Realizable Rational Multiparty Cryptographic Protocols*

John Ross Wallrabenstein and Chris Clifton**

Dept. of Computer Science, Purdue University, USA
{jwallrab,clifton}@cs.purdue.edu

Abstract. In this work, we describe how to realize rational cryptographic protocols in practice from abstract game specifications. Existing work requires strong assumptions about communication resources in order to preserve equilibria between game descriptions and realized protocols. We argue that for real world protocols, it must be assumed that players have access to point-to-point communication channels. Thus, allowing signaling and strategy correlation becomes unavoidable. We argue that ideal world game descriptions of realizable protocols should include such communication resources as well, in order to facilitate the design of protocols in the real world. Our results specify a modified ideal and real world model that account for the presence of point-to-point communication channels between players, where security is achieved through the simulation paradigm.

Keywords: Rational Multiparty Computation, Game Theory, Non-Cooperative Computation.

1 Introduction

The field of *rational cryptography* departs from modeling players as either honest or malicious, and instead models all players as rational utility-maximizing agents: each player chooses those actions that maximize their utility function $\mu(\cdot)$, which expresses their preferences over outcomes. All players may arbitrarily depart from the protocol specification if doing so is a utility-maximizing strategy. This approach to modeling removes the strong assumption of the semi-honest model: that honest players follow the protocol specification, regardless of whether or not it is in their best interest. By considering all players as rational agents, the standard properties of cryptographic protocols (e.g. privacy, correctness and fairness) are modeled through the utility functions of the players. Security of the protocol is then deduced from whether or not the stable equilibrium of the original game specification is reachable given the players' utility functions.

In secure multiparty computation (SMPC), the security of protocols are demonstrated through the *simulation paradigm*. Define an *ideal* protocol for

* The rights of this work are transferred to the extent transferable according to title 17 §105 U.S.C.

** This material is based upon work performed while Dr. Clifton was serving at the National Science Foundation.

R. Poovendran and W. Saad (Eds.): GameSec 2014, LNCS 8840, pp. 134–154, 2014.

computing a functionality f that invokes an incorruptible and universally trusted third party (TTP). Similarly, define a *real* protocol π for computing f where no TTP exists. Security is established if an adversary $\mathcal{A}$ in the real model has no advantage over a simulator $\mathcal{S}$ in the ideal model [1].

A major obstacle when defining security for rational multiparty protocols is the potential for players to form *coalitions*, colluding to undermine the security of the protocol. The strongest result, by Izmalkov et al. [2], allows any function to be computed securely by rational players using the approach of Goldreich et al. [1]. Although a universal result, it relies on strong assumptions including forced actions and physical primitives. A weaker notion, referred to as *collusion-free* computation [3–5], removes the ability of players to communicate additional information subliminally through the protocol communication resources. The result relies on a trusted mediator at the center of a star network topology, where all messages pass through the mediator and are re-randomized in order to prevent steganographic communication between the players. This result relies on adversarial independence, where simulators and adversaries are disallowed communication in the protocol. However, a collusion-free protocol may still cause issues when executed as part of a larger protocol. For example, the collusion-free protocols of Izmalkov et al. [2, 5] provide no guarantees when all players are malicious. This observation led to the work of Alwen et al. [6], where communication restrictions are further weakened to achieve *collusion-preserving* computation, which preserves any potential for collusion present in the original game specification. Although this result removes the requirement of a trusted mediator, it rules out a large class of communication resources (e.g. point-to-point and broadcast channels). Kamara et al. [7] consider a setting where adversaries have the capability to communicate additional information during protocol execution, yet choose to be *non-colluding*. Fuchsbauer et al. [8] give constructions under standard communication channels by forcing parties to send only unique messages as part of the protocol. Thus, collusion within the protocol is avoided, but communication outside of the protocol execution still facilitates collusion.

From this collection of work, addressing the issue of collusion appears to require strong limitations on the type of communication resources granted to players. As the general goal of rational cryptography is to provide a more realistic view of how players behave in cryptographic protocols, we consider what can be achieved when players have access to point-to-point communication channels - an unavoidable aspect in real world applications. Thus, in this work we define a security model where players may communicate information over point-to-point channels both inside and outside the protocol execution.

Our work proposes a new security framework for rational agents that models player access to point-to-point communication channels in the ideal world model. From this, we describe how to demonstrate the security of protocols in a real world model that implements games specified in our modified ideal world model. We note that imposing restrictions on the ideal world to capture unavoidable behavior exists currently in the cryptographic literature: it is a core feature of the malicious model, which extends the semi-honest model to consider more powerful

adversaries. In the malicious setting, the ideal world must capture the ability of an adversary to coordinate the actions and inputs of players it corrupts, and force aborts during protocol execution; these actions are unavoidable in the presence of a monolithic malicious adversary. Our model necessarily limits the class of games that may be modeled in the ideal world formulation of our framework, as point-to-point communication channels must exist in the original game. Our work differs from existing formulations, which attempt to realize all games at the expense of restricting the communication interface available to players.

Throughout the remainder of the introduction, we argue that when point-to-point communication channels are unavoidable, it is meaningful to consider what games are realizable in their presence. We demonstrate that a non-trivial class of games constructed in our modified ideal world model have realizable implementations in the real world model through the Signaling game in Section 1.2, and the classic prisoner's dilemma in Section 2. We give our technical contribution, a security model for realizing protocols from game specifications in the presence of point-to-point communication channels, in Section 3. We demonstrate the power of our model relative to others through a full proof of security for the rational secret sharing protocol of Halpern and Teague [9] in Section 4, which is inadmissible under existing frameworks due to the presence of point-to-point communication channels. These examples demonstrate the key contribution of our model, which is less restrictive than prior work yet is able to correctly model the games' equilibria when played in the real world.

1.1 Local Adversaries

Translating the standard simulation paradigm to the game theoretic setting of rational cryptography requires addressing how adversaries should be modeled. In the original formulation, a centralized semi-honest or malicious adversary corrupts a subset of the players. However, rational cryptography makes no such distinction[1] between honest and corrupted players, and assumes all players are rational and acting to maximize their local utility function. Thus, translating the concept of an adversary is not immediate. Alwen et al. [6] give a collusion preserving framework where each player has an associated *local* adversary. Thus, the monolithic adversary of the standard model is shattered into an adversary for each individual player. Canetti et al. [11] argue that a local adversary should be defined for each ordered pair of players, as this provides a more granular model of the flow of information. Canetti et al. then demonstrate that the *local universal composition* (LUC) model can preserve the incentive structure in games.

We follow this modeling trend of shattering the monolithic adversary $\mathcal{A}$ into a *set* of local adversaries $\mathcal{A} = \{\mathcal{A}_i\}_{i\in[1\ldots n]}$ such that each player $\mathtt{P}_i \in \mathcal{P}$ is associated with adversary $\mathcal{A}_i$. Rather than considering local adversaries that "corrupt" their associated player $\mathtt{P}_i$, we simply require that the adversary selects

[1] A mixed model has been proposed by Lysyanskaya et al. [10] where one subset of players are arbitrarily malicious, and the other subset are utility-maximizing rational agents.

the actions of $\mathtt{P}_i$ to maximize their local utility function μ_i. Thus, we preserve the assumption in rational cryptographic protocols that all players are purely rational and bound to a utility function, rather than remaining honest unless corrupted by a monolithic adversary.

1.2 Communication Resources

A core issue with existing work is how communication resources are modeled in game descriptions. In order to prevent players from signaling information or coordinating their actions, available communication resources are tightly restricted. For example, Izmalkov et al. [2] propose *rational secure computation* where only those equilibria in the game description exist in the realized protocol. However, this result comes at the cost of requiring forced actions and physical primitives such as opaque envelopes and ballot boxes[2]. Although not impossible to realize, in practice it has limited applicability.

In the ideal world model of secure multiparty computation, a protocol is viewed as an interaction between a set of players and a universally trusted third party (TTP). An ideal computation of a function has each player send their private input to the TTP, who computes the function and returns the results to each player. Restricting communication resources is not necessary, as players are assumed to be mutually distrustful. Further, any collusion between players is modeled through a monolithic adversary $\mathcal{A}$ that coordinates the actions of the players it corrupts.

In order to implement *arbitrary* games as protocols, strict notions of privacy preservation and the prevention of signaling and correlation must be satisfied. Arbitrary game specifications may impose restrictions on the communication resources available to players. Thus, the corresponding protocol implementation must not allow players to communicate more information than is possible in the ideal game specification. We briefly review the characteristics a model for implementing arbitrary games must satisfy[3]. We make the argument that even if a protocol satisfies all of these characteristics, it is likely to fall short of satisfying the ideal world model: communication between players outside of the protocol is unavoidable in real world settings. Thus, the model we present is not bound to satisfy these restrictions, and is a more accurate representation of what is achievable for protocols executed in the real world.

Privacy. A protocol π implementing an arbitrary game Γ must preserve both *pre-game privacy* and *post-game privacy* in addition to preserving the equilibrium of Γ. The notion of pre-game privacy ensures that the private input of each party is not revealed, as this will affect the actions of other parties. However, protocols implementing arbitrary games must also preserve the notion of post-game privacy, where nothing beyond the intended result (and what can be

[2] This result is a direct application of the GMW protocol [1].

[3] The ECRYPT summary report [12] on rational cryptographic protocols provides background on modeling techniques used to address privacy, signaling and correlated actions.

inferred from this) is revealed. This notion is necessary so that the equilibria of future games are not perturbed by information revealed in previous games.

Signaling. Similar to the notions of pre- and post-game privacy are the notions of *pre-game signaling* and *post-game signaling*. The ability to signal other players allows protocol participants to coordinate their actions to achieve a higher payoff. For example, consider two players A and B with inputs a and b. The payoff function is defined as $\Pi(\Gamma) \coloneqq a \oplus b$, and described in Table 1:

Table 1. Signaling Game

	A **sets** $a = 1$	A **sets** $a = 0$
B **sets** $b = 1$	(0,0)	(1,1)
B **sets** $b = 0$	(1,1)	(0,0)

If A or B can signal even a single bit to the other, each will receive a payoff of 1 as opposed to an expected payoff of $\frac{1}{2}$. Thus, similar to the restriction on privacy, preventing pre- and post-game signaling is necessary to preserve the equilibria of individual and future games when constructing protocols for *arbitrary* games.

The signaling game specification can be formulated under existing frameworks as a protocol, and demonstrated to preserve the mixed equilibrium of the original game. Yet by ignoring the ability of players to communicate outside of the protocol, the protocol formulation is invalidated in real world settings: players will collude to achieve a payoff of 1, rather than the expected payoff of $\frac{1}{2}$ of the original game specification.

We only consider those game specifications that allow point-to-point communication, as these channels are unavoidable in the real world. Thus, our model correctly predicts a payoff of 1 for players in the signaling game, as point-to-point communication channels allow signaling.

Correlated Actions. Correlated actions are similar to signaling, but allow parties to coordinate actions without exchanging information. This is usually accomplished through a shared value, such as a *common reference string* (CRS). The parties need not distribute information, but rather rely on the shared CRS to coordinate their actions. As with signaling, protocol constructions for arbitrary games must prevent pre- and post-game correlation to preserve equilibria in local as well as future games.

2 Prisoner's Dilemma

As a classic example, we consider the Prisoner's Dilemma[4]: a game between two suspects A and B that have been accused of committing both a principal and

[4] The concept was originally proposed by Flood and Dresher while working at the RAND corporation, and is described in detail by Poundstone [13].

lesser crime. The Authority has sufficient evidence to convict both A and B on the lesser crime, punishable by 1 year in prison. However, there is insufficient evidence to convict A or B on the principal crime. The Authority separates A and B, and offers the following proposal: confess and serve no time while your partner serves 3 years in prison. Players A and B are then subject to the following dilemma:

1. If both A and B remain silent, they will each be convicted on the lesser crime and serve 1 year in prison.
2. If one confesses while the other remains silent, the confessor is set free while the other serves 3 years in prison.
3. If both A and B confess, each will serve 2 years in prison.

Table 2. Prisoner's Dilemma Game

	A **Remains Silent**	A **Confesses**
B **Remains Silent**	(-1,-1)	(0,-3)
B **Confesses**	(-3,0)	(-2,-2)

From the player payoffs listed in Table 2, note that each player maximizes their utility by confessing to the principal crime regardless of the strategy of their partner. We use this example to illustrate the necessity of removing monolithic adversaries, as well as how communication assumptions should be formulated in the ideal game description. Note that the original ideal game specification of the prisoner's dilemma requires that the suspects A and B are physically separated: thus unable to communicate or otherwise coordinate their actions. However, we will construct a modified formulation in the presence of point-to-point communication channels with an *equivalent equilibrium* to the original formulation under our proposed model.

2.1 Monolithic Adversaries

Traditionally, cryptographic protocols are analyzed with respect to their resilience to a monolithic adversary $\mathcal{A}$ corrupting some subset of the players. Protocol resilience to adversarial corruption is quantified by the fraction of players that may be corrupted before the protocol security is violated.

In the game theoretic setting of rational cryptography, this model has been called into question by Alwen et al. [6] and Canetti et al. [11]. The goal of rational cryptography is to model each player as bound to their local utility function, rather than controlled by a monolithic adversary with a global utility function. The monolithic adversary in both of their models is shattered into a set of *local* adversaries unique to each player. Removing the monolithic adversary in favor of

a set of local adversaries is critical to preserving game theoretic equilibria. In the running example of the Prisoner's Dilemma, consider the case where $\mathcal{A}$ corrupts both A and B. As $\mathcal{A}$ controls both players, A and B may be forced to remain silent and achieve payoff $(-1, -1)$. However, consider the case where A (resp. B) has a *local* adversary $\mathcal{A}_{\mathtt{A}}$ (resp. $\mathcal{A}_{\mathtt{B}}$): as $\mathcal{A}_{\mathtt{A}}$ is bound to the utility function $\mu_{\mathtt{A}}(\cdot)$ of A, $\mathcal{A}_{\mathtt{A}}$ maximizes $\mu_{\mathtt{A}}(\cdot)$ by confessing as in the ideal specification of the game. An identical argument holds for $\mathcal{A}_{\mathtt{B}}$ as well. Thus, a monolithic adversary is capable of introducing a stable collusion equilibrium that does not exist in the ideal game specification, whereas the local adversary model preserves the original incentive structure.

2.2 Realistic Communication Model

To prevent pre- and post-game signaling and strategy correlation, many rational cryptographic frameworks impose strong restrictions on the communication resources available to players. This issue is most pronounced in the multiparty setting, where communication resources may enable collusion. To prevent communication resources from perturbing the equilibria of the ideal world game, existing constructions require forced player action and physical primitives [2], trusted mediators and forced broadcast channels [4], as well as the cooperation of adversarial players to deliver messages [6].

While these results provide strong guarantees under restrictive communication resource assumptions, the security guarantees are with respect to the protocol only. That is, assuming players may only interact through the protocol and its communication resources, the equilibria of the ideal world game is preserved. However, we argue that this results in a false sense of security for protocols realized in the real world, where players typically have access to point-to-point communication channels - undermining the strict communication assumptions of the protocol.

Our example of the prisoner's dilemma illustrates a salient point: the necessary and sufficient condition for preserving the equilibrium of the original formulation is the ability of A and B to *privately communicate* with the Authority. The original game specification requires the two players A and B to be physically separated, and thus unable to communicate. However, the key to preserving the equilibrium (confess, confess) of the original game Γ only requires preventing A and B from observing their interaction with the Authority. Consider a modified game $\bar{\Gamma}$ where all players $\{\mathtt{A}, \mathtt{B}, \mathtt{Authority}\} \in \mathcal{P}$ have access to a point-to-point communication resource $\mathcal{R}$. As long as the communication links $\mathcal{R}_{\mathtt{A},\mathtt{Authority}}$, $\mathcal{R}_{\mathtt{B},\mathtt{Authority}}$ are private, the original equilibrium is preserved despite the presence of point-to-point communication channels. In game theoretic terms, communication between A and B through $\mathcal{R}_{\mathtt{A},\mathtt{B}}$ is considered *cheap talk*, as both A and B will claim to play silent, yet as utility maximizing agents they choose to confess, which strictly dominates silent. As neither A nor B can observe the message sent by the other to Authority, the coalition is unstable and disintegrates despite the presence of point-to-point communication channels.

3 Our Contribution

We argue that ideal world protocols should assume that players have the ability to communicate over point-to-point channels. As in the standard SMPC ideal world model, players may not wish to communicate due to mutual distrust. However, the option to do so should be part of the model, as this is unavoidable in the real world. Thus, we present a modified ideal world model capturing the presence of point-to-point communication channels between all players. Specifically, we answer the following questions:

1. How is security formalized when all players are rational and have access to point-to-point communication channels?
2. What benefits result from weakening the security guarantees of the standard malicious model by considering rational players with local adversaries?

3.1 Unstable Coalitions

A powerful aspect of the rational cryptographic setting with local adversaries is the ability to design protocols where coalitions are unstable. As each player has a local adversary that selects their actions in order to maximize a utility function, protocols may be designed to incentivize players to *leave* coalitions [14]. This benefit of modeling each player as an independently rational agent is frequently overlooked, and allows game equilibria to be preserved despite the presence of point-to-point communication channels. We have illustrated the power of unstable coalitions through our example of the prisoner's dilemma. We now consider coalition stability in the setting of *rational secret sharing*, as it is the most familiar example of a rational cryptographic protocol.

Rational Secret Sharing Candidate definitions for achieving security against rational agents should accurately model well-studied problems in rational cryptography. The most familiar rational cryptographic protocol is *rational secret sharing* [8, 15–19]. The goal of *threshold* secret sharing is to split a secret among n parties such that any k shares are sufficient to recover the secret value, using a scheme such as the polynomial interpolation approach proposed by Shamir [20]. Rational secret sharing, introduced by Halpern and Teague [9], is particularly concerned with the process of recovering the secret from the shares[5]. As noted by Halpern et al. [9], rational players' utility functions are assumed to value *exclusivity*, where preference is given to learning the output of the function while preventing other players from doing so. Under this assumption, no party has any incentive to distribute their share to the other parties, which destabilizes coalition formation. The equilibrium is to wait for other players to distribute their shares, as this is the only action that increases a player's utility function.

[5] Maleka et al. [21] consider rational secret sharing in the context of repeated games, and Nojoumian et al. [22] consider the repeated game setting from a socio-rational perspective where player reputation is important.

Thus, a player that does not distribute their share has the potential to be the exclusive player to recover the secret.

The authors demonstrate that this implies no deterministic protocol exists where rational parties are willing to disseminate their shares to other players. Their randomized protocol is a modified game where players are distributed a *set* of shares, where only one share is correct. In each round k, players distribute their shares which evaluate to either the secret or a default value $\perp$. The solution relies on the fact that parties are unaware whether the current round k is terminal (k^*, allowing the secret to be recovered), or merely a "test" round $k \neq k^*$ (where the secret cannot be recovered, but players who do not distribute shares are caught as cheaters). By choosing k^* from a geometric distribution, as in Groce et al. [18], cheating players that choose strategy $\sigma = \perp$ when $k \neq k^*$ are caught and the game may be terminated. Thus, players now have an incentive to distribute their share, as playing $\perp$ only yields positive utility when $k = k^*$.

A candidate security definition should accept this probabilistic protocol for rational secret sharing as secure against rational agents. However, the strong restrictions on communication channels imposed by existing work preclude the above protocol from satisfying their security definitions, despite refinements considering the problem under standard communication models [8, 23–25]. That is, the rational secret sharing protocol of Halpern and Teague [9] assumes players have access to a non-rushing broadcast channel. This clearly violates the assumptions of models assuming physical primitives [2], and even fails to satisfy the weakest security definition that has been proposed: collusion-preserving computation [6]. Ideally, the original rational secret sharing protocol of Halpern and Teague should be demonstrably secure against rational agents under a general security framework. Our framework allows point-to-point communication in the ideal model, and thus is able to accurately model the original solution to rational secret sharing, which we demonstrate in Section 4.

3.2 Adversarial Model

Traditionally, an adversary $\mathcal{A}$ is viewed as a monolithic entity with a specified computational complexity and ability to "corrupt" players in a static or dynamic fashion. In our model, we consider all players to have the ability to act in an adversarial manner. Thus, rather than considering a monolithic adversary $\mathcal{A}$, we endow each player $\mathtt{P} \in \mathcal{P}$ with a local adversary $\mathcal{A}_{\mathtt{P}}$. The adversary is bound to the player's utility function $\mu_{\mathtt{P}}(\cdot)$ and selects actions for $\mathtt{P}$ in order to maximize $\mu_{\mathtt{P}}(\cdot)$. Note that as we bind player actions to a local adversary seeking to maximize a utility function, we cannot bound the number of players that deviate from the protocol. This is an unavoidable consequence of modeling players as rational agents; they select strategies to maximize a local utility function and follow the protocol only when doing so is advantageous. As cryptographic protocols typically require a number of rounds of interaction, we allow the rational players to update their strategy based on observations throughout the game Γ. Thus, we assume each local adversary is *mobile* [26], and may choose to deviate or follow the protocol at each round in a dynamic fashion. Additionally, players

may choose probabilistic strategies[6], so we must introduce a random tape $r_{\mathtt{P}}$ for each player P. Thus, each local adversary is adaptive, mobile, probabilistic, malicious, runs in *probabilistic polynomial-time* (PPT) and is presumed rational: bound to the player's local utility function.

Given the above definition of adversaries, the following actions are unavoidable:

- **Refusal to Participate:** Players may refuse to participate in the protocol. Constructions satisfying our definition thus assume that it is advantageous for players to engage in the protocol, and that this constitutes a utility maximization strategy with respect to their local utility function.
- **Input Substitution:** Players may supply an input to the protocol different from their true input.
- **Premature Abort:** Players may abort the protocol prior to completion.
- **Collusion:** Players may privately communicate over point-to-point communication channels, and collude to influence the protocol execution.

3.3 Ideal World Model

We now formalize the *ideal world* model, under which an ideal game specification Γ is constructed. We assume familiarity with standard game theoretic concepts in our exposition[7]. We first define the game specification of Γ under the *extensive form game* representation. In the game theoretic literature, *normal form game* representation is generally used for single round games where actions are played simultaneously. As cryptographic protocols typically proceed in a series of rounds where actions are played asynchronously, we prefer *extensive form game* representation, where the ideal game specification Γ is represented as a tree. At each node in the game tree, a subset $\mathtt{P} \subseteq \mathcal{P}$ of the players select and simultaneously play an action.

Definition 1. *An* ***extensive form game*** Γ *consists of:*

1. *A finite set* $\mathcal{P} = \{P_i\}_{i=1}^{n}$ *of players.*
2. *A (finite) set of sequences* $\mathcal{H}$ *called the* history. *The empty sequence* $\emptyset$ *is a member of* $\mathcal{H}$. *We let* k *denote the current decision node. If* $(a^k)_{k=1,\ldots,K} \in \mathcal{H}$ *and* $L < K$ *then* $(a^k)_{k=1,\ldots,L} \in \mathcal{H}$. *If an infinite sequence* $(a^k)_{k=1}^{\infty}$ *satisfies* $(a^k)_{k=1,\ldots,L} \in \mathcal{H}$ *for every positive integer* L *then* $(a^k)_{k=1}^{\infty} \in \mathcal{H}$. *A history* $(a^k)_{k=1,\ldots,K} \in \mathcal{H}$ *is a terminal history if it is infinite or if there is no* a^{K+1} *such that* $(a^k)_{k=1,\ldots,K+1} \in \mathcal{H}$. *The set of actions available after the nonterminal history* h *is denoted* $A(h) = \{a : (h, a) \in \mathcal{H}\}$ *and the set of terminal histories is denoted* $\mathcal{Z}$. *We let* $\mathcal{H}^k$ *denote the history through round* k.

[6] In a game theoretic setting, such strategies are referred to as *mixed*.

[7] For a proper introduction to the subject, Katz [27] describes the current effort to combine game theoretic and cryptographic concepts, while Osborne et al. [28] and Fudenberg et al. [29] give a complete introduction to game theory.

3. *A player function P that assigns to each nonterminal history (each member of $\mathcal{H}/\mathcal{Z}$) a member of $\mathcal{P} \cup \{nature\}$. When $P(h) = nature$, then nature determines the action taken after history h.*
4. *For each player $P_i \in \mathcal{P}$ a partition $\mathcal{I}_i$ of $\{h \in \mathcal{H} : P(h) = i\}$ with the property that $A(h) = A(h')$ whenever h and h' are in the same member of the partition. For $I_i \in \mathcal{I}_i$ we denote by $A(I_i)$ the set $A(h)$ and by $P(I_i)$ the player $P(h)$ for any $h \in I_i$. Thus, $\mathcal{I}_i$ is the information partition of player i, while the set $I_i \in \mathcal{I}_i$ is an information set of player i.*
5. *For each player $P_i \in \mathcal{P}$ a preference relation $\precsim_i$ on lotteries[8] over $\mathcal{Z}$ that can be represented as the expected value of a payoff function defined on $\mathcal{Z}$.*

Throughout, we replace the preference relation $\precsim_i$ by a *utility function* $\mu_i : A \to \mathbb{R}$, such that $\mu_i(a) \geq \mu_i(b)$ when $b \precsim_i a$.

We make the following modeling choices:

- **Extensive Form Games:** The ideal game specification Γ is described by a game tree in extensive form representation.
- **Imperfect Information:** A game specification is said to have *imperfect information* if players may have non-singleton information sets $I_i \in \mathcal{I}_i$. That is, at a given round in the game, players may be unaware of the move selected by the previous player(s). Thus, their information set may contain more than one node in the game tree at any given round.
- **Local Simulators:** Each player $\mathtt{P}_i \in \mathcal{P}$ in the ideal model has a *local* simulator $\mathcal{S}_i$ that forces $\mathtt{P}$ to play those actions that maximize $\mu_i(\cdot)$, the utility function of player $\mathtt{P}_i$. Each simulator $\mathcal{S}_i$ has an associated adversary $\mathcal{A}_i$ in the real world execution model, denoted $\mathcal{S}_i = \mathsf{Sim}(\mathcal{A}_i)$.
- **Point-to-Point Communication Resources:** Each player pair $(\mathtt{P}_i, \mathtt{P}_j)_{i \neq j} \in \mathcal{P}$ has a secure point-to-point communication resource $\mathcal{R}_{ij}$.

As we consider all players to be rational agents, we model the ideal world protocol as a game specification Γ that aims to achieve an equilibrium. The ideal game specification is an interaction between a set of n players $\mathcal{P} = \{\mathtt{P}_i\}_{i=1}^n$, their local utility functions $\boldsymbol{\mu} = \{\mu_i\}_{i=1}^n$ and action sets A_i, which contains those actions playable by player $\mathtt{P}_i$. Frequently, a deterministic choice of an action $a \in A_i$ will not yield a Nash equilibrium. Thus, we allow players to choose a *strategy* σ_i: a probability distribution over A_i. The standard equilibrium concept in the rational cryptographic literature is a *computational* Nash equilibrium [24, 25, 30–32], given by Definition 2.

Definition 2. *A **computational Nash equilibrium** of a two-party extensive-form game Γ is an independent strategy profile $\boldsymbol{\sigma^*} = \{\sigma_i^*\}_{i=1}^n$, such that*

1. *$\forall \sigma_i^* \in \boldsymbol{\sigma^*}, \sigma_i^*$ is PPT computable.*

[8] Even if all actions are deterministic, moves by *nature* can induce a probability distribution over the set of terminal histories.

2. *for each player* $\mathtt{P}_i$*, any other PPT computable strategy* $\sigma_i' \neq \sigma_i^*$*, we have* $\mu_i(\sigma_i', \boldsymbol{\sigma^*_{-i}}) \leq \mu_i(\boldsymbol{\sigma^*}) + negl(\lambda)$

where $\sigma_{-i} \stackrel{\text{def}}{=} (\sigma_j)_{j \in [1 \ldots n]/\{i\}}$.

Intuitively, no player $\mathtt{P}_i$ has an incentive to deviate from strategy σ_i given that every other player $\mathtt{P}_j$ selects their equilibrium strategy σ_j. The definition of a computational Nash equilibria adds a negligible term $negl(\lambda)$ with respect to a security parameter λ. This is necessary in the computational setting, as security rests on the premise that breaking cryptographic primitives occurs with only negligible probability. Thus, this notion must be incorporated into the equilibrium definition. Although computational Nash equilibria are the weakest of the equilibrium concepts described in the rational cryptographic literature, preserving only computational Nash equilibria in our framework is sufficient for extensions to more powerful equilibrium concepts.

The standard ideal world model has players interact with an incorruptible trusted third party (TTP) that accepts player inputs, computes the ideal functionality f, and distributes the output to players. In the setting of rational cryptography, we will consider a `Mediator` that enforces the ideal game specification.

Input Distribution:	Each player $\mathtt{P}_i \in \mathcal{P}$ receives its input x_i, random coins r_i and auxiliary input[a] z_i. Each player has the option of inputting a different input $\bar{x}_i \neq x_i$, as this is unavoidable.
Game Execution:	The `Mediator` allows the subset of players $\mathtt{P} \subseteq \mathcal{P}$ specified at each node of the game specification Γ to simultaneously play their actions. Note that games where only a single player moves at each node (asynchronous play) are fully supported, as this is modeled by setting the subset $\mathtt{P} = \{\mathtt{P}_i\}$.
Payoff Assignment:	If the current node k is terminal (i.e. $k \in \mathcal{Z}$), then `Mediator` distributes the payoffs associated with k to all players $\mathtt{P}_i \in \mathcal{P}$.

[a] An auxiliary input is provided to all players to model additional information available to them [33].

Protocol 3.1. Ideal World Game Execution

Definition 3. *Let* Γ *represent the ideal game specification in extensive form representation,* $\mathcal{R}$ *a point-to-point communication resource available between all pairs of players in* $\mathcal{P}$*,* $\mathcal{S}$ *the set of local simulators,* $\boldsymbol{\mu}$ *the set of player utility functions and* z *any auxiliary information provided to a player. We denote by* $\bar{\boldsymbol{x}}$ *the set of inputs for players (which may differ from the set of their true inputs* $\boldsymbol{x}$*) and by* r *the random coins provided to a player. We then define the* i^{th} *output of*

an `ideal world execution` *for players* $\mathcal{P}$ *in the presence of local simulators* $\mathcal{S}$ *as:*

$$\left\{\text{IDEAL}^{(i\in[1\ldots n])}_{\Gamma,\mathcal{R},\mathcal{P},\mathcal{S},\boldsymbol{\mu},z}(\lambda,\bar{\boldsymbol{x}};r)\right\}_{\lambda\in\mathbb{N},\bar{\boldsymbol{x}},r\in\{0,1\}^*} \triangleq \{\boldsymbol{\sigma}^*,\mathcal{I}\}$$

where $\boldsymbol{\sigma}^*$ *is the equilibrium in the ideal game specification* Γ*,* $\mathcal{S}=\{\mathcal{S}_i\}_{i\in[1\ldots n]}$ *is the set of simulators such that* $\mathcal{S}_i=\mathsf{Sim}(\mathcal{A}_i)$*,* $\mathcal{I}$ *is the information partition set for* $\mathcal{P}$*,* $|\bar{x}_i|=|\bar{x}_j|\forall i\neq j$ *and* $|z|=\text{poly}(|\bar{x}_i|)$*.*

This ideal world model necessarily limits the class of games that may be realized, as any game specification that disallows point-to-point communication channels between all parties cannot be modeled in the presence of $\mathcal{R}$. However, we will demonstrate that a broad class of games that initially appear inadmissible under our model are realizable through minor modifications to the game specification, and which preserve the equilibria of the original game.

3.4 Real World Model

We now introduce the real world model protocol Π that implements the ideal game specification Γ. In order to translate ideal game specifications into realizable protocols, we assume the existence of a public key infrastructure (PKI) in the real world model. That is, we must translate the ideal world point-to-point communication resource $\mathcal{R}$ into an implementation allowing point-to-point private communication between all players $\mathtt{P}_i,\mathtt{P}_j\in\mathcal{P}$ during the execution of Π. We denote the real world PKI communication resource by $\mathcal{C}$, where $\forall(\mathtt{P}_i,\mathtt{P}_j)_{i\neq j}\in\mathcal{P},\exists\mathcal{C}_{ij}\in\mathcal{C}$.

In the real world execution, each player $\mathtt{P}_i$ has an associated local adversary $\mathcal{A}_i$, rather than a simulator $\mathcal{S}_i$ as in the ideal world game. The local adversary $\mathcal{A}_i$ selects the actions of $\mathtt{P}_i$ to maximize the player's local utility function μ_i. Similarly, in the real world execution there is no `Mediator`, as the goal is to remove reliance on trusted third parties.

Input Distribution:	Each player $\mathtt{P}_i\in\mathcal{P}$ receives its input x_i, random coins r_i and auxiliary input z_i. Each player has the option of inputting a different input $\bar{x}_i\neq x_i$, as this is unavoidable.
Protocol Execution:	The execution of Π proceeds in a series of rounds, where at each round a subset of players $\mathtt{P}\subseteq\mathcal{P}$ specified at each node play their actions. Each player pair $(\mathtt{P}_i,\mathtt{P}_j)_{i\neq j}\in\mathcal{P}$ is connected by a private authenticated point-to-point communication channel $\mathcal{C}_{ij}$, and may exchange messages throughout the protocol execution.
Payoff Assignment:	If the current node k is terminal (i.e. $k\in\mathcal{Z}$), then each player $\mathtt{P}_i\in\mathcal{P}$ receives its associated payoff.

Protocol 3.2. Real World Protocol Execution

Definition 4. *Let Π represent the real world protocol implementing Π, $\mathcal{C}$ a point-to-point authenticated and private PKI communication resource available between all pairs of players in $\mathcal{P}$, $\mathcal{A}$ the set of local adversaries, $\boldsymbol{\mu}$ the set of player utility functions and z any auxiliary information provided to a player. We denote by $\bar{\boldsymbol{x}}$ the set of inputs for players (which may differ from the set of their true inputs $\boldsymbol{x}$) and by r the random coins provided to a player. We then define the i^{th} output of a* `real world execution` *for players $\mathcal{P}$ in the presence of local adversaries $\mathcal{A}$ as:*

$$\left\{\mathrm{REAL}^{(i\in[1\ldots n])}_{\Pi,\mathcal{C},\mathcal{P},\mathcal{A},\boldsymbol{\mu},z}(\lambda,\bar{\boldsymbol{x}};r)\right\}_{\lambda\in\mathbb{N},\bar{\boldsymbol{x}},r\in\{0,1\}^*} \triangleq \{\boldsymbol{\sigma}^*,\mathcal{I}\}$$

where $\boldsymbol{\sigma}^*$ is the equilibrium in the real world protocol Π, $\mathcal{I}$ is the information partition set for $\mathcal{P}$, $|\bar{x}_i| = |\bar{x}_j| \forall i \neq j$ and $|z| = \mathrm{poly}(|\bar{x}_i|)$.

3.5 Establishing the Security of Realized Protocols

The security of protocols is established by demonstrating that the real and ideal world distribution ensembles are computationally indistinguishable[9]. This guarantees that any attack available to an adversary $\mathcal{A}$ in the real model is also available to a simulator $\mathcal{S}$ in the ideal model.

Definition 5. *(Security against Rational Adversaries) Let Γ be an n-player ideal game specification and Π be an n-party real world protocol. We say that Π securely realizes Γ if there exists a set $\{\mathsf{Sim}_i\}_{i\in[1\ldots n]}$ of PPT transformations admissible in the ideal model such that for all PPT rational adversaries $\mathcal{A} = \{\mathcal{A}_i\}_{i\in[1\ldots n]}$ admissible in the real model, for all $\boldsymbol{x} \in (\{0,1\}^*)^n$ and $\boldsymbol{z} \in (\{0,1\}^*)^n$, and for all $i \in [1\ldots n]$,*

$$\left\{\mathrm{IDEAL}^{(i\in[1\ldots n])}_{\Gamma,\mathcal{R},\mathcal{P},\mathcal{S},\boldsymbol{\mu},z}(\lambda,\bar{\boldsymbol{x}};r)\right\}_{\lambda\in\mathbb{N},\bar{\boldsymbol{x}},r\in\{0,1\}^*} \overset{c}{\equiv} \left\{\mathrm{REAL}^{(i\in[1\ldots n])}_{\Pi,\mathcal{C},\mathcal{P},\mathcal{A},\boldsymbol{\mu},z}(\lambda,\bar{\boldsymbol{x}};r)\right\}_{\lambda\in\mathbb{N},\bar{\boldsymbol{x}},r\in\{0,1\}^*}$$

where $\mathcal{S} = \{\mathcal{S}_i\}_{i\in[1\ldots n]}$ is the set of simulators such that $\mathcal{S}_i = \mathsf{Sim}(\mathcal{A}_i)$, $\mathcal{I}$ is the information partition set for $\mathcal{P}$ and r is chosen uniformly at random.

Thus, to establish the security of a realized protocol Π, we must construct a simulator $\mathcal{S}_i$ for all players $\mathtt{P}_i \in \mathcal{P}$ such that for all probabilistic polynomial-time distinguishers $\mathcal{D}$, the distributions of $\mathcal{S}$ in the ideal world and $\mathcal{A}$ in the real world can only be differentiated with probability negligibly greater than $\frac{1}{2}$.

4 Demonstrating the Model on Rational Secret Sharing

To illustrate the power of our model, we return to the example of rational secret sharing. We demonstrate that, despite the presence of point-to-point communication channels, the original game specification is admissible in our ideal world

[9] That is, any probabilistic polynomial-time (PPT) distinguisher $\mathcal{D}$ cannot distinguish between an execution of Γ in the ideal world model and an execution of Π in the real world model with probability non-negligibly greater than $\frac{1}{2}$.

model, and realizable in the real world model. This violates the assumptions of existing security frameworks, which disallow point-to-point communication either within the protocol execution, outside of the protocol execution, or both.

4.1 Ideal World Game Specification

The ideal world game Γ is an interaction between a set of players $\mathcal{P} = \{\mathtt{P}_i\}_{i\in[1\ldots n]}$, where $\mathtt{P}_i$ has access to a point-to-point communication resource $\mathcal{R}_{\mathtt{P}_i,\mathtt{P}_j} \forall j \neq i$. That is, $\mathtt{P}_i$ may privately communicate with any other player $\mathtt{P}_j$. We now demonstrate that Γ is admissible in our ideal world definition.

Input Distribution:	Each player $\mathtt{P}_i \in \mathcal{P}$ receives its input share x_i, random coins r_i and auxiliary input z_i. Each player has the option of inputting a different share $\bar{x_i} \neq x_i$ or aborting the protocol at any time, as this is unavoidable.
"Cheap Talk":	Player $\mathtt{P}_i$ is free to collaborate with all players $\mathtt{P}_j \in \hat{\mathcal{P}}$ over $\mathcal{R}_{\mathtt{P}_i,\mathtt{P}_j}$, where $\hat{\mathcal{P}}$ is the set of colluding players. Proposition 1 demonstrates that communication over $\mathcal{R}$ is considered "cheap talk" (it *does not affect* the strategy selection of the player), and that the local simulator $\mathcal{S}_i$ for each player will select $a_i = \mathsf{reveal}$, as this maximizes μ_i.
Game Execution:	The `Mediator` instructs $\mathtt{P}_i, \forall i \in n$ to play their action a_i at each round k, where $a_i \in \{\mathsf{silent}^a, \mathsf{reveal}\}$.
Payoff Assignment:	At the terminal round k^* where the shares yield the secret, `Mediator` distributes the payoffs to $\mathtt{P}_i \in \mathcal{P}$.

[a] Note that selecting $a_i = \mathsf{silent}$ is equivalent to aborting.

Protocol 4.1. Ideal World Game Γ Execution

Let Γ be the ideal game specification for rational secret sharing, with player set $\mathcal{P} = \{\mathtt{P}_i\}_{i\in[1\ldots n]}$ and associated set of local simulators $\mathcal{S} = \{\mathcal{S}_i\}_{i\in[1\ldots n]}$ that select actions for players to maximize their local utility functions, resource set $\mathcal{R} = \{\mathcal{R}_{\mathtt{P}_i,\mathtt{P}_j}\}_{\forall i, i\neq j}$, and all players $\mathtt{P}_i \in \mathcal{P}$ have utility functions defined as

$$\mu_i(\sigma_i) \mapsto \begin{cases} (\mu^{++})(p) : \sigma_i = \mathsf{silent}, k = k^* \\ (\mu^{-})(1-p) : \sigma_i = \mathsf{silent}, k \neq k^* \\ (\mu^{+}) : \qquad \sigma_i = \mathsf{reveal} \end{cases} \tag{1}$$

where μ^+ represents positive utility, μ^- represents negative utility, and $\mu^{++} > \mu^+$ as players value exclusivity.

Proposition 1. *For all players $\mathtt{P}_i \in \mathcal{P}$ in Γ with utility function defined as $\mu_i(\sigma_i)$ in Equation 1, strategy $\{\sigma^*_{P_i} = \mathsf{reveal}\}_{\forall i\in n} > \{\sigma_{P_i} = \mathsf{silent}\}_{\forall i\in n}$ when $p < \frac{\mu^+}{\mu^{++}}$.*

Proof. In the original rational secret sharing protocol, the strategy $\boldsymbol{\sigma^*} = \{\sigma^*_{\mathtt{P}_i} = \mathsf{reveal}\}_{\forall i \in n}$ is the only Nash equilibrium, as the true final round k^* (where combining shares reveals the shared secret) is chosen from a geometric distribution. As the probability of correctly guessing the final round k^* is the parameter p, the expected utility for $\sigma_{\mathtt{P}_i} = \mathsf{silent}$ is at most $(\mu^{++})(p)$. We set $\mu^{++} > \mu^+$, as players are assumed to value exclusivity (recovering the secret while preventing other players from doing so). If a player remains silent in any round $k < k^*$, they are caught by the other players as a cheater and excluded from future rounds (receiving negative utility μ^-). By choosing p such that $p < \frac{\mu^+}{\mu^{++}}$, we have $(\mu^{++})(p) < \mu^+$ which implies $\mu_{\mathtt{P}_i}(\mathsf{silent}) < \mu_{\mathtt{P}_i}(\mathsf{reveal})$. Thus revealing the share for each round strictly dominates remaining silent. Players in our ideal model Γ may communicate over $\mathcal{R}$ and attempt to convince other players that they will select silent. This provides a greater degree of exclusivity, as only those colluding players in $\hat{\mathcal{P}} \subseteq \mathcal{P}$ will recover the secret. However, this communication is considered cheap talk, as each player maximizes μ_i by selecting $\sigma_i = \mathsf{silent}$ regardless of the messages sent over $\mathcal{R}$ when $p < \frac{\mu^+}{\mu^{++}}$.

4.2 Real World Protocol Construction

We now translate the ideal game specification Γ to a real world protocol Π, and demonstrate that there exist simulators such that the distribution of the ideal world game is computationally indistinguishable from the distribution of the real world protocol execution.

Input Distribution:	Each player $\mathtt{P}_i \in \mathcal{P}$ receives its input share x_i, random coins r_i and auxiliary input z_i. Each player has the option of inputting a different share $\bar{x_i} \neq x_i$ or aborting the protocol at any time, as this is unavoidable.
"Cheap Talk":	Player $\mathtt{P}_i$ is free to collaborate with all players $\mathtt{P}_j \in \hat{\mathcal{P}}$ over $\mathcal{C}_{\mathtt{P}_i,\mathtt{P}_j}$, where $\hat{\mathcal{P}}$ is the set of colluding players. Proposition 1 demonstrates that communication over $\mathcal{C}$ is considered "cheap talk" (it *does not affect* the strategy selection of the player), and that the local adversary $\mathcal{A}_i$ for each player selects $a_i = \mathsf{reveal}$, as this maximizes μ_i.
Game Execution:	Each player $\mathtt{P}_i \in \mathcal{P}$ selects and plays their action a_i at each round k, where $a_i \in \{\mathsf{silent}^a, \mathsf{reveal}\}$.
Payoff Assignment:	At the terminal round k^* where the shares yield the secret, each player $\mathtt{P}_i \in \mathcal{P}$ receives its associated payoff.

[a] Note that selecting $a_i = \mathsf{silent}$ is equivalent to aborting.

Protocol 4.2. Real World Protocol Π Execution

In the real world model, the communication resource $\mathcal{R}$ is replaced with a public key infrastructure $\mathcal{C}$. Each pair of players $(\mathtt{P}_i, \mathtt{P}_j) \in \mathcal{P}$ has access to a

private and authenticated point-to-point communication channel $\mathcal{C}_{ij}$. Let Π be a real world protocol, with player set $\mathcal{P} = \{\mathtt{P}_i\}_{i\in[1\ldots n]}$ and associated set of local adversaries $\mathcal{A} = \{\mathcal{A}_i\}_{i\in[1\ldots n]}$ that select actions for players to maximize their local utility functions, communication channel set $\mathcal{C} = \{\mathcal{C}_{ij}\}_{\forall i\neq j}$, and all players have identical utility functions defined as in Equation 1.

Clearly Π is admissible under the real world model, as the PKI infrastructure $\mathcal{C}$ facilitates the point-to-point communication channels between all players. The real world protocol Π for rational secret sharing proceeds as in Protocol 4.2. Again, the original equilibrium of $\boldsymbol{\sigma}^* = \{\sigma_{\mathtt{P}_i} = \mathsf{reveal}\}$ is preserved despite the presence of the communication channel $\mathcal{C}$.

4.3 Demonstrating Protocol Π Security

We use the simulation paradigm [33] to demonstrate the security of the construction by proving the distribution of the real world protocol is computationally indistinguishable from the ideal world distribution.

Theorem 1. *(Security of Π against Rational Adversaries) Let Γ be the n-party ideal world game specification of Protocol 4.1 and let Π be the n-party real world execution of Protocol 4.2. There exists a set $\{\mathsf{Sim}_i\}_{i\in[1\ldots n]}$ of PPT transformations admissible in the ideal model such that for all PPT rational adversaries $\mathcal{A} = \{\mathcal{A}_i\}_{i\in[1\ldots n]}$ admissible in the real model, for all $\boldsymbol{x} \in (\{0,1\}^*)^n$ and $\boldsymbol{z} \in (\{0,1\}^*)^n$, and for all $i \in [1\ldots n]$,*

$$\left\{\mathrm{IDEAL}^{(i\in[1\ldots n])}_{\Gamma,\mathcal{R},\mathcal{P},\mathcal{S},\boldsymbol{\mu},z}(\lambda,\bar{\boldsymbol{x}};r)\right\}_{\lambda\in\mathbb{N},\bar{\boldsymbol{x}},r\in\{0,1\}^*} \stackrel{c}{\equiv} \left\{\mathrm{REAL}^{(i\in[1\ldots n])}_{\Pi,\mathcal{C},\mathcal{P},\mathcal{A},\boldsymbol{\mu},z}(\lambda,\bar{\boldsymbol{x}};r)\right\}_{\lambda\in\mathbb{N},\bar{\boldsymbol{x}},r\in\{0,1\}^*}$$

establishing that Π securely realizes Γ.

Proof. To prove the security of Π against rational adversaries $\mathcal{A} = \{\mathcal{A}_i\}_{i\in[1\ldots n]}$ we must construct a set of simulators $\mathcal{S} = \{\mathcal{S}_i\}_{i\in[1\ldots n]}$ whose output in the ideal game specification Γ is indistinguishable from the output of $\mathcal{A}$ in the real world execution.

To achieve this, we construct simulators $\mathcal{S}_i = \mathsf{Sim}(\mathcal{A}_i)$ that simulate all messages and the output of $\mathcal{A}_i$ in the real world execution of Π, and is thus able to return these as its own. The simulated messages and output returned by $\mathcal{S}_i$ must be computationally indistinguishable such that, for all probabilistic polynomial-time distinguishers $\mathcal{D}$, the probability of differentiating the ideal world and real world distributions is at most negligibly greater than $\frac{1}{2}$.

Each simulator $\mathcal{S}_i$ will rely on the private communication resource $\mathcal{R}$ to simulate the messages exchanged and final output produced by $\mathcal{A}_i$ acting to maximize the utility function μ_i for player $\mathtt{P}_i$. The simulator $\mathcal{S}_i$ given in Construction 4.1 holds for all players $\mathcal{P} = \{\mathtt{P}_i\}_{i\in[1\ldots n]}$.

The construction relies on the computational indistinguishability of the real world communication channel $\mathcal{C}$ from the ideal world private and authenticated communication resource $\mathcal{R}$. All messages sent by simulators $\mathcal{S}_i \in \mathcal{S}$ in the ideal

world model are passed over $\mathcal{R}$. In the real world execution, messages are encrypted between players using the PKI communication resource $\mathcal{C}$. Thus, all probabilistic polynomial-time distinguishers $\mathcal{D}$ are able to distinguish the view of the ideal world execution from the real world execution with at most probability negligibly greater than $\frac{1}{2}$ by the security of the PKI communication resource $\mathcal{C}$.

Input Distribution:	The simulator $\mathcal{S}_i \in \mathcal{S}$ is given input share x_i, random coins r_i and auxiliary input z_i
"Cheap Talk":	The simulator $\mathcal{S}_i$ is free to communicate over $\mathcal{R}_{\mathcal{S}_i,\mathcal{S}_j}$ where $i \neq j$. $\mathcal{S}_i, \forall i \neq j$ must simulate the "cheap talk" between the other player's adversary $\mathcal{A}_j$. $\mathcal{S}_i$ uses its random coins r_i to construct a random message m, and sends m over resource $\mathcal{R}_{\mathcal{S}_i,\mathcal{S}_j}$. By definition, $\mathcal{R}$ is a private and authenticated point-to-point communication resource. Thus, the messages sent by the simulator are computationally indistinguishable from those sent in the real world execution, which are encrypted under the public key infrastructure communication resource $\mathcal{C}$. The local simulator $\mathcal{S}_i$ for each player selects $m_i = \mathsf{reveal}$, as this maximizes μ_i regardless of the messages exchanged during this phase.
Game Execution:	The simulator $\mathcal{S}_i$ sends a message m to $\mathcal{S}_j, \forall j \neq i$ over $\mathcal{R}_{\mathcal{S}_i,\mathcal{S}_j}$ with their decision, where $m \in \{\mathsf{silent}, \mathsf{reveal}\}$. By definition, $\mathcal{R}$ is a private and authenticated point-to-point communication resource. Thus, the messages sent by the simulator to $\mathcal{S}_j$ are computationally indistinguishable from those sent in the real world execution, which are encrypted under the public key infrastructure communication resource $\mathcal{C}$.
Payoff Assignment:	After $\mathtt{P}_j \in \mathcal{P}, \forall j \neq i$ has received $m_{\mathcal{S}_i}$, each simulator receives the payoff associated with the outcome.

Construction 4.1. Construction of Simulator $\mathcal{S}_i$

5 Conclusion

In this work, we have proposed a security definition capturing rational cryptographic protocols in the presence of standard point-to-point communication resources. Rather than limit the communication resources available to players, we answer the question of how game specifications admissible in an ideal model allowing point-to-point communication channels may be realized in practice. Thus, the ideal world model necessarily limits the class of games that are admissible and is not a general result. However, we have argued that point-to-point communication channels are unavoidable in real-world settings, and consequently must be incorporated into the definition of security. Further, we have demonstrated that not all game specifications forbidding point-to-point communication

are inadmissible under our model. We presented the transformation for the classic prisoner's dilemma, which disallows point-to-point communication through physical assumptions, into a modified game that is admissible under our model and preserves the original equilibrium. Similarly, we have demonstrated that the signaling game has an expected payoff of 1 when executed in the presence of point-to-point channels, rather than an expected payoff of $\frac{1}{2}$: a distinction not captured by models that disallow communication outside of the protocol execution. Finally, we have presented a full security proof for rational secret sharing under our proposed framework. Although our results are not universal, we have demonstrated a powerful benefit of our model: assigning local adversaries may aid mechanism design in destabilizing the formation of coalitions. Thus, there are tangible benefits from adopting our definition of security against local rational adversaries in the presence of point-to-point communication resources.

References

1. Goldreich, O., Micali, S., Wigderson, A.: How to play any mental game. In: STOC 1987: Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing, pp. 218–229. ACM, New York (1987)
2. Izmalkov, S., Micali, S., Lepinski, M.: Rational secure computation and ideal mechanism design. In: 46th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2005, pp. 585–594 (2005)
3. Alwen, J., Shelat, A., Visconti, I.: Collusion-free protocols in the mediated model. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 497–514. Springer, Heidelberg (2008)
4. Alwen, J., Katz, J., Lindell, Y., Persiano, G., Shelat, A., Visconti, I.: Collusion-free multiparty computation in the mediated model. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 524–540. Springer, Heidelberg (2009)
5. Lepinksi, M., Micali, S., Shelat, A.: Collusion-free protocols. In: Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing, STOC 2005, pp. 543–552. ACM, New York (2005)
6. Alwen, J., Katz, J., Maurer, U., Zikas, V.: Collusion-preserving computation. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 124–143. Springer, Heidelberg (2012)
7. Kamara, S., Mohassel, P., Raykova, M.: Outsourcing multi-party computation. Cryptology ePrint Archive, Report 2011/272 (2011), `http://eprint.iacr.org/`
8. Fuchsbauer, G., Katz, J., Naccache, D.: Efficient rational secret sharing in standard communication networks. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 419–436. Springer, Heidelberg (2010)
9. Halpern, J., Teague, V.: Rational secret sharing and multiparty computation: Extended abstract. In: Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing, STOC 2004, pp. 623–632. ACM, New York (2004)
10. Lysyanskaya, A., Triandopoulos, N.: Rationality and adversarial behavior in multiparty computation. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 180–197. Springer, Heidelberg (2006)

11. Canetti, R., Vald, M.: Universally composable security with local adversaries. In: Visconti, I., De Prisco, R. (eds.) SCN 2012. LNCS, vol. 7485, pp. 281–301. Springer, Heidelberg (2012)
12. Nielsen, J.B., Alwen, J., Cachin, C., Nielsen, J.B., Pereira, O.: Summary report on rational cryptographic protocols (2007)
13. Poundstone, W.: Prisoner's Dilemma: John Von Neumann, Game Theory and the Puzzle of the Bomb, 1st edn. Doubleday, New York (1992)
14. Wallrabenstein, J.R., Clifton, C.: Privacy preserving tatonnement: A cryptographic construction of an incentive compatible market. In: Financial Cryptography and Data Security. LNCS. Springer, Heidelberg (2014)
15. Gordon, S.D., Katz, J.: Rational secret sharing, revisited. Cryptology ePrint Archive, Report 2006/142 (2006), http://eprint.iacr.org/
16. Micali, S., Shelat, A.: Purely rational secret sharing (extended abstract). In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 54–71. Springer, Heidelberg (2009)
17. Groce, A., Katz, J., Thiruvengadam, A., Zikas, V.: Byzantine agreement with a rational adversary. In: Czumaj, A., Mehlhorn, K., Pitts, A., Wattenhofer, R. (eds.) ICALP 2012, Part II. LNCS, vol. 7392, pp. 561–572. Springer, Heidelberg (2012)
18. Groce, A., Katz, J.: Fair computation with rational players. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 81–98. Springer, Heidelberg (2012)
19. Zhang, Z., Liu, M.: Rational secret sharing as extensive games. Science China Information Sciences 56, 1–13 (2013)
20. Shamir, A.: How to share a secret. Commun. ACM 22(11), 612–613 (1979)
21. Maleka, S., Shareef, A., Rangan, C.P.: Rational secret sharing with repeated games. In: Chen, L., Mu, Y., Susilo, W. (eds.) ISPEC 2008. LNCS, vol. 4991, pp. 334–346. Springer, Heidelberg (2008)
22. Nojoumian, M., Stinson, D.R.: Socio-rational secret sharing as a new direction in rational cryptography. In: Grossklags, J., Walrand, J. (eds.) GameSec 2012. LNCS, vol. 7638, pp. 18–37. Springer, Heidelberg (2012)
23. Kol, G., Naor, M.: Games for exchanging information. In: Proceedings of the 40th Annual ACM Symposium on Theory of Computing, STOC 2008, pp. 423–432. ACM, New York (2008)
24. Kol, G., Naor, M.: Cryptography and game theory: designing protocols for exchanging information. In: Canetti, R. (ed.) TCC 2008. LNCS, vol. 4948, pp. 320–339. Springer, Heidelberg (2008)
25. Zhang, Z., Liu, M.: Unconditionally secure rational secret sharing in standard communication networks. In: Rhee, K.-H., Nyang, D. (eds.) ICISC 2010. LNCS, vol. 6829, pp. 355–369. Springer, Heidelberg (2011)
26. Ostrovsky, R., Yung, M.: How to withstand mobile virus attacks (extended abstract). In: Proceedings of the Tenth Annual ACM Symposium on Principles of Distributed Computing, PODC 1991, pp. 51–59. ACM, New York (1991)
27. Katz, J.: Bridging game theory and cryptography: Recent results and future directions. In: Canetti, R. (ed.) TCC 2008. LNCS, vol. 4948, pp. 251–272. Springer, Heidelberg (2008)
28. Osborne, M.J., Rubinstein, A.: A Course in Game Theory. MIT Press Books, vol. 1. MIT Press (1994)
29. Fudenberg, D., Tirole, J.: Game Theory. MIT Press (August 991)

30. Asharov, G., Canetti, R., Hazay, C.: Towards a game theoretic view of secure computation. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 426–445. Springer, Heidelberg (2011)
31. Dodis, Y., Halevi, S., Rabin, T.: A cryptographic solution to a game theoretic problem. In: Bellare, M. (ed.) CRYPTO 2000. LNCS, vol. 1880, pp. 112–130. Springer, Heidelberg (2000)
32. Miltersen, P.B., Nielsen, J.B., Triandopoulos, N.: Privacy-enhancing auctions using rational cryptography. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 541–558. Springer, Heidelberg (2009)
33. Goldreich, O.: Foundations of Cryptography, vol. 2. Cambridge University Press (2004)