

Purely Rational Secret Sharing (Extended Abstract)

Silvio Micali¹ and abhi shelat²

¹ MIT CSAIL (silvio@csail.mit.edu)

² U. Virginia (abhi@virginia.edu)

Abstract. Rational secret sharing is a fundamental primitive at the intersection of cryptography and game theory. In essence, a dealer wishes to engineer a communication game that, when rationally played, guarantees that each of the players learns the dealer’s secret. Yet, all solutions so far were *quite inefficient* and relied on the players’ *beliefs* and not just on their rationality.

After providing a more complete definition of this problem, we exhibit a very efficient and purely rational solution to it with a verifiable trusted channel.

1 Introduction

In [LMPS04], Lepinski, Micali, Peikert and shelat put forward the notion and the first implementation of *Fair Secure Function Evaluation*. This is a communication protocol extending the traditional notion of secure function evaluation [GMW87]. In essence, a Fair SFE is an SFE in which either (1) all players learn the result of evaluating a given function on their secret inputs (but no other information about their inputs) or (2) none of them learns anything. The first outcome is reached when all players want it, and the second one when at least one of the players wants it. The difficulty lies in the fact that such objectives must be reached no matter what the function may be and no matter how many the player are, provided that at least one of the players is *honest*, that is sticking to his communication instructions in all cases.

In [HT04], Halpern and Teague put forward the notion of *rational secret sharing*, aiming at distilling separately, and in purely game theoretic terms, the last stage of a Fair SFE (where the players attempt to reconstruct the specified output from their shares of it). We believe this to be a very valuable contribution, but we also believe that the meaningfulness of a rational secret would benefit from a more stringent definition.

In this extended abstract we shall solely deal with the two-player version of the notion, arguably the best way to highlight the novel and most poignant aspects of the problem.³

³ Our approach easily extends to n players, where the dealer wishes that n out of n of them learn the secret. The k -out-of- n definition of traditional secret sharing is very relevant for robustness, and protects against the potential loss of shares, but is quite orthogonal to the rationality problem at hand, and is quite a distraction. Indeed, in a “ k -out-of- n ” rational secret sharing, any k players could prevent the others from learning their secret so as to maximize their utility (assuming that the fewer the players knowing the secret, the more value to them). Is this the natural wish of the dealer? Moreover, when k players try to reconstruct the secret among themselves they will face the real setting of a k -out-of- k rational secret sharing.

1.1 Rational Secret Sharing as a Special Form of Mechanism Design

The Intuitive Notion. At the verge of dyeing, a *dealer* possessing a secret string S wishes to ensure that two players will later on be able, *together*, to learn S . To this end, he provides each player i with his share of the secret, a string S_i . Each share is individually meaningless (i.e., its distribution is independent of S), while collectively the two shares reveal S . If the players were both honest, the dealer's goal could be trivially achieved. Unfortunately, honesty is not an available commodity: each player is assumed to be *rational* (i.e., always trying to maximize his own utility), and the utilities that the players attach to the possible ways of learning S are quite problematic. In particular, each player prefers best to be the only one to learn S , less to learn S along with some other players, and even less not to learn S at all. Accordingly, the dealer wishes to choose the shares such that,

for a suitable communication channel, there exists a communication game that, when rationally played, yields the secret to both players.

It is thus worth quickly recalling the traditional notions of solving a game.

Game Solution Concepts. Given a game G , a solution concept essentially is a way of predicting how G will be played. From the cryptographic perspective of the authors, traditional "solution concepts" are only partially meaningful, as they are stated from the perspective of *individual* players, disregarding *collusion* altogether. Nonetheless their meaningfulness is intact for the problem at hand, if restricted to just two players. (Which is indeed a main reason to focus on this restriction.)

The strongest, traditional solution concept is that of solvability in *dominant strategies*. Here, each player has a strategy σ_i that is best of him, no matter what strategies the other players may use. In such a case predicting that each player i will play σ_i is indeed the strongest form of prediction. Note that in choosing such a σ_i , each player i does not need to rely on the rationality of the other players, but just be rational himself! Unfortunately, not all games admit dominant-strategy solutions.

The next best solution concept is *dominant solvability*. Let us explain (again in the two-player setting). In a game G , a strategy a for player A is said to weakly dominate another strategy a' for A if (1) for all possible strategies b of player B , A 's utility under a is greater than or equal to his utility under a' ; and (2) for at least one strategy b' of B , A 's utility under a is strictly greater than his utility under a' . This being the case, a rational A should remove strategy a' and all of his weakly dominated strategies from consideration. And a rational B should do the same on his side. Trusting that B has done so, A should then eliminate from his remaining strategies all those that now become weakly dominated (relative to the strategies left over to B). And so on, until neither player can eliminate any more strategies. If at that point player A is left with a single strategy a and player B with a single strategy b , then G is called *dominant solvable*, and (a, b) is a very strong prediction for the way in which G will be rationally played. Notice, however, that this second solution concept is weaker than the first one, since each player must rely not only on his rationality, but also on that of his opponent. Again, not all games admit such a strong solution concept.

The next solution concept in line is a Nash equilibrium. This is a pair of strategies (a, b) , a for A and b for B , such that a is the best strategy for A if he believes that

B will play b (and symmetrically for B). The good news is that each game admits such equilibria, but the bad news is this is a *very distant third* among these solution concepts. The meaningfulness of a Nash equilibrium in fact depends not only on the rationality of both players, but also on their *beliefs*. Typically a game has a plurality of equilibria, often having symmetric payoffs, making it very uncertain to predict which of them will be played. Furthermore, the game may easily not end up in equilibrium at all. If A believes that equilibrium (a, b) will be played, while B believes (a', b') , then the strategy profile ultimately played may be (a, b') which needs not to be an equilibrium at all!

Mechanism Design. Very roughly said the goal of mechanism design is to engineer a game so that, “at equilibrium”, a given property P is guaranteed to hold. The quality of such design therefore crucial depends on the solution concept adopted: it is very meaningful when the game has dominant strategy solution, it is very meaningful when the game is “dominant solvable”, and has only very limited meaningfulness when the game is “Nash solvable.” The latter is indeed the case even if P is *guaranteed to hold at each of possible Nash equilibria of the game*. In a sense, if the game has k Nash equilibria, then —due mismatched beliefs— it roughly has k^2 (k^n if there are n players) possible ways not to end in any equilibrium.

Another measure of quality in mechanism design is the amount of knowledge about the players (e.g., knowledge about their utilities) required to engineer the game. Indeed, since precise knowledge about the players may not be available or too expensive to gather, the lesser the knowledge required from the designer the better.

Rational Secret Sharing and Mechanism Design. We propose to view *rational secret sharing as a special mechanism-design problem*. The property to guarantee in a rational play is “all players learn the secret”. And the engineered game should be a purely communicational one. Namely, (1) the only actions the players can take consist of exchanging messages over a special channel, (2) no trusted party can be involved in the game, and (3) no exogenous punishments, fines, etc. can be triggered by the final outcome: the players’ utilities must solely depend on who has or has not learned the secret.

This point of view enables us to extend to rational secret sharing the same quality analysis applicable to mechanism design, thus enabling a more meaningful comparison among various RSS protocols primarily focusing on their game theoretic properties.

1.2 Prior Solutions

In their quoted paper, Halpern and Teague present a protocol for the 3-out-of-3 case (and then show how to modify it for the m -out-of- n case, where $3 \leq m < n$). Their protocol guarantees that all players learn the secret at a Nash equilibrium whose strategies survive the iterated elimination of weakly dominated strategies (IEWDS for short). Rather than the *swap channel* of Lepinski, Micali, Peikert, and Shelat, they rely on *simultaneous-broadcast channels*, and prove that no rational secret sharing protocol can be fixed-round with such channels. A main limitation of their protocol is that the trusted dealer continues to be an active participant. (In most settings, such a dealer could directly inform the players of what the secret is.)

Gordon and Katz [GK06] present a protocol for just two players that dismisses the need for the periodic involvement of the dealer. Their protocol too relies on simultaneous-broadcast channels, and guarantees that all players learn the secret at a Nash equilibrium whose strategies survive IOWDS. Abraham, Dolev, Gonen, and Halpern [ADGH06] present a similar protocol, but focus on defining (and protecting against) coalitions of rational players.

Lysyanskaya and Triandopoulos [LT06], with the same channels and implementation type, consider a model with a mix of rational and malicious players.

Kol and Naor [KN08b] present a quite different protocol with simultaneous-broadcast channels, which guarantees that all players learn the secret at a “strict Nash equilibrium”, a locally stronger version of a Nash equilibrium. (In essence, any player deviating from his own equilibrium strategy expects to receive a strictly smaller utility.)

A Separate Protocol. We wish to mention an interesting and recent protocol of Ong, Parkes, Rosen, and Vadhan [OPRV08]. Their protocol however works in a quite different model. On one side, it does not require any special channels (that is, it relies on ordinary broadcast channels rather than simultaneous-broadcast ones). On the other, it relies on the *honesty* of a few players. (As we focus solely on rational players, we shall not include this protocol in any future discussion or comparison.)

1.3 Weaknesses of Prior Solutions

Protocol Inefficiency and Excessive Designer Knowledge. The prior protocols share the following logical structure. The players interact in several rounds, using some special channels. The protocol has a special round R , unknown to the players because it is secretly selected by the dealer according to a given distribution. If no player “cheats” then all players learn the secret. A player can successfully cheat only if he correctly guesses R . If a player i erroneously guesses R , then no one learns the secret (which gives i utility u_i). But if i guesses R correctly (and acts appropriately), then he is the only one to learn the secret (which gives him utility U_i).

In essence, therefore, to hope that it is rational to stick to the protocol’s prescribed strategies without cheating, letting p be the probability of successfully guessing R , p needs to be so small that $p \cdot u_i \leq (1-p)U_i$. This shows two separate weaknesses of these protocols. First, because properly engineering the game implies properly selecting p , the designer needs to know the u_i ’s and the U_i ’s quite accurately. (Thus, from a mechanism design perspective, this diminishes the quality of these approaches. In particular because they cannot be used when a good knowledge of the players’ utilities is not available.) Second, this implies that all prior protocols run in exponential time. In fact, independent of the distribution according to which R is selected, the expected number of rounds of the prior protocols must be exponential in k , assuming as it is natural that all players utilities are presented in binary, and that their length is k .

This inefficiency alone would be sufficient reason to seek new protocols. But another problem lurks in the background.

Limitation of Prior Designs. Prior solutions ensure that the property “all players learn the secret” holds at a given Nash equilibrium NE of the engineered game. As discussed

for general mechanism design, however, this assurance is far from guaranteeing our property for two separate reasons: *equilibrium selection* and *equilibrium absence*. Let us discuss the first reason first. Even if one were certain that the engineered game will end up in an equilibrium, he could not be certain of which equilibrium would be actually selected. And since, in the engineered games of the previous works, our property was guaranteed only at one of the possible Nash equilibria, there is a real possibility that the equilibrium ultimately selected is not one in which all players learn the dealer’s secret. Let us now discuss the second reason. The meaningfulness of any Nash equilibrium is inextricably to a given set of beliefs of the players. Thus, even if all players learned the secret at each Nash equilibrium of the engineered games, there is no guarantee at all that the engineered game ends up in equilibrium. Because the players’s beliefs may not properly match, the first player may believe that Nash equilibrium (a, b) will be played out, and thus rationally (based on this belief!) plays strategy a . The second player may believe that equilibrium (a', b') will be played out, and thus chooses to play strategy b' . And the final outcome (a, b') may not be an equilibrium, let alone guarantee that all players learn the secret.

To be sure the prior protocols ensured that all players learned the dealer’s secret not just at a generic Nash equilibrium, but at one whose strategies survived IEOWDS. But while of value, this additional quality of the Nash does not alleviate the above concerns at all. As long as multiple Nash survive IEOWDS (which is the case in prior protocols), then equilibrium selection and equilibrium absence will continue to poison the landscape.

To be sure too, some of the prior RSS protocols guaranteed that all players learned the secret at a an even stronger type of equilibrium, such as the strict Nash of [KN08b]. But these equilibria are in a sense only “locally stronger.” That is, if the players believe that a strict Nash equilibrium E will be played out, they would have “even less incentives” —so to speak— of deviating, but this does not mean that the players have “stronger beliefs” —so to speak again— that E will be played out. Accordingly, strict Nash equilibria too offer little protection against the main problems of equilibrium selection and equilibrium absence.

In sum, *in all prior works, the property that all players learned the dealer’s secret did not solely depend on the players’ rationality, but also on their beliefs*. And thus all prior protocols suffered from all the weakness entailed by the latter dependency.

1.4 Our Contributions

Our contributions can be summarized as follows.

- *Modeling*. We put forward a more complete modeling of the RSS problem. In particular: we highlight the inputs available to the designer of protocol; provide a more comprehensive set of utilities —including the possibility of learning the wrong secret—; highlight the necessity of modeling RSS as a potentially infinite communication game; provide a very general definition of a communication channel; highlight the necessity of worrying about other channels even in a communication game designed for a specific channel; provide the first rationalization of aborting in a communication game; and bring to the fore the necessity of including bargaining into the definition of RSS.

- *Purely Rational Implementation.* Our RSS protocol is an *implementation in surviving strategies*, as put forward by Chen and Micali [CM08]. In essence, such an implementation is “equilibrium-less.” It guarantees that the desired property holds for any combination of strategies surviving IEOWDS. Implementation in surviving strategies thus implies that *the desired property is guaranteed based solely on the rationality of the players, and not on their beliefs*. In a sense, as long no player chooses a dumb strategy the property holds.

Actually, our protocol satisfies a stronger notion of implementation: namely, the surviving strategy of each player is essentially *unique*.⁴ That is, in our RSS protocol, after iteratively deleting all weakly dominated strategies, essentially a single strategy survives for each player, and playing these two strategies guarantees that both players learn the secret. That is, *our RSS protocol essentially is a dominant solvable game*.

Note that IEOWDS often eliminates very few strategies (a fact that has been used to argue that Nash equilibria that survive IEOWDS is a solution concept not really better than an ordinary Nash). Thus it is even more remarkable that our protocol is such that, for any player, all but one strategy is “rationally credible.”

Note too that, in general, which strategy survive IEOWDS depends on the order in which the strategies are eliminated. In our case, however, the (essentially) unique surviving strategy of a player is the same irrespective of any possible elimination order.

In sum, our solution concept is indeed very strong.

- *Communication Channel and Security.* Our communication channel uses only ordinary envelopes (as a way of temporarily and perfectly hiding a secret value) and the dealer’s public key.

The security depends on the ability of envelopes to perfectly hide their content and unforgeable digital signatures.

- *Operational Efficiency.* Ours is the first polynomial-time RSS protocol, fully accounting for all inputs. In fact, each surviving strategy requires a total of $10k$ envelope operations, $4kL$ bit operations, plus the time of verifying two signatures relative to k -bit public keys. Here L is an upperbound to the length of the binary representation (of the absolute value) of any of the players’ utilities, and k is a security parameter. The security parameter k controls the probability that something goes wrong. (The probability of something going wrong is guaranteed to be exponentially small in k .)

The total number operations required by the dealer consist of $4kL$ bit operations, the time necessary to generate matching public and secret keys of a digital signa-

⁴ The reason that we do not say unique is that, as we shall argue, a pure communication game G should be modeled as a possibly infinite sequence of the same sub-game g . Thus, a strategy of any player in G actually consists of a sequence of strategies, $\sigma_1, \sigma_2, \dots$, where σ_j is the player’s intended strategy for j th copy of g , if reached. By saying that each player has an essentially unique surviving strategy in G we mean that any of his surviving strategies is of the form s, σ_2, \dots , where s is fixed; that is first sub-strategy is the same for any surviving strategy of the player. And when all players play their first such strategies, G terminates.

ture with security parameter k , and the time of producing two signatures relative to k -bit public keys.

- *Round Efficiency.* A play of our surviving strategies involves only 6 rounds (1 for the players, and 5 for the channel).

2 Selected Modeling Issues

Dealer Secret, Player Outputs, Player Utilities, and Designer Knowledge. For concreteness, we model the secret as a uniformly selected string of n bits. (Our protocol of course works for all kinds of other distributions as well.)

We assume that, upon termination, each player outputs either an n -bit string (interpretable as the player’s guess for the dealer’s secret) or the special symbol “?” (interpretable as the player’s having no information about the secret). The protocol terminates when a prespecified stage is publicly reached, or when either one of the players *aborts*, that is stops communicating and for ever takes no further action —after setting his own output.⁵

We define an outcome of an RSS protocol to consists of the outputs of both players, distinguishing three types of outputs for each player: (1) the correct secret string of the dealer, (2) the symbol “?”, and (3) an incorrect string. We assume that each player prefers them in this order, and prefers the inverse order for the output of the other player. That is, denoting by K_i (for “ i knows the secret”) an output of the first type, and by W_i (for “ i wrongly learns the secret) an output of the third type, and by u_i the utility function of player i , we assume that the utilities of the first player over the possible 9 outcomes are as follows:

$$\begin{aligned} u_1(K_1, W_2) &\geq u_1(K_1, ?) \geq u_1(K_1, K_2) \geq \\ u_1(? , W_2) &\geq u_1(? , ?) = 0 \geq u_1(? , K_2) \geq \\ u_1(W_1, W_2) &\geq u_1(W_1, ?) \geq u_1(W_1, K_2). \end{aligned}$$

Player 2’s utilities are symmetrically defined. (Setting the players’ utilities to 0 when both of them have no information about the secret is somewhat arbitrary, but concretely useful to fix our thoughts.) All of the above inequalities can be strict. But for our analysis it suffices that $u_1(K_1, ?) > u_1(K_1, K_2) > u_1(? , ?) > u_1(W_1, W_2)$, and symmetrically for player 2. That is, each player prefers learning the secret alone to learning together with the other player, prefers the latter to not learning the secret, and prefers the latter to learning the wrong secret.⁶ It is also useful to assume that a player’s expected utility when randomly guessing the secret is negative. (Alternatively, we must

⁵ That is, we explicitly assume that one players’ aborting is detectable by the other player. (After all, stopping all communications should be “eventually detectable” in practical settings, and immediately detectable in synchronous ones.) Alternatively, each player may keep track of his current output at all times (rather than producing his output at termination). This way if a player aborts without the “courtesy” of informing the other player, the latter’ output is properly set.

⁶ Indeed, if the secret were the combination of a safe with money and a bomb inside, and the safe exploded when the wrong combination were entered, learning the wrong secret could have truly negative utility for a player!

ensure that the utility of random guessing is less than that of learning the secret together with the other player. Else, a player would not have any incentive to participate in an RSS protocol.)

This structure of the utility is assumed to be known to the dealer/designer. And so is an upperbound to the number of bits necessary to write down the largest of the 16 possibilities of the players. (In other words, it suffices for the designer to know the players' utilities within an exponential accuracy, rather than the linear accuracy of the prior works.)

Ensuring the Rationality of Abort. Rational secret sharing is a purely informational game, over a proper communication channel, where the dealer provides the players with their shares of a secret S . Our protocol, if a special point in which a player i has not yet learned the secret is reached, calls for him to abort. By so doing, of course, the player loses any hope of learning the secret. Thus, in order to guarantee that the suggested strategy survives IOWDS, we need to ensure that, at that point, the player no longer has any rational hope of learning the secret (whether alone or together with the other player). What should this mean? In particular, of course, it should mean that i 's expected utility when continuing the current execution of the protocol is worse than that of aborting outputting "?". But it *should not* mean just that. The dealer who has provided the players with their shares is now dead, and can no longer control what the players do from his grave. Thus, not only have the players all the information they need to continue any given execution of our protocol (if they so want), but they also have the ability of starting another execution from scratch. (For instance, they may use their same shares, but different coin tosses for their strategies, if probabilistic. Alternatively, if reusing the old shares is not "rationally advisable," they may first resort to a secure function evaluation to "compute new, equivalent and, independently selected shares from their old ones, and then execute our protocol again. The possible alternatives abound.) Better yet, perhaps, they also have the ability to start a totally different RSS protocol using the same communication channel. More generally yet (unless one were ready to make the outlandish assumption that no other channel exists), they have the ability to execute a totally different RSS protocol with a totally different channel! In sum,

*To rationalize player i 's aborting in an RSS protocol, we should **prove** that any chance for i to learn the secret, by any envisaged or un-envisaged means, is gone for ever.*

Realizing, formalizing, and indeed delivering this property is a main contribution of our work.

Modeling Special Channels. As discussed already, all RSS protocols with rational players have been designed for some special communication channel, such as a swap channel, or a simultaneous-broadcast channel. Since we have just argued that a proper analysis of RSS should include the possibility of running a different protocol over a different channel, it becomes imperative to model any possible special channel of communication. We do so by letting special channels consist of "mildly trusted parties in abundant supply." Let us explain. If some party T could be totally and universally trusted, then many problems (including rational secret sharing) would be trivialized. For instance, the dealer might as well confide his secret to T and ask him to reveal it when all

the designated players show up together. Thus “mild trust” became imperative. As for abundant supply we mean that there is not a unique mildly trusted party in the world. (If this were the case, one might ask T to interact only once with a given group of players for a given task, and simplify a lot of things too.) By contrast, to model the fact that a special communication channel (if it exists at all) is indeed a commodity purchasable at any store, we envisage that there is a plurality of mildly trusted parties, not aware of—or not in contact with—each other.

Accordingly, following [ILM08], we model a mildly trusted party in abundant supply as a *verifiable trusted party* (VTP for short) with no memory. By verifiable we mean that every one can see the actions a VTP takes and verify that they are the prescribed ones. That is, a VTP is not trusted to keep, nor to correctly make any secret actions. A VTP knows nothing and acts publicly, so that he is trusted only to the extent that he will indeed publicly perform his prescribed public actions.

For example, a VTP can trivially implement a swap channel between two parties as follows. First each of A and B seals his message for the other into an opaque envelope and publicly gives it to the VTP. Then the VTP publicly hands A ’s envelope to B and B ’s envelope to A .

As for another example, a VTP can implement a simultaneous-broadcast channel as follows. First, A and B seal their respective messages for the other in two envelopes and publicly hand them to T , then T publicly opens both of them.

In sum, VTPs can be viewed as a formalization of a legal system. One may not want to trust his secrets to—say—a judge, but should at least trust a judge to carry out under public scrutiny a specified sequence of totally public actions. Since typically there are multiple judges to choose from, the analogy with the legal system makes it clear that the players can always walk to a new judge to execute their protocol one more time. The analogy also makes it clear that if one type of channel is available, then indeed other types are likely to be available too. Whether or not, as *functions*, the “swap channel” is reducible to the “simultaneous-broadcast-channel” (or viceversa), from the VTP perspective, both exist. (Indeed any judge can, with envelopes, implement both channels and a host of similar ones.) This highlights the point that when a player is asked to abort, then it really must be the case that no hope to resurrect the secret exists for him, no matter what other protocol and channel might be considered.

Adding Costs to the Model. Consider a cryptographic rational secret sharing protocol in which the dealer also announces an encryption E of the secret S . Then, a player, in addition to any other strategy, also has available a computational-attack one: namely, abort and try to decrypt E . A computational-attack strategy is also possible in our protocol, but in a more complex way. Indeed, successfully forging a given value enables a player to learn the secret alone, and force the other to learn a false secret. Thus we too need to argue that computational-attack strategies are not rational. One way to do so is to define a computationally bounded version of rational secret sharing. A preferable way is to *attach cost to computation* so as make it preferable for a player to play honestly our protocol rather than try to attack the signature scheme and then, if the attack is successful, getting an advantage in the protocol. Details will be provided in the final version. (In any case, as argued by Halpern and Pass [HP08] considering computational

costs may be meaningful even for more traditional —i.e., non-cryptographic— game theoretic settings.)

We also associate a small additive cost of γ to each use of the channel. (E.g., every one has the right to access the legal system, but incurs a fixed cost in doing so.)

We note that additive (or multiplicative) discounts of the players final utilities are quite standard in game theoretical models in which the players could go on interacting (possible even for ever), typically by executing a given sub-game.⁷

The Issue of Bargaining. Finally, let us bring to the reader attention a point totally neglected so far. Traditionally, to guarantee the dealer’s wish that all players learn the secret (at least when everyone behaves rationally), the only restrictions envisaged for the utilities are *local to each player* (e.g., each player must prefer reconstructing the secret alone to reconstructing it together with the other player, etc.). That is, the utilities of an individual player must be “compatible with each other,” but not with those of other players. We wish to point out, however, that it is necessary to consider *inter-player* restrictions on the players’ utilities, or be ready live with the consequences relative to the dealer’s wishes. Let us explain.

A dealer providing players with shares of his secret S automatically enables them to *bargain*. In a bargaining situation, one player may get a better deal than others *without any failure of rationality*. For instance, in an RSS context, Player 1 may simply insist that unless everyone plays a protocol in which he learns the secret *alone* 99 times out of 100, he is not going to cooperate. (In a sense, if to Bill Gates learning the secret together with you and me is worth \$1K, but learning it alone is worth \$1B, then he would be wasting time and opportunity costs in participating with you and me in a “fair” reconstruction of the secret. Therefore, he may successfully bargain for a higher probability of learning the secret.) Now, if the dealer indeed has come up with shares and channels enabling the players to rationally reconstruct the secret together using a given special communication channel, then we should also expect that —whether with the same or with a different channel— the players can use their same shares to skew the payoffs so as to suit their bargaining needs. Truly unbelievable assumptions must be made to prevent the shares to be used in this alternative manner (especially in light of the result of [ILM08], that essentially enables the players to do rationally almost anything, although not too conveniently). Thus, either one must make the additional assumption that the players utilities are such that their bargaining game has a unique solution (e.g., some form of symmetry), or the dealer must be ready to die in peace with the comfort that either all players (if rational) will learn the secret, or that he has put all of them on a *technically* equal bargaining position.

The reader is free to pick the assumption he prefers. But always guaranteeing that all players together learn the secret may not be possible. For the rest of this extended abstract let us assume that the utilities are such that there is a unique bargaining solution.

⁷ For instance, if a given contract is executed after i days of negotiation it is worth less to the players than executing the same contract as $i - 1$ days of negotiation.

3 Our Enriched Solution

It is actually simpler to explain first our protocol assuming that a VTP can operate not just on ordinary envelopes and digital signature schemes, but also on a special, dealer-sealed, envelopes, and then to show how these special envelopes can be replaced by ordinary ones and digital signatures. A dealer-sealed envelope is such that anyone can verify that it has been sealed by the dealer, and thus that its content is what the dealer wanted it to be. Any attempt to break the dealer's seal is guaranteed to be detectable by anyone.

Notice that, if such special envelopes were available, then a trivial solution to the RSS problem exists. In essence, the Dealer creates two random shares of the secret s , that is two random strings s_A, s_B such that $s = s_A \oplus s_B$, and then provides player A (respectively B) with infinitely many pink (respectively, blue) dealer-sealed envelopes, each containing s_A (respectively s_B).

Players A and B interact with a VTP as follows. First, each player, simultaneously with the other, gives the VTP one of his dealer-sealed envelopes. Then, if the VTP receives both a pink and a blue dealer-sealed envelope, he publicly opens both of them. Else (e.g., one of the envelopes is ordinary, or has a broken seal), he destroys all envelopes received. In either case, the players incur a positive cost for this interaction.

It is not hard to prove that the above indeed is a RSS protocol, and that it actually is game solvable in dominant-strategies. The problem, however, is that we see no way of replacing dealer-sealed envelopes with ordinary ones and digital signatures in the simple protocol above. We thus now describe a more complex protocol for which we can "simulate" dealer-sealed envelopes.

In order to guarantee implementation in surviving strategies, our protocol critically introduces an asymmetry in the way the players are treated.

3.1 Dealer's Instructions:

On input an ℓ -bit secret s and a security parameter k' , do:

1. Choose a random string $\sigma \in \{0, 1\}^\ell$ and compute $s' \leftarrow s \oplus \sigma$.
2. Choose a value k such that for all i
 - (a) $u_i(\text{both}) > (2^{-k/2}) u_i(\text{alone}) + (1 - 2^{-k/2}) u_i(\text{empty})$
3. For $i = 1, 2, \dots, k$, repeat the following
 - (a) Randomly select a four-tuple (a_0, a_1, b_0, b_1) such that a_0, b_0 are a random \oplus -sharing of the secret s' and a_1, b_1 are random and independent values of the same length as s .
 - (b) Pick two random bits $e_1, e_2 \leftarrow \{0, 1\}$.
 - (c) Player 1's share is (a_{e_1}, a_{1-e_1}) and Player 2's share is (b_{e_2}, b_{1-e_2}) .
 - (d) Player 1's check value is $C_{1,i} = (e_2, b_1)$ and player 2's check value is $C_{2,i} = (e_1, a_1)$.
 - (e) Place value a_j into envelope $E_{1,i,j}$ and place value b_j into envelope $E_{2,i,j}$ for $j \in \{0, 1\}$.
4. Let C be the $k(\ell+1)$ -bit number corresponding to the check values $C_{2,1}, \dots, C_{2,k}$. Choose random values $\alpha, \beta \in \mathbb{Z}_k$ and compute the message authentication code $\gamma = \alpha \cdot C + \beta$.

5. Place into an envelope $E_{1,0}$ the values $(C_{1,1}, \dots, C_{1,k}, \alpha, \beta)$ and into an envelope $E_{2,0}$ the values $(C_{2,1}, \dots, C_{2,k}, \gamma)$. Seal the envelope $E_{1,0}$.
6. Place into an envelope $E_{p,\sigma}$ the value σ for $p \in \{0, 1\}$.
7. Send the player 1 the envelopes $E_{1,0}, E_{1,\sigma}$ and $E_{1,i,j}$ for $i \in [1, k]$ and $j \in \{0, 1\}$. Send to player 2 the envelopes $E_{1,0}, E_{1,\sigma}$ and $E_{1,i,j}$ for $i \in [1, k]$ and $j \in \{0, 1\}$.

3.2 Reconstruction Instructions

Recall that a player's strategy consists of a Turing machine that on input a history h outputs either a special symbol \perp to indicate abort, an output string s , or a sequence of $2k + 1$ strings to place into envelopes that are submitted to the VTP. We use the symbol ε to denote the initial history consisting of only the envelopes received from the dealer.

Player p instructions $T(h)$:

1. If $h = \varepsilon$, then submit envelopes $E_{p,0}$ and $E_{p,i,j}$ for $i \in [1, k]$ to the VTP. If the VTP destroy the envelopes, output \perp and stop. Else, after the VTP completes all of its steps, reconstruct n candidates of s by xor'ing the non-check values that have been opened. Let s' be the majority candidate. If no majority exists, then output \perp . Otherwise, privately open envelope E_σ and output $s' \oplus \sigma$.
2. For all other histories, output \perp (i.e. do not invoke the VTP).

VTP Instructions :

1. Publicly verify envelope $E_{1,0}$. If the envelope's seal does not verify, then destroy all envelopes. Otherwise, publicly open the envelope to reveal the values $(C_{1,1}, \dots, C_{1,k})$ and α, β .
2. Publicly open envelope $E_{2,0}$ to reveal values $C = (C_{2,1}, \dots, C_{2,k})$ and γ . If $\gamma \neq \alpha \cdot C + \beta$, then destroy all envelopes.
3. Open the check envelopes (left or right) of player two indicated by $C_{1,i}, \dots, C_{1,k}$. If there exists an opened envelope $E_{2,i,j}$ that does not match its stated value in $C_{1,i}$, the check fails: destroy all envelopes.
4. Open the check envelopes (left or right) of player one indicated by $C_{2,i}, \dots, C_{2,k}$. If there exists an opened envelope $E_{1,i,j}$ that does not match its stated value in $C_{2,i}$, the check fails: destroy all envelopes.
5. If all k checks succeed, open the remaining $2k$ envelopes (corresponding to shares of the secret s').

3.3 Analysis

Theorem 1. *The strategy profile (T, T) for players 1 and 2 constitute a profile that uniquely survives the iterated deletion of weakly dominated strategies in the given VTP model.*

The main idea of the proof. Unless the first envelope submitted by the first player is sealed correctly, the VTP destroys envelopes. Once the one-and-only sealed envelope $E_{1,0}$ is opened, the second player knows which of her share values are check values, and which are values that are used in the sharing of s' . If the VTP succeeds in the same use that $E_{1,0}$ is opened, then both players learn the secret. If it does not, then some check envelope has failed and therefore no share value has yet been opened. In subsequent uses of the VTP, the second player can then modify all of her share values by XORing a random string r to them. This action is undetectable by the first player. Moreover, this action is the weakly dominant response for player 2 since player 2 prefers to learn the secret alone. Therefore, the first player has no hope to recover the secret (since any future opened share values will be independent of the real secret s'). Thus, the first player will abort in every subsequent use of the VTP. As a result, it is best for the second player to submit the envelopes received from the dealer on the first use (since either her envelopes are never opened, or they are opened in the first and only rational opportunity there will be to recover the secret). In this case, the first player should follow T since each use of the VTP incurs a small cost. Then finally, the second player should also play T .

Definition 1. A revealing history h is a history in which the envelope $E_{1,0}$ has been opened and verified in some use of the VTP, but in every use of the VTP, all envelopes have been destroyed.

Let X_1 be the set of all player-one strategies, and X_2 be the set of all player-two strategies. Notice that for all $\sigma \in X_1$, $u_2(\sigma, T) \geq u_2(\text{nolearn})$ and for all $\tau \in X_2$, $u_1(T, \tau) \geq u_1(\text{nolearn})$. Therefore in the first step of removal, all guessing strategies that have expected utility less than $u_p(\text{nolearn})$ can be removed.

For any player-two strategy τ , define $\Gamma(\tau)$ as the following strategy:

1. For the first use of the VTP, follow $\tau(\varepsilon)$. If the first use of the VTP results in all envelopes being opened, output the same as strategy τ .
2. If the first use of the VTP does not result in all envelopes being opened, for the subsequent uses of the VTP, follow strategy τ with the following exception: for any revealing history h , compute which of player 2's envelopes are non-check envelopes, choose a random value r and XOR r to each of these non-check values. Use these new non-check envelope values in place of the original non-check values received from the dealer to compute $\tau(h)$ for all subsequent histories h . If in this use or any subsequent use of the VTP, all envelopes are opened, compute the output O as per τ using the original non-check envelope values.

Claim. For any player-two strategy $\tau \in X_2$ such that $\tau \neq \Gamma(\tau)$, $\Gamma(\tau)$ weakly dominates τ .

For any player-one strategy σ , the player-two strategies τ and $\Gamma(\tau)$ are the same for the first use of the VTP, and thus result in similar utilities in any execution that succeeds.

For any revealing history, $\Gamma(\tau)$ never does worse than τ since $\Gamma(\tau)$ is both perfectly indistinguishable from τ to player one, and the share values produced by $\Gamma(\tau)$ do not have any information about the secret s' . Since $\Gamma(\tau) \neq \tau$, then there is some σ and some execution for which $\Gamma(\tau)$ will be strictly better than τ .

Set $X_2^1 = \{\Gamma(\tau)\}_{\tau \in X_2}$. For any player-one strategy σ , let $\Pi(\sigma)$ be the strategy that does the following: If the input history h is not revealing, then follow $\sigma(h)$. If input history h is revealing, then (a) never use the VTP in any subsequent round and (b) if $\sigma(h)$ outputs a string s , then output s and otherwise output \perp .

Claim. For any strategy σ for which (1) $\sigma(\varepsilon)$ submits the sealed envelope $E_{1,0}$, and (2) there exists $\tau' \in X_2^1$ such that (σ, τ') produces a revealing history h with positive probability and $\sigma(h) \neq \perp$, the strategy $\Pi(\sigma)$ weakly dominates σ .

Consider any profile (σ, τ) where $\tau \in X_2^1$. The strategies σ and $\Pi(\sigma)$ are equivalent on the first use of the VTP and therefore result in the same history h . If h is successful, then both σ and $\Pi(\sigma)$ result in reconstructing the secret. Similarly, if h is not successful and also not revealing, then the two strategies are equivalent. If h is revealing, but $\sigma(h)$ produces an output, then both are equivalent. Finally, if h is a revealing history and $\sigma(h)$ uses the VTP again, then $\Pi(\sigma)$ is strictly better. This follows because τ survives the first step of removal, and therefore τ produces envelopes for the second (and future) uses of the VTP that are independent of the secret s' . This upper-bounds player 1's utility $u_1(\sigma, \tau)$ by $-\epsilon + u_1(\text{nolearn})$. However, $u_1(\Pi(\sigma), \tau') = u_1(\text{nolearn})$ which is strictly greater. (Similar analysis for the case when σ outputs s instead of \perp .)

The second condition of the claim ensures this situation occurs for some τ , and therefore therefore $\Pi(\sigma)$ weakly dominates σ .

Set X_1^1 to be the set of player-one strategies in which after the sealed envelope is submitted, the VTP is never used again. Let $\Theta^i(\tau)$ be the player-two strategy that plays $T(\varepsilon)$ in the first i uses of the VTP, and follows τ for all subsequent uses.

Claim. If $\tau \neq \Theta^1(\tau)$, then $\Theta^1(\tau)$ weakly dominates τ .

Consider any player-one strategy $\sigma \in X_1^1$.

For those executions of σ in which player 1 submits an unsealed envelope in the first use of the VTP, all envelopes are immediately destroyed and therefore it holds that $u_2(\sigma, \Theta(\tau)) = u_2(\sigma, \tau)$ since both strategies are equivalent for all second and subsequent uses of the VTP.

We now consider those executions of σ in which the sealed $E_{1,0}$ is submitted. (This can only happen once.) Let $p_{\sigma, \tau}$ be the probability that under profile (σ, τ) , the first use of the VTP results in destroyed envelopes. Observe that $p_{\sigma, \tau} \geq p_{\sigma, \Theta(\tau)}$ for all σ . Since $\sigma \in X_1^1$, the VTP is never used again by σ , and therefore $u_2(\sigma, \tau) = p_{\sigma, \tau} u_2(\text{learn})$ which is less than or equal to $p_{\sigma, \Theta(\tau)} u_2(\text{learn}) = u_2(\sigma, \Theta(\tau))$. The condition that $\Theta(\tau) \neq \tau$ implies that the inequality is strict for some player one strategy σ which establishes the claim. Induction can be used to show that the claim holds for all i .

Set $X_2^2 = \{\Theta(\tau)\}_{\tau \in X_1^1}$.

Claim. The player-one strategy T weakly dominates every surviving strategy σ .

Observe that $u_1(T, \tau) = u_1(\text{both})$ for any $\tau \in X_2^2$. Any other player one strategy has a positive probability of causing the VTP to destroy all envelopes, and therefore incurring a cost of $-\epsilon$.

A similar argument with Π can be used to show that every player-two strategy, in any use of the VTP that reveals the dealer-received envelope E_2 , no longer uses the

VTP. This implies that the player-two strategy T weakly dominates every surviving strategy.

Acknowledgements

Many thanks to Sergei Izmalkov for his characteristically generous and insightful help.

References

- [ADGH06] I. Abraham, D. Dolev, R. Gonen, and J. Halpern. Distributed Computing Meets Game Theory: Robust Mechanisms for Rational Secret Sharing and Multiparty Computation. In *PODC '06*.
- [BP98] Elchanan Ben-Porath. Correlation without mediation: Expanding the set of equilibria outcomes by “cheap” pre-play procedures. *J. of Economic Theory*, 80:108–122, 1998.
- [CM08] Jing Chen and Silvio Micali. Resilient Mechanisms For Truly Combinatorial Auctions. MIT-CSAIL-TR-2008-067, November 2008.
- [GK06] S. Dov Gordon and Jonathan Katz. Rational secret sharing, revisited. In *SCN'06*.
- [GMW87] Oded Goldreich and Silvio Micali and Avi Wigderson. How to play any mental game. In *STOC'87*.
- [HT04] J. Halpern and V. Teague. Rational secret sharing and multiparty computation. In *STOC'04*, 2004.
- [HP08] J. Halpern and R. Pass. Game Theory with Costly Computation. Manuscript, 2008.
- [IML05] S. Izmalkov, S. Micali, and M. Lepinski. Rational Secure Computation and Ideal Mechanism Design. In *FOCS '05*.
- [ILM08] S. Izmalkov, M. Lepinski, and S. Micali. Verifiably secure devices. In *TCC'08*.
- [KN08a] G. Kol and M. Naor. Cryptography and Game Theory: Designing Protocols for Exchanging Information In *TCC'08*.
- [KN08b] G. Kol and M. Naor. Games for Exchanging Information. In *STOC'08*.
- [LMPS04] M. Lepinski, S. Micali, C. Peikert, and A. Shelat. Completely Fair SFE and Coalition-Safe Cheap Talk. In *PODC '04*.
- [LT06] Anna Lysyanskaya and Nikos Triandopoulos. Rationality and adversarial behavior in multi-party computation. In *CRYPTO'06*.
- [OPRV08] Shien Jin Ong and David Parkes and Alon Rosen and Salil Vadhan. Fairness with an Honest Minority and a Rational Majority. On Eprint, 2008/097.
- [OR] Martin J. Osborne and Ariel Rubinstein. A Course in Game Theory. MIT Press, 1994.

A The Ballot-Box Model

Ballot-box mechanisms are extensive-form, imperfect-information mechanisms with Nature. Accordingly, to specify them we must specify who acts when, the actions and the information available to the players, when the play terminates, and how the outcome is determined upon termination.

A ballot-box mechanism ultimately is a mathematical abstraction, but possesses a quite natural physical interpretation. The physical setting is that of a group of players, seated around a table, acting on a set of *ballots*. Within this physical setting, one has considerable latitude in choosing reasonable actions available to the players. In this paper, we make a specific choice, sufficient for our present goals.

A.1 Intuition

Ballots. Externally, all ballots of the same kind are identical. (Unlike [ILM08], we do not need super-envelopes here.) An envelope may contain a symbol from a finite alphabet. An envelope perfectly hides and guarantees the integrity of the symbol it contains until it is opened. Initially, all ballots are empty and in sufficient supply.

Ballot-Box Operations. We only need 3 classes of ballot-box operations. Each operation except for the first type is referred to as a *public action*, because it is performed in plain view, so that all players know exactly which action has been performed. These classes are: (1) writing a symbol on a piece of paper and sealing it into a new, empty envelope; (2) publicly opening an envelope to reveal its content to all players; (3) publicly destroying a ballot; and (4) do nothing.

Public Information. Conceptually, the players observe which actions have been performed on which ballots. Formally, (1) we associate to each ballot a unique identifier, a positive integer that is common information to all players (these identifiers correspond to the order in which the ballots are placed on the table for the first time or returned to the table —e.g., after being ballot-boxed); and (2) we have each action generate, when executed, a public string of the form “ A, j, k, l, \dots ”; where A is a string identifying the action and j, k, l, \dots are the identifiers of the ballots involved. The *public record* is the concatenation of the public strings generated by all actions executed thus far.

A.2 Formalization

Basic Notation. We denote by Σ the alphabet consisting of English letters, arabic numerals, and punctuation marks; by Σ^* the set of all finite strings over Σ ; by \mathbb{S}_k the group of permutations of k elements; by $x := y$ the operation that assigns value y to variable x ; by $p := \text{rand}(\mathbb{S}_k)$ the operation that assigns to variable p a randomly selected permutation in \mathbb{S}_k ; and by \emptyset the empty set.

If S is a set, by S^0 we denote the empty set, and by S^k the Cartesian product of S with itself k times. If x is a sequence, by either x^i or x_i we denote x 's i th element,⁸ and by $\{x\}$ the set $\{z : x^i = z \text{ for some } i\}$. If x and y are sequences, respectively of length j and k , by $x \circ y$ we denote their concatenation (i.e., the sequence of $j + k$ elements whose i th element is x^i if $i \leq j$, and y^{i-j} otherwise). If x and y are strings (i.e., sequences with elements in Σ), we denote their concatenation by xy .

If A is a probabilistic algorithm, the distribution over A 's outputs on input x is denoted by $A(x)$. A probabilistic function $f : X \rightarrow Y$ is *finite* if X and Y are both finite sets and, for every $x \in X$ and $y \in Y$, the probability that $f(x) = y$ has a finite binary representation.

⁸ For any given sequence, we shall solely use superscripts, or solely subscripts, to denote all of its elements.

Ballots and Actions. An *envelope* is a triple $(j, c, 0)$, where j is a positive integer, and c a symbol of Σ . A *ballot* is an envelope. If (j, c, L) is a ballot, we refer to j as its *identifier*, to c as its *content*, and to L as its *level*.

A set of ballots B is *well-defined* if distinct ballots have distinct identifiers. If B is a well-defined set of ballots, then I_B denotes the set of identifiers of B 's ballots. For $j \in I_B$, B_j (or the expression *ballot j*) denotes the unique ballot of B whose identifier is j . For $J \subset I_B$, B_J denotes the set of ballots of B whose identifiers belong to J .

Relative to a well-defined set of ballots B : if j is an envelope in B , then $\text{cont}_B(j)$ denotes the content of j ; if $x = j^1, \dots, j^k$ is a sequence of envelope identifiers in I_B , then $\text{cont}_B(x)$ denotes the concatenation of the contents of these envelopes, that is, the string $\text{cont}_B(j^1) \cdots \text{cont}_B(j^k)$.

A *global memory* consists of a pair (B, R) , where

- B is a well defined set of ballots; and
- R is a sequence of strings in Σ^* , $R = R^1, R^2, \dots$

We refer to B as the *ballot set*; to R as the *public record*; and to each element of R as a *record*. The *empty global memory* is the global memory for which the ballot set and the public record are empty. We denote the set of all possible global memories by GM .

Ballot-box actions are functions from GM to GM . The subset of ballot-box actions available at a given global memory gm is denoted by \mathcal{A}_{gm} . The actions in \mathcal{A}_{gm} are described below, grouped in 8 classes. For each $a \in \mathcal{A}_{gm}$ we provide a formal identifier; an informal reference (to facilitate the high-level description of our constructions); and a functional specification. If $gm = (B, R)$, we actually specify $a(gm)$ as a program acting on variables B and R . For convenience, we include in R the auxiliary variable ub , the *identifier upper-bound*: a value equal to 0 for an empty global memory, and always greater than or equal to any identifier in I_B .

1. (NEWEN, c) —where $c \in \Sigma$.
“Make a new envelope with content c .”
 $\text{ub} := \text{ub} + 1$; $B := B \cup \{(\text{ub}, c, 0)\}$; and $R := R \circ (\text{NEWEN}, c, \text{ub})$.
2. (OPENEN, j) —where j is an envelope identifier in I_B .
“Publicly open envelope j to reveal content $\text{cont}_B(j)$.”
 $B := B \setminus \{B_j\}$ and $R := R \circ (\text{OPENEN}, j, \text{cont}_B(j), \text{ub})$.
3. (DESTROY, j) —where j is a ballot identifier in I_B .
“Destroy ballot j ”
 $B := B \setminus \{B_j\}$ and $R := R \circ (\text{DESTROY}, j, \text{ub})$.
4. (DONOTHING).
“Do nothing”
 $B := B$ and $R := R \circ (\text{DONOTHING}, \text{ub})$.

Remarks.

- All ballot-box actions are deterministic functions.
- The variable ub never decreases and coincides with the maximum of all identifiers “ever in existence.” Notice that we never re-use the identifier of a ballot that has left, temporarily or for ever, the table. This ensures that different ballots get different identifiers.

Definition 2. A global memory gm is feasible if there exists a sequence of global memories gm^0, gm^1, \dots, gm^k , such that gm^0 is the empty global memory; $gm^k = gm$; and, for all $i \in [1, k]$, $gm^i = a^i(gm^{i-1})$ for some $a^i \in \mathcal{A}_{gm^{i-1}}$.

If (B, R) is a feasible memory, we refer to R as a feasible public record.

Notice that if $gm = (B, R)$ is feasible, then \mathcal{A}_{gm} is easily computable from R alone. Indeed, what ballots are in play, which ballots are envelopes and which are super-envelopes, *et cetera*, are all deducible from R . Therefore, different feasible global memories that have the same public record also have the same set of available actions. This motivates the following definition.

Definition 3. If R is a feasible public record, by \mathcal{A}_R we denote the set of available actions for any feasible global memory with public record R .

B The Notion of a Public Ballot-Box Mediator (VTP in our language)

Definition 4. Let \mathcal{P} be a sequence of K functions. We say that \mathcal{P} is a public ballot-box mediator (of length K) if, for all $k \in [1, K]$ and public records R , $P^k(R)$ is a public ballot-box action in \mathcal{A}_R .

An execution of \mathcal{P} on an initial feasible global memory (B^0, R^0) is a sequence of global memories $(B^0, R^0), \dots, (B^K, R^K)$ such that $(B^k, R^k) = a^k(B^{k-1}, R^{k-1})$ for all $k \in [1, K]$, where $a^k = P^k(R^{k-1})$.⁹

If e is an execution of \mathcal{P} , by $B^k(e)$ and $R^k(e)$ we denote, respectively, the ballot set, the public record, and the private history profile of e at round k . By $R_{\mathcal{P}}^k(e)$ we denote the last k records of $R^k(e)$ (i.e., “the records appended to R^0 by executing \mathcal{P} ”).

Remarks.

- Note that the above definition captures our intuitive desideratum that no special trust is bestowed on a public mediator. Because he performs a sequence of public ballot-box actions, any one can verify that
 - (i) he performs the right sequence of actions;
 - (ii) he does not choose these actions; and
 - (iii) he does not learn any information that is not publicly available.
- Note too that if $\mathcal{P} = P^1, \dots, P^K$ and $\mathcal{Q} = Q^1, \dots, Q^L$ are public mediators, then their concatenation, that is, $P^1, \dots, P^K, Q^1, \dots, Q^L$ is a public mediator too.

⁹ Note that the executions of \mathcal{P} are, in general, random since $P^k(R)$ may return an action of Nature.