our purpose an event (at round j) will be a set of possible configurations of the protocol, i.e., a subset $E \subset T_1 \times S_1^{(<j)} \times T_2 \times S_2^{(<j)}$. Conditioned utility is just the expected utility given that the event occurs.

Definition A.2 (View). For a game Γ , a strategy σ and a common prior C , let $\text{view}(\sigma, \Gamma, C)$ denote the random variable described as follows: sample $(t_1, s_1, t_2, s_2) \leftarrow C$, sample $(t_1, s_1^{(<m)}, t_2, s_2^{(<m)}) \leftarrow \text{Rnd}_{\Gamma, \sigma^{(m)}}^{(m-1)}(\dots \text{Rnd}_{\Gamma, \sigma^{(2)}}^{(2)}(\text{Rnd}_{\Gamma, \sigma^{(1)}}^{(1)}(t_1, s_1, t_2, s_2)) \dots)$, sample $(a_1, a_2) \leftarrow \text{Rnd}_{\Gamma, \sigma^{(m)}}^{(m)}(t_1, s_1^{(<m)}, t_2, s_2^{(<m)})$, and output $(t_1, s_1^{(<m)}, a_1, t_2, s_2^{(<m)}, a_2)$.

Definition A.3 (Conditional Utility). Let Γ be a game, let σ be a strategy for Γ and let C be a common prior for Γ . An event in Γ is a subset $E \subset T_1 \times S_1^{<m} \times A_1 \times T_2 \times S_2^{<m} \times A_2$. For fixed Γ and C and for $i = 1, 2$ we define a conditional utility function $u_i \wedge E : \Sigma \rightarrow \mathbb{R}$, by letting $u_i(\sigma \wedge E)$ be the expected value of u_i in $(u_1, u_2) = \alpha u(a_1, t_2, a_2, t_1)$ when $(t_1, s_1^{(<m)}, a_1, t_2, s_2^{(<m)}, a_2) \leftarrow \text{view}(\sigma, \Gamma, C)$ and conditioned on $(t_1, s_1^{(<m)}, a_1, t_2, s_2^{(<m)}, a_2) \in E$, where $\alpha = \Pr[(t_1, s_1^{(<m)}, a_1, t_2, s_2^{(<m)}, a_2) \in E]$. If $\alpha = 0$, then we let $u_i(\sigma \wedge E) = 0$. When we need to make the game and the common prior explicit we write $u_i(\sigma, \Gamma, C \wedge E)$.

A.6 Nash Equilibrium

Definition A.4 (ε -Nash Equilibrium). Let Γ be a CTSM game and let C be a common prior for Γ . Let $\varepsilon \in \mathbb{R}$. We say that $\sigma \in \Sigma$ is an ε -NE for (Γ, C) if for both $i = 1, 2$ and all strategies $\sigma_i^* \in \Sigma_i$ it holds that $u_i(\sigma_i^*, \sigma_{i-1}) \leq u_i(\sigma) + \varepsilon$. We use $\text{NE}^{(\varepsilon)}(\Gamma, C)$ to denote the set of ε -NE for (Γ, C) .

A.7 Empty-Threat Freeness

We want to refine the notion of NE by requiring that one cannot use empty threats for stability. The underlying assumption is that empty threats will be called if the other player would gain from you not carrying through the threat, and hence a NE with an empty threat would not be stable. Traditionally the notion of sub-game perfect equilibrium has been used for ruling out empty threat, but it is too strong for this purpose and is problematic to define in a computational setting. We therefore go for an explicit notion of empty-threat freeness.

One cannot threaten in a 1-round simultaneous move game, as a threat is a future action meant to deter a currently possible action of your opponent. The only reasonable notion of “threat” in a 1-round simultaneous move game would be to threaten, prior to the game, that you will play in a particular way, as to make your opponent respond optimally to your claimed play. The actual play being simultaneous, your threat will, however, be empty if the resulting two strategies are not in equilibrium: in the actual play you would deviate to your optimal strategy instead, and your opponent knows this. We will therefore equate the empty-threat free equilibria with the NE when we consider 1-round simultaneous move games.

Definition A.5 (Empty-Threat Free). We say that $\sigma \in \Sigma$ is an ε -ETFE for 1-round CTSM game Γ and common prior C for Γ if σ is an ε -NE for (Γ, C) . We use $\text{ETFE}^{(\varepsilon)}(\Gamma, C) = \text{NE}^{(\varepsilon)}(\Gamma, C)$ to denote the set of ε -ETFE for (Γ, C) .

To handle games with several rounds, we first define a notion of sub-game, where the first cheap talk rounds are swallowed by the common prior. Namely, if two players reach some internal round in a game, they still have the same types and they did not yet pick actions for the simultaneous move game. I.e., they are essentially in a CTSM game—only their signals are changed by prior messages. Then we recursively define what an empty threat is and then what empty-threat freeness is.

Definition A.6 (Sub-Game). For an m -round CTSM game Γ with $m > 1$, define an \bar{m} -round game $\Gamma^{(\geq 2)} = \bar{\Gamma}$, where $\bar{m} = m - 1$, $\bar{T}_i = T_i$, $\bar{A}_i = A_i$, $\bar{u} = u$, $\bar{S}_i = S_i^{(<2)}$, $\bar{b} = 3 - b$, $\bar{R}_i^{(j)} = R_i^{(j+1)}$, and $\bar{S}_i^{(j)} = S_i^{(j+1)}$, and $\bar{\Sigma}_i^{(j)} = \Sigma_i^{(j+1)}$. In general, for $\rho \geq 3$, let $\Gamma^{(\geq \rho)} = (\Gamma^{(\geq \rho-1)})^{\geq 2}$.

Definition A.7 (Sub-Strategy). For an m -round strategy profile σ for an m -round CTSM game with $m > \rho$, let $\sigma^{(\geq \rho)} = (\sigma^{(\rho)}, \dots, \sigma^{(m)})$.

Note that if σ is a strategy profile for Γ , then $\sigma^{(\geq 2)}$ is a strategy profile for $\Gamma^{(\geq 2)}$.

Definition A.8 (Sub-Common Prior). For an m -round CTSM game Γ with $m > 1$, a common prior $C \in \mathcal{D}[T_1 \times S_1 \times T_2 \times S_2]$ and a round function $R = R_{\Gamma, \sigma^{(1)}}^{(1)} : T_1 \times S_1 \times T_2 \times S_1 \rightarrow T_1 \times S_1^{(<2)} \times T_2 \times S_2^{(<2)}$ for the first round, let $\sigma^{(1)}(C) = R_{\Gamma, \sigma^{(1)}}^{(1)}(C)$ be the common prior from $\mathcal{D}[T_1 \times S_1^{(<2)} \times T_2 \times S_2^{(<2)}]$ given by sampling $(t_1, s_1, t_2, s_2) \leftarrow C$, sampling $(t_1, s_1^{(<2)}, t_2, s_2^{(<2)}) \leftarrow R(t_1, s_1, t_2, s_2)$ and outputting $(t_1, s_1^{(<2)}, t_2, s_2^{(<2)})$. In general, let $\sigma^{(\leq 1)} = \sigma^{(1)}$ and for $\rho > 1$, let $\sigma^{(\leq \rho)}(C) = \sigma^\rho(\sigma^{\leq \rho-1}(C))$.

Definition A.9 (Conditional Common Prior). For a common prior $C \in \mathcal{D}[T_1 \times S_1 \times T_2 \times S_2]$ and an event $E \subseteq T_1 \times S_1 \times T_2 \times S_2$ we use $\alpha = \Pr[E|C]$ to denote the probability that $C \in E$ and if $\alpha > 0$ we use $C|_E$ to denote the distribution of C given E .

Note that if C is a common prior for Γ and σ is a strategy profile for Γ , then $\sigma^{(1)}(C)$ is a common prior for $\Gamma^{(\geq 2)}$.

We now discuss and motivate the upcoming formal definition of empty threat. For our purpose an empty threat posed by me in a NE is a part of my future strategy which I do not currently play in the NE and which I would not play should you call my bluff by deviating in a way making the threatening strategy active. Basically, you would demonstrate the existence of such a future empty threat posed by me by demonstrating a deviation by you which would make me deviate from playing the supposedly empty threat. We will use this as a definition by saying that an empty threat exists iff you can come up with such a constructive demonstration that it exists. In [9] the definition for games with imperfect information is only hinted. It is suggested that it would be reasonable to require that the deviation used to demonstrate the existence of a future empty threat be observable by the other party such that the strategic response can be done on basis of observing your deviation. This appears reasonable, but is rather cumbersome to define. And, it turns out that an essentially equivalent and much simpler definition can be given. We will only require that a demonstrating deviation only is allowed to generate a strategic response when it actually occurs happens, i.e., I'm not allowed to react to a possibly unobservable deviation in a way which would make we change strategy even when it did not occur. These definitions are almost equivalent in our settings: in a communication protocol I can always choose to send you a single extra bit whenever I do my deviation, and have you react only when you see this bit. The sending of the bit would be a deviation in it self, so you would act according to our definition, and the bit would at the same time make the deviation observable by the other party as suggested by [9].

The above definition still requires some qualification though. Consider the case where you make me deviate from playing the supposedly empty threat by posing a yet more future empty threat punishing me if and when I play my supposedly empty threat. This would not be a credible demonstration, as I could conceivably in turn call your bluff instead of abstaining from executing my threat. We therefore require that your demonstration that I will abstain from the threat leads to an empty-threat free future play, which is why we need to do a recursive definition.

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 this deviation.

All in all, a credible demonstration that I'm posing an empty threat on you would therefore be a deviation by you which has the property that it leads to an empty-threat free future play in which you

have higher utility. This is not a full definition if your deviation could lead to distinct empty-threat free continuations in which you have different expected utilities. Following Gradwohl *et al.* [9] we require that you have higher expected utility in *all* empty-threat free continuations which can result from your deviation. Think of this as allowing me to respond strategically to your deviation by picking an empty-threat free punishment if one exists. This seems reasonable, as the whole setting of avoiding empty threats assumes that I will respond strategically to your deviations by possibly modifying my future play. Note that a consequence of this definition is that a seemingly empty threat is not considered empty if there exists a credible threat with the same effect: If I have two buttons, one which blows us both up, and one which only blows up you, with no cost to me, then a threat that I blow us both up if you do not pay me a dollar is not an empty threat, as the play where I blow up only you has the same effect.

Definition A.10. Let $\sigma_i^{(j)}, \delta_i^{(j)} : T_i \times S_i^{(<j)} \times R^{(j)} \rightarrow S^{(j)}$ be strategies for player i for round j and let $E \subseteq T_i \times S_i^{(<j)} \times R^{(j)}$ be an event observable by player i in round j . We say that $\sigma_i^{(j)}$ and $\delta_i^{(j)}$ play identically given E if $\sigma_i^{(j)}(e) = \delta_i^{(j)}(e)$ for all $e \in E$. We write $\sigma_i^{(j)} =_E \delta_i^{(j)}$.

Definition A.11 (Empty-Threat Free). We say player 1 is facing a first round ε -empty threat in an m -round CTSM game setting (Γ, C, σ) if there exists an event E for player 1 for round 1 and a deviation $\sigma_1^{(1)*} \in \Sigma_1^{(1)}, \sigma_1^{(1)*} =_{\bar{E}} \sigma_1^{(1)}$ such that for all $\tilde{\sigma} \in \text{ETFE}^\varepsilon(\Gamma^{(\geq 2)}, (\sigma_1^{(1)*}, \sigma_2^{(1)})(C)|_E)$, i.e., in the ETF plays in the sub-game where E occurred, it holds that player 1 gets at least ε more than when E occurs in σ , both of them weighed by the probability that E actually occurs. Formally, let

$$\begin{aligned} \text{FET}_1(\Gamma, C) = \{ \sigma \in \Sigma(\Gamma) \mid \exists E, \sigma_1^{(1)*} \in \Sigma_1^{(1)}(\Gamma) : \sigma_1^{(1)*} =_{\bar{E}} \sigma_1^{(1)} \wedge \\ \forall \tilde{\sigma} \in \text{ETFE}^\varepsilon(\Gamma^{(\geq 2)}, \sigma^{(1)*}(C)|_E) : \\ \Pr[E | \sigma^{(1)*}(C)] u_1(\tilde{\sigma}, \Gamma^{(\geq 2)}, \sigma^{(1)*}(C)|_E) > u_1(\sigma, \Gamma, C \wedge E) + \varepsilon, \\ \text{where } \sigma^{(1)*} = (\sigma_1^{(1)*}, \sigma_2^{(1)}) \}. \end{aligned} \quad (\text{A.1})$$

We define that player 2 is facing a first round ε -empty threat symmetrically, and we let $\text{FET}_2(\Gamma, C)$ denote the strategies where player 2 faces a first round ε -empty threat. We say that (Γ, C, σ) is an ε -first-round-empty-threat free NE if it is a NE and it is not the case that a player is facing a first round ε -empty threat. Let

$$\begin{aligned} \text{FETF}^\varepsilon(\Gamma, C) &= \Sigma(\Gamma) \setminus (\text{FET}_1(\Gamma, C) \cup \text{FET}_2(\Gamma, C)), \\ \text{FETFE}^\varepsilon(\Gamma, C) &= \text{FETF}^\varepsilon(\Gamma, C) \cap \text{NE}^\varepsilon(\Gamma, C). \end{aligned} \quad (\text{A.2})$$

We say that (Γ, C, σ) is an ε -empty-threat free NE if it is empty-threat free in all rounds, i.e., if all sub-games are empty-threat free in their first round. For $m > 2$, let

$$\sigma \in \text{ETFE}^\varepsilon(\Gamma, C) \equiv \sigma \in \text{FETFE}^\varepsilon(\Gamma, C) \wedge \sigma^{(\geq 2)} \in \text{ETFE}^\varepsilon(\Gamma^{(\geq 2)}, \sigma^{(1)}(C)). \quad (\text{A.3})$$

Notice that since we define ETFE^ε for m -round games via ETFE^ε for $(m-1)$ -round games and we have 1-round games as basis, our notion is indeed well-defined.

It is instructive to compare the notions ETFE and NE. In a NE, a player will not have incentive to deviate, assuming the other player keeps playing according to the NE. In an ETFE, a player can have no incentive to deviate, even if he believes that the other party will respond strategically to his deviation.

A.8 Computational Version

We want to consider games playable by computers and therefore set the messages and randomizers to be bit strings. To allow the use of cryptography, we introduce a security parameter κ , and we allow that the strategies depend on κ . We need to restrict the players to efficient strategies, i.e., the running time

of (implementing) the strategies should be polynomial in κ . Formally we capture this by restricting the strategy space to strategies which are efficient.

A polynomial family of CTSM games is a tuple

$$\Gamma = (T_1, T_2, A_1, A_2, u : T_1 \times T_2 \times A_1 \times A_2 \rightarrow \mathbb{R}^2, m : \mathbb{N} \rightarrow \mathbb{N}, b \in \{1, 2\}, c : \mathbb{N} \rightarrow \mathbb{N}),$$

where m and c are monotonously increasing and bounded by a polynomial. For a specific value $\kappa \in \mathbb{N}$ of the security parameter, the family defines a CTSM game

$$\Gamma(\kappa) = (T_1, T_2, A_1, A_2, u, S_1, S_2, m(\kappa), b, \{R^{(j)}, S^{(j)}, \Sigma^{(j)}\}_{j \in [m-1]}, R_1^{(m)}, \Sigma_1^{(m)}, R_2^{(m)}, \Sigma_2^{(m)}),$$

where $S_1 = S_2 = R^{(j)} = S^{(j)} = R_1^{(m)} = R_2^{(m)} = \{0, 1\}^*$ and $\Sigma^{(j)}$, $\Sigma_1^{(m)}$ and $\Sigma_2^{(m)}$ are the subsets of strategies which are computable in complexity $c(\kappa)$. We here fix the complexity measure to be that the given function can be computed by a Boolean circuit of size $c(\kappa)$, but the definition readily applies to other complexity measures.

Definition A.12. We use Σ to denote the function which maps $\kappa \in \mathbb{N}$ to the strategy space of $\Gamma(\kappa)$. I.e., $\Sigma(\kappa)$ is the strategy space of $\Gamma(\kappa)$. A family of strategies for Γ is function σ on \mathbb{N} , where $\sigma(\kappa) \in \Sigma(\kappa)$.

Definition A.13. An efficient common prior for (T_1, T_2) is a family of common priors $C : \mathbb{N} \rightarrow \mathcal{D}[T_1 \times \{0, 1\}^* \times T_2 \times \{0, 1\}^*]$, which can be sampled in non-uniform polynomial time. An efficient common prior for a game Γ is an efficient common prior for the type space (T_1, T_2) of Γ .

Definition A.14. Let Γ be a polynomial family of CTSM games, let σ be a strategy for Γ and let C be an efficient common prior for Γ . Let $\varepsilon : \mathbb{N} \rightarrow \mathbb{R}$. We say that σ is an ε -ETFE for Γ and C if it holds for all κ that $\sigma(\kappa)$ is an $\varepsilon(\kappa)$ -ETFE for $\Gamma(\kappa)$ and $C(\kappa)$.

It is customary in cryptography to require that the complexity of a proposed protocol must be some fixed polynomial, but that it should tolerate attacks which can be implemented in any polynomial complexity. To lift our definition to this setting we need to be able to see a strategy for a given complexity restriction also as a strategy for a more liberal complexity restriction. At the same time we will also consider a strategy as a strategy for a game where more rounds are allowed. The reason is that we do not want the stability of a strategy to depend too strongly on the exact number of rounds available. As an example, a protocol which is stable if there are exactly 12 rounds of cheap talk available, but not if 13 rounds of cheap talk are available, would probably not be stable in practice, where there is no small *a priori* bound on the number of rounds of communication possible, even if an exogenous deadline for the game is given.

Definition A.15. Let $\Gamma = (T_1, T_2, A_1, A_2, u, m, b, c)$ be a polynomial family of CTSM games. We call $\Gamma' = (T_1, T_2, A_1, A_2, u, m', b, c')$ an *extension* of Γ , and write $\Gamma' \geq \Gamma$, if $m'(\kappa) \geq m(\kappa)$ and $c'(\kappa) \geq c(\kappa)$ for all κ . Given a strategy $\sigma \in \Sigma(\Gamma)$, we can consider it as a strategy $\sigma' \in \Sigma(\Gamma')$ which computes the action a_i after round $m(\kappa)$ and then just sends the empty string in the extra rounds and ignores the messages from the other party in the extra rounds. More formally, we have that $\sigma_i'^{(j)}(\kappa) = \sigma_i^{(j)}(\kappa)$ for $i = 1, 2$ and $j = 1, \dots, m(\kappa) - 1$. For the extra rounds $j = m(\kappa), \dots, m'(\kappa) - 1$ where $P(j) = P_i$, the strategy $\sigma_i'^{(j)}(\kappa)$ simply sends the empty string ε . In the final round $\sigma_i'^{(m')}(\kappa)$ plays as $\sigma_i^{(m)}(\kappa)$, using as input to $\sigma_i^{(m)}(\kappa)$ only the recall of the rounds $< m$, i.e., $(\sigma_i'^{(m')}(\kappa))(\cdot, s_i^{(< m')}, \cdot) = (\sigma_i^{(m)}(\kappa))(\cdot, s_i^{(< m)}, \cdot)$, where $s_i^{(< m')} = (s_i^{(1)}, \dots, s_i^{(m-1)}, s_i^{(m)}, \dots, s_i^{(m'-1)})$ and $s_i^{(< m)} = (s_i^{(1)}, \dots, s_i^{(m-1)})$.

We think of computational CTSM game as a polynomial family of games where the exact complexity are left open, only requiring them to be polynomials. All we need to specify is who speaks first.

Definition A.16 (Computational Game). A computational CTSM game is a tuple $\Gamma = (T_1, T_2, A_1, A_2, u : T_1 \times T_2 \times A_1 \times A_2 \rightarrow \mathbb{R}^2, b \in \{1, 2\})$.

Definition A.17 (Computational ETFE). Let $\Gamma = (T_1, T_2, A_1, A_2, u, b)$ be a computational CTSM game. Let C be a efficient common prior for Γ . We say that σ is a computational ETFE for Γ and C if there exist c and m and negligible ε such that σ is an ε -ETFE for the polynomial family of games $\Gamma = (T_1, T_2, A_1, A_2, u, m, b, c)$ and C and for all $\Gamma' \geq \Gamma$ there exist a negligible ε' such that σ is an ε' -ETFE for Γ' . We write $\sigma \in \text{CETFE}(\Gamma, C)$.

A.9 Finite, Type-Free Games

We define a class of CTSM games, which we call finite, type-free CTSM (FTFCTSM) games. These are just games where the parties have no types and where they have a finite set of actions in the simultaneous move games. Such a game is specified by $\Gamma = (A_1, A_2, u, b)$ with $|A_1|, |A_2| \in \mathbb{N}$, and specifies the computational CTSM game $\Gamma = (T_\emptyset, T_\emptyset, A_1, A_2, u, b)$, where $T_\emptyset = \{\top\}$. To study such games we can restrict our study to $b = 0$, as we can always transpose the utility function, so a FTFCTSM is given by just the strategic game $\Gamma = (A_1, A_2, u)$. The only difference is that it specifies an extensive-form game where some cheap talk is included before the simultaneous move game is played.

B Some Relations to Traditional Game Theory

In this section we prove some relations to traditional game theory. We will prove that in a computational cheap talk game with no communication and no common prior, the notion of computational NE is very closely related to the traditional notion of NE, in that a computational NE must be negligibly close to a NE in terms of statistical distance on the distribution of the play. We also show that in a computational cheap talk game with no communication and an unrestricted common prior, a computational NE lies negligibly close to a CE. The intuition is that the common prior is the mediator mechanism of the CE. Finally, we show that in a computational cheap talk game with no communication and a common prior which only contains common information, a computational NE lies negligibly close to a CHNE. All of these results are used in later sections, but are factored out here as their share a lot of details in their proofs.

Recall the traditional notions of NE, CE and CHNE defined in Sect. 3. We define similar notions for families of games and a slack parameter. We require that this slack parameter goes to 0.

Definition B.1 (Correlated Equilibrium). An ε correlated equilibrium for a strategic game (A_1, A_2, u) with utility profile (v_1, v_2) is a probability distribution γ on $A_1 \times A_2$ where $U(\gamma) = (v_1, v_2)$ and where $i \in \{1, 2\}$ and for every $a_i, a_i^* \in A_i$ it holds that $\sum_{a_{-i} \in A_{-i}} \gamma(a_i, a_{-i}) u_i(a_i^*, a_{-i}) \leq \sum_{a_{-i} \in A_{-i}} \gamma(a_i, a_{-i}) u_i(a_i, a_{-i}) + \varepsilon$. A CE for (A_1, A_2, u) is a 0-CE.

Definition B.2 (Nash Equilibrium). An (ε) -NE σ for a strategic game (A_1, A_2, u) with utility profile (v_1, v_2) is a (ε) -CE σ , where σ is a product distribution on $A_1 \times A_2$.

Definition B.3 (Convex Hull Nash Equilibrium). A (ε) -CHNE γ for a strategic game (A_1, A_2, u) with utility profile (v_1, v_2) is a (ε) -CE γ , where γ is a convex combination of (ε) -NE.

Definition B.4 $((\varepsilon, \delta)$ -CE). An (ε, δ) -CE for a strategic game (A_1, A_2, u) with utility profile (v_1, v_2) is an infinite sequence $\{\gamma_\kappa\}_{\kappa=1}^\infty$ such that each γ_κ is a probability distribution on $A_1 \times A_2$, for all κ it holds that $U(\gamma_\kappa)$ is within distance $\delta(\kappa)$ from (v_1, v_2) and that $\forall i \in \{1, 2\} \forall a_i, a_i' \in A_i$ it holds that $\sum_{a_{-i} \in A_{-i}} \gamma_\kappa(a_i, a_{-i}) u_i(a_i, a_{-i}) \geq \sum_{a_{-i} \in A_{-i}} \gamma_\kappa(a_i, a_{-i}) u_i(a_i', a_{-i}) - \varepsilon(\kappa)$.

Definition B.5 $((\varepsilon, \delta)$ -NE). An (ε, δ) -NE for a strategic game (A_1, A_2, u) with utility profile (v_1, v_2) is an (ε, δ) -CE for (A_1, A_2, u) with utility profile (v_1, v_2) , where each γ_κ is a product distribution.

Definition B.6 ((ε, δ) -CHNE). An (ε, δ) -CHNE for a strategic game (A_1, A_2, u) with utility profile (v_1, v_2) is an (ε, δ) -CE for (A_1, A_2, u) with utility profile (v_1, v_2) , where each γ_κ is a convex combination of product distributions.

Notice an important difference between the definition of ε -CHNE and (ε, δ) -CHNE: we require that a ε -CHNE is a convex combination of ε -NE, and we only require that a (ε, δ) -CHNE is a convex combination of product distributions, which is potentially much weaker. For proof purposes it turns out that it is convenient to have a notion even stronger than the two above.

Definition B.7 ((ε, δ) -near CHNE). An (ε, δ) -near CHNE for a strategic game (A_1, A_2, u) with utility profile (v_1, v_2) is an (ε, δ) -CE for (A_1, A_2, u) with utility profile (v_1, v_2) , where each γ_κ is a convex combination of strategies which have statistical distance at most ε to a NE, i.e., there exists a set S_κ , a probability distribution p_κ on S_κ and ε -NEs $\sigma_\kappa(s)$ such that

$$\gamma_\kappa = \sum_{s \in S_\kappa} p_\kappa(s) \sigma_\kappa(s) .$$

Lemma B.8. *If A_1 and A_2 are finite, then, if there exists an (ε, δ) -CE, (ε, δ) -NE, or (ε, δ) -CHNE for $\Gamma = (A_1, A_2, u)$ with utility profile (v_1, v_2) , where $\lim_{\kappa \rightarrow \infty} \varepsilon(\kappa) = \lim_{\kappa \rightarrow \infty} \delta(\kappa) = 0$, then there also exists a CE, NE, or CHNE, respectively, for (A_1, A_2, u) with utility profile (v_1, v_2) .*

Proof. We first show it for CE. We have an infinite sequence $\gamma_1, \dots, \gamma_\kappa, \dots$, each a probability distribution on $A_1 \times A_2$. Since the set of probability distributions on $A_1 \times A_2$ form a compact space, any infinite sequence in the set contains a convergent sub-sequence, converging to a point in the set. Let $\lambda_1, \dots, \lambda_\kappa, \dots$ denote this sub-sequence. It is easy to see that $\lambda_1, \dots, \lambda_\kappa, \dots$ is again an (ε', δ') -CE for (A_1, A_2, u) with utility profile (v_1, v_2) for ε' and δ' such that $\lim_{\kappa \rightarrow \infty} \varepsilon'(\kappa) = \lim_{\kappa \rightarrow \infty} \delta'(\kappa) = 0$. Let λ be the probability distribution on $A_1 \times A_2$ to which $\lambda_1, \dots, \lambda_\kappa, \dots$ converge. It is easy to see that λ is a CE for (A_1, A_2, u) with utility profile (v_1, v_2) : it is certain a distribution on $A_1 \times A_2$, and the utility of any switch from a_i to a'_i will go to 0, so the utility of any switch from a_i to a'_i in λ will be 0.

The proof for NE goes as above. All that has to be checked is that if a sequence of product distributions converge, then it converges to a product distribution. This can be seen by looking at the distance from the points in the sequence to the space of product distributions. It is always 0. This will therefore be true also at the limit. Since the space of product distributions is closed, it follows that the limit point is a product distribution.

The proof for CHNE goes like above, but there are some extra complications. We first show the result for (ε, δ) -near CHNE and then reduce to this case.

Claim B.9. *Assume that there exists an (ε, δ) -near CHNE γ with utility profile (v_1, v_2) , with $\lim_{\kappa \rightarrow \infty} \varepsilon(\kappa) = \lim_{\kappa \rightarrow \infty} \delta(\kappa) = 0$. Then there also exist a CHNE for (A_1, A_2, u) with utility profile (v_1, v_2) .*

As above, we can without loss of generality assume that γ actually converges to some CE H with utility profile (v_1, v_2) , otherwise use compactness to pick an infinite sub-sequence with this property. What remains is to show that H is a CHNE, i.e., a convex combination of NE. Since γ converges to H we know there exists α such that $\lim_{\kappa \rightarrow \infty} \alpha(\kappa) = 0$ and such that the statistical distance from H to γ_κ is at most $\alpha(\kappa)$. We use this later, but first have to derive another distance bound.

By definition, we have that each $\sigma_\kappa(s)$ can be written as $\sigma_\kappa(s) = (1 - \varepsilon(\kappa))N_\kappa(s) + \varepsilon(\kappa)A_\kappa(s)$, where $N_\kappa(s)$ is a NE and where $A_\kappa(s)$ is some arbitrary strategy profile and where ε goes to 0. Hence

$$\begin{aligned} \gamma_\kappa &= \sum_s p_\kappa(s) \sigma_\kappa(s) \\ &= \sum_s p_\kappa(s) ((1 - \varepsilon(\kappa))N_\kappa(s) + \varepsilon(\kappa)A_\kappa(s)) \\ &= (1 - \varepsilon(\kappa)) \sum_s p_\kappa(s) N_\kappa(s) + \varepsilon(\kappa) \sum_s p_\kappa(s) A_\kappa(s) . \end{aligned}$$

By definition $\sum_s p_\kappa(s)N_\kappa(s)$ is a CHNE, so the statistical distance from γ_κ to the space of CHNE is at most $\varepsilon(\kappa)$.

By combining the above two bounds, we get that the statistical distance from H to the space of CHNE is at most $(\alpha + \varepsilon)(\kappa)$. Since $\lim_{\kappa \rightarrow \infty} (\alpha + \varepsilon)(\kappa) = 0$, it follows that the statistical distance from H to the space of CHNE is 0. Since CHNE is a closed space, it follows that H is a CHNE.

We can then conclude the proof by showing the following:

Claim B.10. *If there exists an (ε, δ) -CHNE γ with utility profile (v_1, v_2) and $\lim_{\kappa \rightarrow \infty} \varepsilon(\kappa) = \lim_{\kappa \rightarrow \infty} \delta(\kappa) = 0$. Then there exists an (v, ψ) -near CHNE γ with utility profile (v_1, v_2) and $\lim_{\kappa \rightarrow \infty} v(\kappa) = \lim_{\kappa \rightarrow \infty} \psi(\kappa) = 0$.*

As above, we can without loss of generality assume that γ actually converges to some CE H with utility profile (v_1, v_2) . By assumption we have that each γ_κ is a convex combination of product distributions, i.e.,

$$\gamma_\kappa = \sum_{s \in S_\kappa} p(s) \sigma_\kappa(s)$$

for some set S_κ and some probability distribution p on S_κ and each $\sigma_\kappa(s)$ being a product distribution. What we need is that each $\sigma_\kappa(s)$ is actually statistically close to a NE. This is not always the case, but it turns out we can massage γ to get this property without changing the utility profile. We first get rid of every $\sigma_\kappa(s)$ which is not ε' -NE for some some ε' . Then we use that being ε' -NE means that you are close to a NE, when ε' gets small enough.

Define the bad subset $B_\kappa \subset S_\kappa$ to be the $s \in S_\kappa$ for which $p(s) > 0$ and for which $\sigma_\kappa(s)$ is not a $\sqrt{\varepsilon(\kappa)}$ -NE. It is easy to see that $\Pr[s \in B_\kappa] \leq \sqrt{\varepsilon(\kappa)}$, as $\Pr[s \in B_\kappa] > \sqrt{\varepsilon(\kappa)}$ would imply that γ_κ is not a $\sqrt{\varepsilon(\kappa)}$ -NE, a contradiction. So, if we define D_κ to be γ_κ , where we let $D_\kappa(s) = \sigma_\kappa(s)$ for $s \notin B_\kappa$ and $D_\kappa(s) = N$ for $s \in B_\kappa$, for some fixed NE N of Γ , then each D_κ has statistical distance at most $\sqrt{\varepsilon(\kappa)}$ to γ_κ . Hence D is a (ϕ, ψ) -CHNE for $\Gamma = (A_1, A_2, u)$ with utility profile (v_1, v_2) , for $\phi = \sqrt{\varepsilon}$, $\psi = \delta + 2c\sqrt{\varepsilon}$, where $c = \max_{i=1,2; (a_1, a_2) \in A_1 \times A_2} |u_i(a_1, a_2)|$ is a constant. Note that $\lim_{\kappa \rightarrow \infty} \phi(\kappa) = \lim_{\kappa \rightarrow \infty} \psi(\kappa) = 0$. Furthermore, D has the extra property that $D_\kappa(s)$ is a $\phi(\kappa)$ -NE for all s and all large enough κ .

Now note that if for all $v > 0$, there exists $\phi > 0$ such that if a strategy profile E is a ϕ -NE for Γ , then the statistical distance from E to the nearest NE is at most v , then D is a (v, ψ) -near CHNE for $\Gamma = (A_1, A_2, u)$ with utility profile (v_1, v_2) with $\lim_{\kappa \rightarrow \infty} v(\kappa) = \lim_{\kappa \rightarrow \infty} \psi(\kappa) = 0$, which would mean we would be done with the proof.

Our claim has the form $\forall v > 0 \exists \phi > 0 \forall E (E \notin \phi\text{-NE}(\Gamma) \vee E \in N_v(\text{NE}(\Gamma)))$, where $N_v(\text{NE}(\Gamma))$ is the strategy profiles with statistical distance at most v to a NE. The negation is therefore equivalent to $\exists v > 0 \forall \phi > 0 \exists E (E \in \phi\text{-NE}(\Gamma) \wedge E \notin N_v(\text{NE}(\Gamma)))$. Now pick the $v > 0$ given by this formula and for $i = 1, 2, 3, \dots$ pick some E_i given by this formula at $\phi = 1/i$, i.e., $E_i \in \phi\text{-NE}(\Gamma)$ and $E_i \notin N_v(\text{NE}(\Gamma))$. It follows that $E = \{E_\kappa\}$ does not converge to a NE, as the points keep having distance v the space of NE. Yet, $E = \{E_\kappa\}$ is an ϕ -NE for $\phi(\kappa) = 1/\kappa$, which clearly goes to 0, so we know via above arguments that E converges to a NE, a contradiction. \square

We then relate to computational cheap talk games. Let (A_1, A_2, u) be a finite strategic game and let Γ be the corresponding cheap talk game. We call a strategy σ for Γ *silent* if both parties send the empty string in all rounds at all security levels. We call a common prior C_\emptyset for Γ *empty* if it always outputs $(\top, \varepsilon, \top, \varepsilon)$, where ε denotes the empty string. We say that a common prior C_{com} for Γ is common information if it always outputs the same to both parties, i.e., it is a probability distribution over strings (\top, s, \top, s) .

Theorem B.11. *Let (A_1, A_2, u) be a finite strategic game and let Γ be the corresponding cheap talk game. If there exist a common prior C for Γ and a silent strategy σ for (Γ, C) such that σ is a computational NE for (Γ, C) and has utility profile (v_1, v_2) , then there exist a CE γ for (A_1, A_2, u) with utility profile (v_1, v_2) .*

Furthermore, if C is empty, then γ is a NE, if C is common information, then γ is a CHNE, and if a player has an open signal, then γ is a OSCE.

Proof. In general, it is easy to check that when σ is a computational NE for (Γ, C) , then the sequence $\{\text{Play}_{\Gamma, \sigma(\kappa)(C(\kappa))}\}_{\kappa=1}^{\infty}$ of distributions on $A_1 \times A_2$ is an (ε, δ) -CE for (A_1, A_2, u) for negligible ε and δ . If it was not, there would be a deviation for one of the parties, which would give non-negligible utility. This deviation could also be used in the computational cheap talk game, to get the same non-negligible utility, a contradiction. We can therefore use Lemma B.8 to get the result for CE. To get the result for NE, all we have to check is that each element $\text{Play}_{\Gamma, \sigma(\kappa)(C(\kappa))}$ is a product distribution, which is easy when σ is silent and $C(\kappa)$ is empty. To get the result for CHNE, all we have to show is that each element $\text{Play}_{\Gamma, \sigma(\kappa)(C(\kappa))}$ is a convex combination of product distributions, which is trivial when σ is silent and $C(\kappa)$ is common information. \square

B.1 For any (ETF) computational NE, a (NE-punishable) CE

We have already proven in Theorem B.11 that if the cheap talk game corresponding to (A_1, A_2, u) has a computational NE with utility profile (v_1, v_2) , then the strategic game (A_1, A_2, u) has a CE with utility profile (v_1, v_2) . We now give an analogue of Theorem B.11 for ETF computational NE of cheap talk games and NE-punishable CE of the underlying strategic game.

Theorem 5.1. *Let $\Gamma = (A_1, A_2, u)$ be a strategic game and let $\tilde{\Gamma}$ be the corresponding computational 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) .*

Proof. Assume that the cheap talk game $\tilde{\Gamma}$ corresponding to Γ has a computational ETFE σ with utility profile (v_1, v_2) . Then σ is in particular a computational NE, so we can get a CE γ for Γ constructed as in the proof of Theorem B.11. It will clearly have utility profile (v_1, v_2) , so it suffices to prove that this γ is a NE-punishable CE. We assume that it is not, and use this to conclude that then σ is not a computational ETFE, proving the theorem by contradiction.

If γ is not NE-punishable, then there exist $i \in \{1, 2\}$ and an action $a_i^\dagger \in A_i$ played with non-zero probability such that $U_i(\gamma|a_i^\dagger) < U_i(\sigma_i^\dagger)$, where σ_i^\dagger is the worst NE for player i . Assume without loss of generality that $i = 2$. By construction of γ there exists an infinite subset $K \subset \mathbb{N}$ such that the distribution on $A_1 \times A_2$ played by the sequence σ_κ for $\kappa \in K$ converges to γ .

By Def. A.17 it is sufficient for us to specify a strategy space Γ such that σ is not an ε -ETF for Γ for any negligible ε . Let Γ be any strategy space for σ such that there is at least one empty round of communication and such that player 2 is the player to send the message in the last round of the strategy space, and such that the size of the message is large enough that the entire view of an execution of σ can be sent in one message. By Def. A.11 it is sufficient to give an event E and a deviation σ_2^* of player 2 in the face of E such that player 2 gets noticeably more in all ETF plays in the sub-game defined by E occurring, when σ_2^* was played—we will be more precise below. For now, let E be the event that the execution of σ makes σ_2 output the *bad* $a_2^\dagger \in A_2$ for which $U_2(\gamma|a_2^\dagger) < U_2(\sigma_2^\dagger)$. Let σ_2^* be that player 2 after observing E waits until the last round of the strategy space and then sends its entire view of the execution of σ to player 1. We argue that these choices finish the proof of the theorem.

We have to show that for all negligible ε it holds for all ε -ETF strategies $\tilde{\sigma}$ for the last round of the strategy space, where actions are picked, that

$$\Pr[E|\tilde{C}]u_2(\tilde{\sigma}, \Gamma^{(\geq m)}, \tilde{C}|_E) > u_2(\sigma, \Gamma^{(\geq m-1)}, \bar{C} \wedge E) + \varepsilon,$$

where $\bar{C} = \sigma^{1, \dots, m-2}(C)$ and $\tilde{C} = \sigma^{(m-1)*}(\bar{C})$.

If this is not the case, then there exists a negligible ε and an ε -ETF strategy $\tilde{\sigma}$ for the last round of the strategy space such that

$$\Pr[E|\tilde{C}]u_2(\tilde{\sigma}, \Gamma^{(\geq m)}, \tilde{C}|_E) \leq u_2(\sigma, \Gamma^{(\geq m-1)}, \bar{C} \wedge E) + \varepsilon.$$

Let α be the probability that the bad a_2^\dagger is played in γ . It is non-zero, or a_2^\dagger could not be bad. This means that the probability that a_2^\dagger is played by σ_κ for $\kappa \in K$ converges to α , in particular, it will at some point remain over $\alpha/2$ for all sufficiently large $\kappa \in K$. This gives us that there exists negligible ε and negligible $\varepsilon' = \varepsilon/\alpha$ such that

$$u_2(\tilde{\sigma}, \Gamma^{(\geq m)}, \tilde{C}|_E) \leq u_2(\sigma, \Gamma^{(\geq m-1)}, \bar{C} \wedge E) / \Pr[E|\tilde{C}] + \varepsilon'.$$

We have that $u_2(\sigma, \Gamma^{(\geq m-1)}, \bar{C} \wedge E) / \Pr[E|\tilde{C}]$ is the expected utility of player 2 given that a_2^\dagger is played, which converges to $U_2(\gamma|a_2^\dagger)$ for $\kappa \in K$. From $U_2(\sigma_2^\dagger)$ and $U_2(\gamma|a_2^\dagger)$ being constants and $U_2(\gamma|a_2^\dagger) < U_2(\sigma_2^\dagger)$ it therefore follows that there exists a non-zero constant β , say $\beta = U_2(\sigma_2^\dagger) - U_2(\gamma|a_2^\dagger)$, such that

$$u_2(\tilde{\sigma}, \Gamma^{(\geq m)}, \tilde{C}|_E) < U_2(\sigma_2^\dagger) - \beta,$$

which means that player 2 is getting strictly less than in his worst NE. This clearly implies that the utility profile $(v_1, v_2) = u(\tilde{\sigma}, \Gamma^{(\geq m)}, \tilde{C}|_E)$ is a Cryptomania utility profile, i.e., it cannot be achieved by 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.

We then conclude the proof by showing that $(u, \Gamma^{(\geq m)}, \tilde{C}|_E)$ implements a CHNE profile. Notice that in the game $(u, \Gamma^{(\geq m)}, \tilde{C}|_E)$ there is no communication, so by Def. A.5 we get that $\tilde{\sigma}$ is an ε -NE (for a negligible ε) for this game, as we have assumed that $\tilde{\sigma}$ is an ε -ETFE for this game. From Theorem B.11 we then get that there exists a CE ψ for (A_1, A_2, u) with utility profile (v_1, v_2) . We now conclude that ψ is a CHNE, from the fact that $\tilde{\sigma}$ implements ψ given the common prior $\tilde{C}|_E$ and no communication, where the crucial property we use of $\tilde{C}|_E$ is that it is of the form that player 1 knows the signal of player 2.

We show that there exists a function $\xi : \mathbb{N} \rightarrow \mathbb{R}$ such that $\lim_{\kappa \rightarrow \infty} \xi = 0$ and that there exists a common information common prior D such that there exists a ξ -NE for (u, Γ^∞, D) with a utility profile ξ -close to ψ , where Γ^∞ is Γ extended to allow both players unbounded computing time. By the proof of Theorem B.11, this show that ψ is a CHNE profile.⁹

The proof goes as follows. At this point, rename $(u, \Gamma^{(\geq m)}, \tilde{C}|_E)$ to (u, Γ, C) and rename $\tilde{\sigma}$ to σ , all for notational convenience. We have assumed that σ is an ε -NE for (u, Γ, C) , where Γ allows no communication rounds and in the common prior player 1 knows the signal of player 2.

Change C into a related common prior D , which works as follows: first it samples $(s_1, s_2) \leftarrow C$. Then it outputs (D, D) , where $D = D(s_2)$ is the probability distribution of the action $a_2 \in A_2$ of player 2 in σ when he receives s_2 (write this as $\sigma_2(s_2)$). This probability distribution is represented as a vector of $|A_2|$ probabilities p_a for $a \in A_2$. Each p_a is the true probability of playing a written as a binary number, but truncated to precision $\log_2(2|A_2|v\kappa)$, where v is the absolute distance between the smallest utility in u to the highest utility in u , and with the probability of some fixed action a_1 rounded up to make the probabilities sum to 1. Notice and remember that the statistical distance between $\sigma_2(s_2)$ and $D(s_2)$ is at most $(v\kappa)^{-1}$. Also, notice and remember that $D(s_2)$ can be written down using at most $|A_2| \log_2(2|A_2|v\kappa)$ bits, so there is at most $2^{|A_2|} 2|A_2|v\kappa = O(\kappa)$ possible values of $D(s_2)$.

We then define a strategy γ for (u, Γ^∞, D) . The strategy $\gamma_2(D)$ for player 2 is to play an action according to the distribution D . The strategy $\gamma_1(D)$ for player 1 is to sample $(s_1, s_2) \leftarrow C$ until $D(s_1, s_2) = D$ and then play according to $\sigma_1(s_1)$. Compare this to playing σ in (u, Γ, C) . Here (s_1, s_2) are sampled, and player 1 plays $\sigma_1(s_1)$ and player 2 plays $\sigma_2(s_2)$. Notice that player 1 plays exactly the same distribution in the two cases. Note also that player 1 plays distributions which are at most at statistical distance $(v\kappa)^{-1}$ from each other. Let $\delta = 4\kappa^{-1}$. We claim that γ is an $(\varepsilon + \delta)$ -NE for (u, Γ^∞, D) . We prove by contradiction that no party has a deviation giving better utility than $\xi = \varepsilon + \delta$.

Assume first that player 1 can make a deviation γ_1^* giving utility better than $(\varepsilon + \delta)$ for (u, Γ^∞, D) . Then there is also a poly-time deviation γ_1^\dagger of player 1 giving utility better than $(\varepsilon + \delta)$ -NE, which follows

⁹In the proof ε and δ are assumed to be negligible, but all that is needed is that they go to 0, and Γ is assumed to be bounded to polynomial time computation, but this is not used in the proof, which is purely analytic.

from the fact that when the distribution D of how player 2 plays is fixed and known to player 1, then player 1 can compute an optimal strategy in poly-time, by solving some simple linear equations. Observe then that we have that $|u(\gamma_1^\dagger, D) - u(\gamma_1^\dagger, \sigma_2)| \leq \kappa^{-1}$, as $D(s_2)$ and $\sigma_2(s_2)$ have statistical distance at most $(\nu\kappa)^{-1}$. We have that $|u(\gamma_1, D) - u(\gamma_1, \sigma_2)| \leq \kappa^{-1}$ for the same reason. So, the deviation γ_1^\dagger gives extra utility at least $(\varepsilon + \delta) - 2\kappa^{-1} \geq 2\kappa^{-1}$ against σ_2 , a contradiction as $2\kappa^{-1}$ is non-negligible.

Assume then that player 1 can make a deviation γ_2^* giving utility better than $(\varepsilon + \delta)$ -NE for (u, Γ^∞, D) . Then there is also a poly-time deviation γ_2^\dagger of player 2 giving utility better than $(\varepsilon - \delta - \kappa^{-1})$ -NE, namely for each of the polynomially many values of D hard-code into γ_2^\dagger a distribution with statistical distance at most $(2\kappa\nu)^{-1}$ from the optimal reply to the strategy of player 1 given that the common prior is D , and use a reasoning as above. Using the same reasoning as above this then gives that the deviation γ_2^\dagger gives extra utility at least $(\varepsilon + \delta - \kappa^{-1}) - 2\kappa^{-1} \geq \kappa^{-1}$ against σ_1 , a contradiction as κ^{-1} is non-negligible.

So, γ is an ξ -NE for (u, Γ^∞, D) and $\lim_{\kappa \rightarrow \infty} \xi = 0$, for $\xi = \varepsilon + \delta$. It remains to show that the utility profile is ξ -close to that of ψ , but this follows using a similar reasoning, as we have moved the strategy at most $(\kappa\nu)^{-1}$ and hence changed the utility at most κ^{-1} . \square

C The Three-Card Trick

A deck of three cards and the ability to do a perfect shuffle is complete for cryptography. Specifically, it implies random Rabin OT. Random Rabin OT is a two-party protocol where player 1 has an output $m_1 \in \{0, 1\}$ and player 2 has an output $m_2 \in \{0, 1, \perp\}$. The bit m_1 is uniformly random. The probability that $m_2 = \perp$ is $\frac{1}{2}$, and, if $m_2 \neq \perp$, then $m_2 = m_1$. As for privacy, player 1 gets on information on whether $m_2 = m_1$ or $m_2 = \perp$. And, if $m_2 = \perp$, then player 2 has no information on m_1 .

We show how to implement semi-honest random Rabin OT given three cards. Say the deck consists of A_\spadesuit , K_\spadesuit and Q_\spadesuit . The trick proceeds as follows:

1. Player 1 shuffles the deck.
2. Player 2 shuffles the deck.
3. Player 1 takes the top card, c_1 , of the deck, hiding the value from player 2.
4. Player 2 takes the top card, c_2 , of the remaining deck, hiding the value from player 2.
5. A player i having $c_i \in \{A_\spadesuit, K_\spadesuit\}$ sets $b_i = 1$. A player i having $c_i = Q_\spadesuit$ sets $b_i = 0$.
6. Player 1 sends $c = m_1 \oplus b_1$ to player 2.
7. If $b_2 = 0$, then player 1 outputs $m_2 = c \oplus 1$. If $b_2 = 1$, then player 1 outputs $m_2 = \perp$.

It is easy to see that the joint output is distributed as follows

prob.	P ₁	P ₂
$\frac{1}{3}$	0	1
$\frac{1}{3}$	1	0
$\frac{1}{3}$	1	1

The parties run the above trick twice. If player 1 ends up with identical outputs in the two runs, then he calls a rerun. This goes on until player 1 has different outputs in the two runs. Whether player 1 gets outputs 01 or 10 is equiprobable. When player 1 has outputs 01, then player 2 has outputs 10 or 11. When player 1 has outputs 10, then player 2 has outputs 01 or 11. The distribution is as follows:

prob.	P ₁	P ₂
$\frac{1}{4}$	01	11
$\frac{1}{4}$	01	10
$\frac{1}{4}$	10	01
$\frac{1}{4}$	10	11

If player 1 takes his output to be the output in the first run and player 2 takes her output to be her output in the second run when the outputs are different and \perp when they are identical, then the output distribution is

prob.	P ₁	P ₂
$\frac{1}{4}$	0	\perp
$\frac{1}{4}$	0	0
$\frac{1}{4}$	1	1
$\frac{1}{4}$	1	\perp

Furthermore, the parties clearly has no information extra to these outputs, so they implemented a random Rabin OT.

C.1 Handling skew Three-Card Trick distributions

A skew Three-Card Trick distribution is a distribution on joint outputs as follows:

prob.	P ₁	P ₂
α	0	1
β	1	0
γ	1	1

where $\alpha, \beta, \gamma > 0$ (and $\alpha + \beta + \gamma = 1$). If the players do not have any information extra to their outputs, all skew Three-Card Trick distributions imply random Rabin OT.

For a starter, assume that joint output is distributed as follows:

prob.	P ₁	P ₂
β	0	1
β	1	0
γ	1	1

for $\beta, \gamma > 0$ and $2\beta + \gamma = 1$.

If $\beta = \gamma$, then this is the exact *Three-Card Trick* distribution, which we know implies random Rabin OT. We look at the two other case.

If $\beta > \gamma$, then consider the following protocol: The parties generate the above distribution. A player with output $b_i = 0$ will announce an abort with probability $1 - \gamma/\beta$. If a player announces an abort, then then both players output \perp . This gives the output distribution:

prob.	P ₁	P ₂
γ	0	1
γ	1	0
γ	1	1
$1 - 3\gamma$	\perp	\perp

with $\gamma > 0$. So, by rerunning in case of abort, they can generate the output distribution:

prob.	P ₁	P ₂
$\frac{1}{3}$	0	1
$\frac{1}{3}$	1	0
$\frac{1}{3}$	1	1

If $\beta < \gamma$, then consider the following protocol: The parties generate the above distribution. A player with output $b_i = 1$ will announce an abort with probability $< \delta = 1 - \frac{\beta}{1-2\beta}$. If a player announces an abort, then then both players output \perp . This gives the output distribution:

prob.	P ₁	P ₂
$(1 - \delta)\beta$	0	1
$(1 - \delta)\beta$	1	0
$(1 - \delta)^2\gamma$	1	1
1 - sum of the above	\perp	\perp

where it can be checked that $(1 - \delta)\beta = (1 - \delta)^2\gamma$ and that $(1 - \delta)\beta > 0$, so by rerunning in case of abort, they can generate the *Three-Card Trick* distribution.

Assume then that the joint output is distributed as follows:

prob.	P ₁	P ₂
α	0	1
β	1	0
γ	1	1

If $\alpha = \beta$, then we already showed how to handle this distribution, so assume without loss of generality that $\alpha > \beta$. Now, let player 2 call an abort with probability $1 - \frac{\beta}{\alpha}$ when $b_2 = 1$. This gives the joint output distribution:

prob.	P ₁	P ₂
β	0	1
β	1	0
γ'	1	1
$\alpha - \beta$	\perp	\perp

for $\gamma' = \frac{\gamma\beta}{\alpha}$. By rerunning this protocol until it does not abort, the output distribution becomes

prob.	P ₁	P ₂
β'	0	1
β'	1	0
γ''	1	1

for $\gamma'' = \frac{\gamma'}{1-\alpha+\beta}$ $\beta' = \frac{\beta}{1-\alpha+\beta}$. Since $\beta', \gamma'' > 0$ we already know how to handle such distributions.

C.2 Computational Case

We then consider the case where extra to the actions, each player also gets to see a random variable dependent on the outputs of both parties, and where we restrict the parties to be poly-time. Let view_i be the random variable seen by player i .

If we further require that

$$[\text{view}_1 | b_1 = 1 \wedge b_2 = 1] \approx [\text{view}_1 | b_1 = 1 \wedge b_2 = 0]$$

and

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

then we can implement a computational random Rabin OT using the exact same reduction as above. The proof is via a standard hybrids argument. The same holds true for all the skew Three-Card Trick distributions, as long as all probabilities are positive constants. It would even hold if the probabilities go to 0 as an inverse polynomial. We cannot allow the probabilities to go too faster to 0, as the run time of the reductions would not be polynomial.

D Details for Sect. 3

We give a formal proof for Claim 3.6.

Claim 3.6 (restated). *For any $k \in \mathbb{N}$, such that k If Λ_k and γ_k are defined as in Lemma 3.5, then γ_k is a correlated equilibrium of Λ_k .*

Proof. Clearly, γ_k satisfies the trivial constraints for being a CE, i.e., $\sum_{(a,b) \in A \times B} \gamma_k(a,b) = 1$, and $0 \leq \gamma_k(a,b) \leq 1$ for every action profile $(a,b) \in A \times B$. We need to check that no player has an incentive to deviate from any private advice distributed according to γ_k .

We show that given the advice a_1 , player A has no incentive to deviate. The expectation given advice a_1 is $U_A(\gamma_k|a_1) = \gamma_k(a_1, b_1)u_A(a_1, b_1) + \gamma_k(a_1, b_2)u_A(a_1, b_2) = \gamma_k(a_1, b_1)(f + g)$, that is strictly larger than $\gamma_k(a_1, b_1)(c + f)$ obtained by switching to a_2 , than $\gamma_k(a_1, b_1)(c + g)$ obtained by switching to a_k , or than $2c \cdot \gamma_k(a_1, b_1)$ gained by selecting any other action. The verification is analogous for any advice that player A (player B) gets such that the two action profiles selected with non-zero probability in the corresponding row(column) are equiprobable.¹⁰

Now consider the advice a_{k-1} of player A . It should hold that

$$\gamma_k(a_{k-1}, b_{k-1})f + \gamma_k(a_{k-1}, b_k)e \geq \gamma_k(a_{k-1}, b_{k-1})g + \gamma_k(a_{k-1}, b_k)c.$$

This can be transformed into

$$\gamma_k(a_{k-1}, b_k)(e - c) \geq \gamma_k(a_{k-1}, b_{k-1})(g - f).$$

It follows from the way $\gamma_k(a_{k-1}, b_{k-1})$ and $\gamma_k(a_{k-1}, b_k)$ are defined that the sides of the above inequality are identical. The recommendation b_1 of player B is the last one to consider, and the corresponding constraint is

$$\gamma_k(a_1, b_1)g + \gamma_k(a_{k-1}, b_1)d \geq \gamma_k(a_1, b_1)f + \gamma_k(a_{k-1}, b_1)c.$$

This inequality is also satisfied due to the initial condition that $c < d < e < f < g$. Hence, our selection of the parameters $c, d, e, f, g \in \mathbb{R}$ ensures that γ_k indeed is a correlated equilibrium. \square

Lemma 3.7. *There is no 2×2 strategic game that satisfies the extensibility criterion.*

Proof. It is easy to see that if Γ has a pure NE, then there is no CE γ outside CHNE strictly better than the best NE for every player. Any pure NE must lay on the weakly Pareto optimal boundary of the polygon defined by the four payoff profiles in Γ , moreover no convex combination of this payoff profiles can be strictly improving to the weakly Pareto optimal boundary. Thus, we need to only consider games without pure NE.

Moulin and Vial [15] provide a classification of 2×2 games w.r.t. the number of pure Nash equilibria. They show that a 2×2 game with no pure Nash equilibrium must have a unique totally mixed NE which cannot be improved by any CE. \square

D.1 A minimal example

We include also a minimal example for a game with a payoff profile in the polytope of correlated equilibria payoffs that cannot be achieved by any NE-punishable CE. Moreover, this payoff profile is strictly better for both players than their respective worst Nash equilibria.

¹⁰The verification in the case of advice b_k of player B goes through because of the requirement $g - f < e - c$.

	A	B	C	D
a	9, 9	-25, -25	-25, -25	-25, -25
b	-25, -25	9, 10	10, 9	-100, -100
c	-25, -25	-100, -100	9, 10	4, 9
d	-25, -25	10, 3	-100, -100	3, 4

Figure 6: A game with utility profiles not achievable by any NE-punishable CE.

Consider game Γ given by the payoff matrix in Fig. 6. The strategy profile $\sigma = (a, A)$ is a unique NE of Γ which is also the worst NE for both players with the utility profile $(9, 9)$. Let γ be a probability distribution over $A_1 \times A_2$ such that $\gamma(b, B) = 0.99$, $\gamma(b, C) = 9 \cdot 10^{-3}$, $\gamma(c, D) = 9 \cdot 10^{-4}$, $\gamma(c, C) = 9 \cdot 10^{-5}$, and $\gamma(d, D) = \gamma(d, B) = 5 \cdot 10^{-6}$. It can be verified that γ satisfies all the conditions for being a correlated equilibrium. The payoff profile achieved by γ is $(9.004475, 9.984635)$ that is strictly better for both players than the utility profile achieved by σ . However, given the advice B, the expected utility $U_2(\gamma|D)$ of player 2 is smaller than 9, hence γ is not a NE-punishable CE.

We argue that the utility profile $U(\gamma)$ cannot be achieved by any NE-punishable CE. Clearly, a CE γ' achieving payoff higher than 9 for player 2 must have in its support either (b, B) or (c, C) . If (b, B) is in the support of γ' , then (b, C) must be in its support as well for strategic reasons. Otherwise player 1 would deviate to playing d given advice b. Using similar reasoning, (c, C) and consequently (c, D) must be in its support. This already gives us that such γ' is not NE-punishable CE because the expected utility of player 1 given the advice c must be strictly smaller than 9.

One should note that the payoff profile $U(\gamma)$ is not the only one strictly dominating $U(\sigma)$ that cannot be achieved by any NE-punishable CE. One can for example easily come up with a CE for which the gap between its payoff and the payoff of σ is for the players more symmetric. The reason why we selected $U(\gamma)$ is because of the ease of enumeration of the achieved utilities and the probabilities in γ .

E Implementing Correlated Equilibria using Cryptographic Protocols

E.1 Implementing any CHNE using one-way functions

For completeness, we restate the result of Gradwohl *et al.* [9], who realized that all CHNE are NE-punishable and gave a protocol to implement any weakly Pareto optimal CHNE.

Theorem E.1 (Gradwohl *et al.* [9]). *Let Γ be a strategic game. If one-way functions exist, then for every γ , a weakly Pareto optimal CHNE of Γ , there exists an empty-threat free computational NE of (Γ, C_0) achieving the same utility profile.*

Take the CHNE γ . It can be written as a probability distribution over finitely many NE. For security parameter κ , take a probability distribution which is $2^{-\kappa}$ -close to γ and which can be sampled from a random string of length $\text{poly}(\kappa)$. Use coin-flipping to flip a random string of length $\text{poly}(\kappa)$: commit, send random string, open, take xor. If any party deviates, by sending more or less information than specified by the protocol, then punish with the worst NE. For a proof that it is indeed empty-threat free to punish with the worst NE, observe that any CHNE clearly is NE-punishable and then use the proof of the below theorem.

E.2 Implementing any NE-punishable CE using OT

Assuming the existence of OT allows us to use the full power of active-secure two-party computation. The players can thus securely implement the mediation device for sampling the NE-punishable CE.

Theorem E.2. *Let Γ be any strategic game and let Γ' be the computational cheap talk extension of Γ . If OT exists, then for any weakly Pareto optimal NE-punishable correlated equilibrium γ of Γ there exists*

an ETF computational Nash equilibrium σ of Γ' , such that the payoffs for both players are the same in σ and γ .

Proof (sketch). The players can use a protocol securely implementing the mediator that hands out advice according to γ . If any of the players deviates from the prescribed behavior, then she will be punished by the other player playing according to the worst NE for the deviating player. Specifically, the other player will use the strategy which ignores all messages sent from the other player and then at the end it will play according to the worst NE.

Since γ is NE-punishable, the expectation of every player from playing according to the assigned advice is strictly larger than her expectation in the worst NE. Therefore, the strategy profile in which each player follows the protocol and plays according to the obtained advice, or plays according to the worst NE for the other player in case the other player deviates is an ETF computational NE of Γ .

All that has to be checked is that it is empty-threat free to ignore the messages sent by the other player and then play according to the worst NE. Denote by r the round in which the punishing player adopts the strategy to ignore the messages sent by the punished player and then play according to the worst NE of the punished player at the end. We do the proof by reduction. If r is the last round of the strategy space, then there is no more communication rounds, so the only possible deviation of the punished player is to unilaterally change his action, which cannot give more utility as the parties are playing a NE. I.e., when played in the last round of the strategy space, any NE is also an empty-threat free NE. Assume then that r is not the last round of the strategy space and assume for the sake of contradiction that the punishing strategy is not empty-threat free when played from round r . In that case, by the definition of not being empty-threat free, there exists a future round $r' > r$ and a deviation of the punished player in round r' such that the punished player does better than in its worst NE in all empty-threat free plays starting with that deviation in round r' . It cannot be that r' is the last round of the strategy space, as then the deviation is again a unilateral deviation from a NE, which cannot give extra utility. But since r' is not the last round, it follows that one of the possible continuations from the deviation of the punished players is that the punishing player plays the original punishing strategies from round $r' + 1$ and on, and we can, by induction, assume that this is empty-threat free when played from round $r' + 1 > r$, as we have shown it to be empty-threat free when played in the last round. So, one of the empty-threat free continuations give the punished player the utility of his worst NE, so clearly it is not the case that all empty-threat free continuations give the punished player more utility than the play he deviated from, which was exactly a punishment to his worst NE. \square