Chapter 4

The Composition Theorem

4.1 Composition of Commitment Schemes

4.1.1 Introduction

In this chapter, we present one of the central results of this thesis: the composition theorems. We define a *composition operation* in Definition 4.3 which constructs a *composed scheme* out of two *component schemes*. This is achieved as follows: The provers use the first scheme to commit, but instead of opening the commitment, they instead commit to the opening information using the second scheme. They then open this second commitment, which allows the verifier to obtain the opening information of the first scheme and thus open the first commitment.

The composition operation cannot be applied to any two schemes – we define a notion of *eligible pairs* (Definition 4.1) that it can be applied to. Some of the requirements in that definition are necessary for the composition operation to make sense. For example, the composition operation requires that the first scheme is structured so that in the opening phase, a prover sends some *opening information* to the verifier who then computes the result. Other requirements are necessary for our proofs of the composition theorems to work.

The composition theorems show that if the first scheme is ε-binding (according to some binding definition) and the second one is δ-binding (according to the same definition), then the composed scheme is $(\epsilon + \delta)$-binding.

While this is what one would expect from this composition operation, it is non-trivial to formally prove. It is unclear how one would prove this result using the sum-binding definition directly, rather than our newly-introduced definitions.

Together with our analysis of the CHSH scheme in Section 3.2.6, the composition theorem implies that the binding parameter of the Lunghi *et al.* scheme decays linearly in the number m of rounds, rather than double-exponentially,

as [LKB+15] suggests.

4.1.2 Eligible Pairs and the Composition Operation

We consider two 2-prover commitment schemes $\mathcal{S}$ and $\mathcal{S}'$ of a restricted form, and we compose them to a new 2-prover commitment scheme $\mathcal{S}'' = \mathcal{S} \star \mathcal{S}'$ in a well-defined way; our composition theorem then shows that $\mathcal{S}''$ is secure (against classical attacks) if $\mathcal{S}$ and $\mathcal{S}'$ are. We start by specifying the restriction to $\mathcal{S}$ and $\mathcal{S}'$ that we impose.

Definition 4.1. *Let $\mathcal{S}$ and $\mathcal{S}'$ be two 2-prover string-commitment schemes with domains D and D', respectively. We call the pair $(\mathcal{S}, \mathcal{S}')$* eligible *if the following two properties hold, or they hold with the roles of P and Q exchanged.*

1. *The commit phase of $\mathcal{S}$ is a protocol com_{PQV} that involves communication between P and V only, and the opening phase of $\mathcal{S}$ is a protocol $\mathsf{open}_{QV} = (\mathsf{open}_Q, \mathsf{open}_V)$ between Q and V only. In other words, com_Q and open_P are both trivial and do nothing.*[1] *Similarly, the commit phase of $\mathcal{S}'$ is a protocol com'_{PQV} that involves communication between Q and V only (but both provers may be active in the opening phase).*

2. *The opening phase open_{QV} of $\mathcal{S}$ is of the following simple form: Q sends some $y \in D'$ to V, and V computes s deterministically as $s = \mathsf{Extr}(y, c)$, where c is the commitment.*[2] *We call this message y the* opening information*.*

Furthermore, we specify that the possible attacks on $\mathcal{S}$ are so that P and Q do not communicate during the course of the entire scheme, and the possible attacks on $\mathcal{S}'$ are so that P and Q do not communicate during the course of the commit phase but there may be limited communication during the opening phase.

An example of an eligible pair of 2-prover commitments is $(\mathcal{CHSH}^q, \mathcal{XCHSH}^q)$, where $\mathcal{XCHSH}^q$ coincides with scheme $\mathcal{CHSH}^q$ except that the roles of P and Q are exchanged.

Remark 4.2. *For an eligible pair $(\mathcal{S}, \mathcal{S}')$, it will be convenient to understand open_Q and open_V as* non-interactive *algorithms, where open_Q produces y as its* output*, and open_V takes y as additional* input *(rather than viewing the pair as a protocol with a single one-way communication round).*

We now define the composition operation. Informally, committing is done by means of committing using $\mathcal{S}$, and to open the commitment, Q uses open_Q

[1] Except that com_Q may output state information to the opening protocol open_Q, e.g., in order to pass on the commit phase randomness.

[2] Our composition theorem also works for a randomized Extr, but for simplicity, we restrict to the deterministic case.

to locally compute the opening information y and he commits to y with respect to the scheme $\mathcal{S}'$, and then this commitment is opened (to y), and V computes and outputs $s = \mathsf{Extr}(y, c)$. Formally, this is captured as follows (see also Fig. 4.1).

Definition 4.3. *Let $\mathcal{S} = (\mathsf{com}_{PV}, \mathsf{open}_{QV})$ and $\mathcal{S}' = (\mathsf{com}'_{QV}, \mathsf{open}'_{PQV})$ be an eligible pair of 2-prover commitment schemes. Then, their composition $\mathcal{S} \star \mathcal{S}'$ is defined as the 2-prover commitment scheme consisting of $\mathsf{com}_{PV} = (\mathsf{com}_P[\xi_{PQ}], \mathsf{com}_V)$ and*

$$\mathsf{open}''_{PQV} = (\mathsf{open}'_P,\ \mathsf{open}'_Q \circ \mathsf{com}'_Q \circ \mathsf{open}_Q,\ \mathsf{open}_V \circ \mathsf{open}'_V \circ \mathsf{com}'_V)\,,$$

where we make it explicit that com_P and open_Q use joint randomness, and so do com'_Q and open'_P.

When considering attacks against the binding property of the composed scheme $\mathcal{S} \star \mathcal{S}'$, we declare that the possible deterministic attacks[3] are those of the form $(\overline{\mathsf{com}}_P, \overline{\mathsf{open}}'_{PQ} \circ \mathsf{ptoq}_{PQ} \circ \overline{\mathsf{com}}'_Q)$, where $\overline{\mathsf{com}}_P$ is an possible deterministic commit strategy for $\mathcal{S}$, $\overline{\mathsf{com}}'_Q$ and $\overline{\mathsf{open}}'_{PQ}$ are possible deterministic commit and opening strategies for $\mathcal{S}'$, and ptoq_{PQ} is the one-way communication protocol that communicates P's input to Q (see also Fig. 4.2).[4]

Remark 4.4. *We point out that the composition $\mathcal{S} \star \mathcal{S}'$ can be naturally defined for a* larger *class of pairs of schemes (e.g. where* both *provers are active in the commit phase of both schemes), and the above intuition still holds. However, our proof only works for this restricted class of pairs of schemes. Extending the composition result in that direction is an open problem.*

Remark 4.5. *We observe that if $(\mathcal{S}, \mathcal{XS})$ is an eligible pair, where $\mathcal{XS}$ coincides with $\mathcal{S}$ except that the roles of P and Q are exchanged, then so is $(\mathcal{XS}, \mathcal{S} \star \mathcal{XS})$. As such, we can then compose $\mathcal{XS}$ with $\mathcal{S} \star \mathcal{XS}$, and obtain yet another eligible pair $(\mathcal{S}, \mathcal{XS} \star \mathcal{S} \star \mathcal{XS})$, etc. We write $\mathcal{S}_m$ for the m-fold composition of $\mathcal{S}$ with itself, i.e., $\mathcal{S}_{\Updownarrow} = \mathcal{S} \star \mathcal{XS} \star \mathcal{S} \star \ldots$ for m terms. Applying this to the schemes $\mathcal{S} = \mathcal{CHSH}^q$, we obtain the multi-round scheme from Lunghi et al. [LKB+15]. As such, our composition theorem below implies security of their scheme — with a* linear *blow-up of the error term (instead of double exponential).*

We point out that formally we obtain security of the Lunghi et al. scheme as a 2-prover commitment scheme *under an* abstract restriction *on the provers' communication: in every round, the active prover cannot access the message that the other prover received in the previous round. As such, when the rounds of the protocol are executed fast enough so that it is ensured that there is no*

[3] The possible *randomized* attacks are then naturally given as those that pick one of the deterministic attacks according to some distribution.

[4] This one-way communication models that in the relativistic setting, sufficient time has passed at this point for P to inform Q about what happened during com_P.

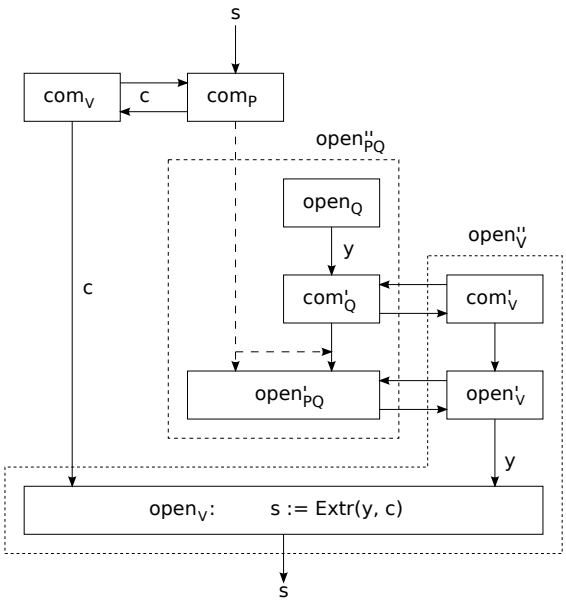

Figure 4.1: The composition of $\mathcal{S}$ and $\mathcal{S}'$ (assuming single-round commit phases). The dotted arrows indicate communication possible to the dishonest provers.

time for the provers to communicate between subsequent rounds, then security as a relativistic commitment scheme *follows immediately.*

Remark 4.6. *In Section 2.2.2, we noted that the verifier in a relativistic commitment scheme also needs to be convinced that the provers keep the appropriate distance from each other. In some schemes, this can be achieved by splitting up the verifier into two entities and placing them at the same locations as the provers. If $(\mathcal{S}, \mathcal{XS})$ is an eligible pair, then $\mathcal{S} \star \mathcal{XS}$, and more generally $\mathcal{S}_m$, allow for the verifier to be split up in this way. This holds because in each round before the last, the verifier initiates a new commitment.*

Remark 4.7. *It is immediate that $\mathcal{S} \star \mathcal{S}'$ is a commitment scheme in the sense of Definition 2.8, and that it is complete if $\mathcal{S}$ and $\mathcal{S}'$ are, with the error parameters adding up. It is intuitively clear that $\mathcal{S} \star \mathcal{S}'$* should be *binding if $\mathcal{S}$ and $\mathcal{S}'$ are: committing to the opening information y and then opening the commitment allows the provers to* delay *the announcement of y (which is the whole point of the exercise), but it does not allow them to* change *y, by the binding property of $\mathcal{S}'$; thus, $\mathcal{S} \star \mathcal{S}'$ should be (almost) as binding as $\mathcal{S}$. This intuition is confirmed by our composition theorem below.*

4.1.3 Hiding Property for Composed Schemes

The hiding property is obviously inherited from $\mathcal{S}$, i.e., $\mathcal{S} \star \mathcal{S}'$ is δ-binding in the sense of Definition 2.13 if and only if $\mathcal{S}$ is δ-hiding. However, this definition is not suitable for schemes with a multi-round opening phase: a scheme that reveals the committed string s in the first round of the opening phase would still satisfy Definition 2.13, but clearly, doing so defeats the entire purpose of a multi-round commitment scheme. Recall that, using the terminology used in context of relativistic commitments, the rounds of the opening phase up to before the last are referred to as the *sustain phase*, and only the last round is considered the opening phase proper. We show in this section that $\mathcal{S} \star \mathcal{S}'$ is hiding up to before the last round, with the error parameters adding up.

Definition 4.8. *Let $\mathcal{S} = (\mathsf{com}_{PQV}, \mathsf{open}_{PQV})$ be a commitment scheme. We write v for the verifier's view immediately before the last round of communication in open_{PQV}. We say that a scheme is* ε-hiding until the last round *if for any (possibly dishonest) verifier V and any two inputs s_0 and s_1 to the honest provers, we have $d(p(v|s_0), p(v|s_1)) \leq \varepsilon$.*

Theorem 4.9. *Let $\mathcal{S}$ be a ε-hiding commitment scheme and $\mathcal{S}'$ a scheme that is δ-hiding until the last round. If $(\mathcal{S}, \mathcal{S}')$ is eligible, then the composed scheme $\mathcal{S}'' = \mathcal{S} \star \mathcal{S}'$ is $(\varepsilon + \delta)$-hiding until the last round.*

Proof. Fix a strategy against the hiding-until-the-last-round property of $\mathcal{S}''$. We consider the distribution $p(v, y, v'|s)$ where s is the string that the provers commit to, v the verifier's view after $\overline{\mathsf{com}}_{PQV}$ has been executed, y the opening

information to which Q commits using the scheme $\mathcal{S}'$, and v' the verifier's view immediately before the last round of communication. We need to show that $d(p(v'|s_0), p(v'|s_1)) \leq \varepsilon + \delta$ for any s_0 and s_1.

First, note that $p(v'|v, y, s_b) = p(v'|v, y)$ since v' is produced by P, Q and V acting on y and v only. From any strategy against $\mathcal{S}''$, we can obtain a strategy against $\mathcal{S}'$ by fixing v. Thus, by the hiding property of $\mathcal{S}'$, for any y_0 and y_1, we have $d(p(v'|v, y = y_0), p(v'|v, y = y_1)) \leq \delta$ and it follows by the convexity of the statistical distance in both arguments that

$$p(v'|v, s_0) = \sum_y p(y|v, s_0)p(v'|v, y) \approx_\delta \sum_y p(y|v, s_1)p(v'|v, y) = p(v'|v, s_1)$$

where we use $\approx_\delta$ to indicate that the two distributions have statistical distance at most δ. Since we have $d(p(v|s_0), p(v|s_1)) \leq \varepsilon$ by the hiding property of $\mathcal{S}$, it follows that

$$\begin{aligned} p(v'|s_0) = p(v, v'|s_0) = p(v|s_0)p(v'|v, s_0) &\approx_\delta p(v|s_0)p(v'|v, s_1) \\ &\approx_\varepsilon p(v|s_1)p(v'|v, s_1) \\ &= p(v, v'|s_1) \\ &= p(v'|s_1) \end{aligned}$$

where the first and last equality hold because v' contains v since v' is the view of V at a later point in time. □

4.2 The Composition Theorems

4.2.1 A Composition Theorem for (Strongly) Binding Schemes

Before stating and proving the composition theorem, we need to single out one more relevant parameter.

Definition 4.10. *Let $(\mathcal{S}, \mathcal{S}')$ be an eligible pair, which in particular means that V's action in the opening phase of $\mathcal{S}$ is determined by a function* Extr. *We define $k(\mathcal{S}) := \max_{c,s} |\{y \,|\, \mathsf{Extr}(y, c) = s\}|$.*

I.e., $k(\mathcal{S})$ counts the number of y's that are consistent with a given string s in the worst case. Note that $k(\mathcal{CHSH}^q) = 1$: for every $a, x, s \in \mathbb{F}_q$ there is at most one $y \in \mathbb{F}_q$ such that $x + y = a \cdot s$.

In the following composition theorems, we take it as understood that the assumed respective binding properties of $\mathcal{S}$ and $\mathcal{S}'$ hold with respect to a well-defined respective classes of possible attacks. We start with the composition theorem for the fairly-binding property, which is easier to prove than the one for the fairly-weak-binding property.

Theorem 4.11. *Let $(\mathcal{S}, \mathcal{S}')$ be an eligible pair of 2-prover commitment schemes, and assume that $\mathcal{S}$ and $\mathcal{S}'$ are respectively ε-fairly-binding and δ-fairly-binding. Then, their composition $\mathcal{S}'' = \mathcal{S} \star \mathcal{S}'$ is $(\varepsilon + k(\mathcal{S}) \cdot \delta)$-fairly-binding.*

Proof. We first consider the case $k(\mathcal{S}) = 1$. We fix an attack $(\overline{\mathsf{com}}_P, \overline{\mathsf{open}}''_{PQ})$ against $\mathcal{S}''$. Without loss of generality, the attack is deterministic, so $\overline{\mathsf{open}}''_{PQ}$ is of the form $\overline{\mathsf{open}}''_{PQ} = \overline{\mathsf{open}}'_{PQ} \circ \mathsf{ptoq}_{PQ} \circ \overline{\mathsf{com}}'_Q$.

Note that $\overline{\mathsf{com}}_P$ is also a commit strategy for $\mathcal{S}$. As such, by the fairly-binding property of $\mathcal{S}$, there exists a function $\hat{s}(c)$, only depending on $\overline{\mathsf{com}}_P$, so that the property specified in Definition 3.2 is satisfied for every opening strategy $\overline{\mathsf{open}}_Q$ for $\mathcal{S}$. We will show that it is also satisfied for the (arbitrary) opening strategy $\overline{\mathsf{open}}''_{PQ}$ for $\mathcal{S}''$, except for a small increase in ε: we will show that $p(\hat{s}(c) \neq s \wedge s = s_\circ) \leq \varepsilon + \delta$ for every fixed target string $s_\circ$. This then proves the claim.

To show this property on $\hat{s}(c)$, we "decompose and reassemble" the attack strategy $(\overline{\mathsf{com}}_P, \overline{\mathsf{open}}'_{PQ} \circ \mathsf{ptoq}_{PQ} \circ \overline{\mathsf{com}}'_Q)$ for $\mathcal{S}''$ into an attack strategy $(\overline{\mathsf{com}}'_Q, \overline{\mathsf{newopen}}'_{PQ})$ for $\mathcal{S}'$ with $\overline{\mathsf{newopen}}'_{PQ}$ formally defined as

$$\overline{\mathsf{newopen}}'_{PQ}[c](\overline{state}'_Q) := \overline{\mathsf{open}}'_{PQ}\big(\overline{state}_P(c) \| (\overline{state}_P(c), \overline{state}'_Q)\big)$$

where

$$(\overline{state}_P(c) \| c) \leftarrow \big(\overline{\mathsf{com}}_P || \mathsf{com}_V\big)\,.$$

Informally, this means that ahead of time, P and Q *simulate* an execution of $(\overline{\mathsf{com}}_P(\emptyset) || \mathsf{com}_V(\emptyset))$ and take the resulting communication/commitment[5] c as shared randomness, and then $\overline{\mathsf{newopen}}'_{PQ}$ computes $\overline{state}_P$ from c as does $\overline{\mathsf{com}}_P$, and runs $\overline{\mathsf{open}}'_{PQ}$ (see Fig. 4.2).[6] It follows from the fairly-binding property that there is a function $\hat{y}(c')$ of the commitment c' so that $p(\hat{y}(c') \neq y \wedge y = y_\circ(c)) \leq \delta$ for every function $y_\circ(c)$.

The existence of $\hat{y}$ now gives rise to an opening strategy $\overline{\mathsf{open}}_Q$ for $\mathcal{S}$; namely, simulate the commit phase of $\mathcal{S}'$ to obtain the commitment c', and output $\hat{y}(c')$. By Definition 3.2, for $\tilde{s} := \mathsf{Extr}(\hat{y}(c'), c)$ and every $s_\circ$, $p(\hat{s}(c) \neq \tilde{s} \wedge \tilde{s} = s_\circ) \leq \varepsilon$.

We are now ready to put things together. Fix an arbitrary target string $s_\circ$. For any c we let $y_\circ(c)$ be the unique string such that $\mathsf{Extr}(y_\circ(c), c) = s_\circ$ (and some default string if no such string exists); recall, we assume for the moment that $k(\mathcal{S}) = 1$. Omitting the arguments in $\hat{s}(c), \hat{y}(c')$ and $y_\circ(c)$, it follows that

$$\begin{aligned}
p(\hat{s} \neq s \wedge s = s_\circ) &\leq p(\hat{s} \neq s \wedge s = s_\circ \wedge s = \tilde{s}) + p(s = s_\circ \wedge s \neq \tilde{s}) \\
&\leq p(\hat{s} \neq \tilde{s} \wedge \tilde{s} = s_\circ) + p\big(\mathsf{Extr}(y, c) \neq \mathsf{Extr}(\hat{y}, c) \wedge \mathsf{Extr}(y, c) = s_\circ\big) \\
&\leq p(\hat{s} \neq \tilde{s} \wedge \tilde{s} = s_\circ) + p(y \neq \hat{y} \wedge y = y_\circ) \\
&\leq \varepsilon + \delta.
\end{aligned}$$

[5] Recall that by convention (Remark 2.10), the commitment c equals the communication between V and, here, P.

[6] We are using here that Q is inactive during $\overline{\mathsf{com}}_{PQ}$ and P during $\overline{\mathsf{com}}'_{PQ}$, and thus the two "commute".

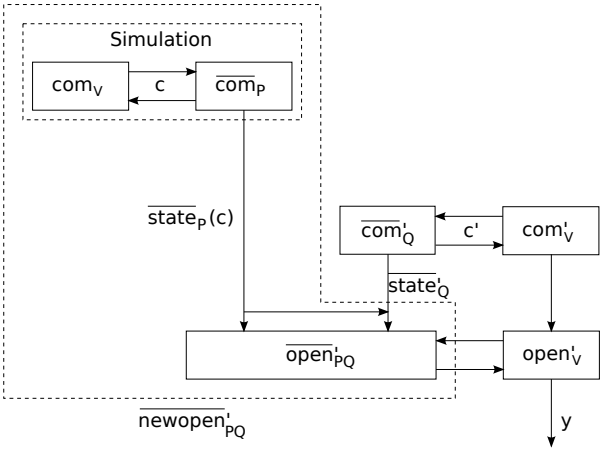

Figure 4.2: Constructing the opening strategy $\overline{\mathsf{newopen}}'_{PQ}$ against $\mathcal{S}'$.

Thus, $\hat{s}$ is as required.

For the general case where $k(\mathcal{S}) > 1$, we can reason similarly, except that we then list the $k \leq k(\mathcal{S})$ possibilities $y_\circ^1(c), \ldots, y_\circ^k(c)$ for $y_\circ(c)$, and conclude that $p(s \neq \tilde{s} \wedge s = s_\circ) \leq \sum_i p\big(y \neq \hat{y} \wedge y = y_\circ^i\big) \leq k(\mathcal{S}) \cdot \delta$, which then results in the claimed bound. □

Remark 4.12. *Putting things together, we can now conclude the security (i.e., the binding property) of the Lunghi et al. multi-round commitment scheme. Corollary 3.22 ensures the fairly-binding property of $\mathcal{CHSH}^q$, i.e., the Crépeau et al. scheme as a string-commitment scheme, with parameter $2\sqrt{q^{-1}}$. The composition theorem (Theorem 4.11) then guarantees the fairly-binding property of the m-fold composition as a string-commitment scheme, with parameter $(m+1) \cdot 2\sqrt{q^{-1}}$. Finally, Proposition 3.6 implies that the m-fold composition of $\mathcal{CHSH}^q$ with itself is a ε_m-binding bit-commitment scheme with error parameter $\varepsilon_m = (m+1) \cdot 4\sqrt{q^{-1}}$ as claimed in the introduction, or, more generally, and by taking Remark 3.7 into account, a $(m+1) \cdot 2^{k+1}\sqrt{q^{-1}}$-binding k-bit-string-commitment scheme.*

4.2.2 Composition Theorem for Weakly Binding Schemes

We now show the composition theorem for the weak version of the binding property. Since this notion makes sense also against quantum attacks, we emphasize the restriction to classical attacks — extending the theorem to quantum attacks is an open problem. See Chapter 5 for some partial progress in this direction.

Theorem 4.13. *Let $(\mathcal{S}, \mathcal{S}')$ be an eligible pair of 2-prover commitment schemes, and assume that $\mathcal{S}$ and $\mathcal{S}'$ are respectively ε-fairly-weak-binding and δ-fairly-weak-binding against classical attacks. Then, their composition $\mathcal{S}'' = \mathcal{S} \star \mathcal{S}'$ is a $(\varepsilon + k(\mathcal{S}) \cdot \delta)$-fairly-weak-binding 2-prover commitment scheme against classical attacks.*

Proof. We first consider the case $k(\mathcal{S}) = 1$. We fix an arbitrary deterministic attack $(\overline{\mathsf{com}}_P, \overline{\mathsf{open}}''_{PQ})$ against $\mathcal{S}''$, where $\overline{\mathsf{open}}''_{PQ}$ is of the form $\overline{\mathsf{open}}''_{PQ} = \overline{\mathsf{open}}'_{PQ} \circ \mathsf{ptoq}_{PQ} \circ \overline{\mathsf{com}}'_Q$. Let a be V's randomness in com_V. Then, c is a function $c(a)$ of a, and the distribution $p(a, y)$ is well defined. Since $\overline{\mathsf{com}}_P$ is also an attack strategy against $\mathcal{S}$, there exists a distribution $p(\hat{s})$ (only depending on $\overline{\mathsf{com}}_P$) such that Definition 3.9 is satisfied for every opening strategy $\overline{\mathsf{open}}_Q$ for $\mathcal{S}$.

Similar to the proof of Theorem 4.11, we decompose and reassemble the attack strategy $(\overline{\mathsf{com}}_P, \overline{\mathsf{open}}'_{PQ} \circ \mathsf{ptoq}_{PQ} \circ \overline{\mathsf{com}}'_Q)$ for $\mathcal{S}''$ into an attack strategy $(\overline{\mathsf{com}}'_Q, \overline{\mathsf{newopen}}'_{PQ})$ for $\mathcal{S}'$. Concretely, for every fixed choice of a, we obtain a deterministic opening strategy $\overline{\mathsf{newopen}}'_{PQ,a}$ given by

$$\overline{\mathsf{newopen}}'_{PQ,a}(\overline{state}'_Q) := \overline{\mathsf{open}}'_{PQ}\big(\overline{state}_P(c(a)) \| (\overline{state}_P(c(a)), \overline{state}'_Q)\big)\,,$$

and the distribution of the verifier's output y when the provers use $\overline{\mathsf{newopen}}'_{PQ,a}$ is $p(y|a)$. It follows from the fairly-weak-binding property of $\mathcal{S}'$ that there exists a distribution $p(\hat{y})$, only depending on $\overline{\mathsf{com}}'_Q$, so that for every choice of a there exists a consistent joint distribution $p(\hat{y}, y|a)$ so that $p(\hat{y} \neq y \wedge y = y_\circ|a) \leq \delta$ for every fixed target string $y_\circ$. Note that here, consistency in particular means that $p(\hat{y}|a) = p(\hat{y})$. This joint conditional distribution $p(\hat{y}, y|a)$ together with the distribution $p(a)$ of a then naturally defines the distribution $p(a, \hat{y}, y)$, which is consistent with $p(a, y)$ considered above.

The existence of $p(\hat{y})$ now gives rise to an opening strategy $\overline{\mathsf{open}}_Q$ for $\mathcal{S}$; namely, sample $\hat{y}$ according to $p(\hat{y})$ and output $\hat{y}$. Note that the joint distribution of a and $\hat{y}$ in this "experiment" is given by

$$p(a) \cdot p(\hat{y}) = p(a) \cdot p(\hat{y}|a) = p(a, \hat{y})\,,$$

i.e., is consistent with the distribution $p(a, \hat{y}, y)$ above. By Definition 3.9, we know there exists a joint distribution $p(\hat{s}, \tilde{s})$, consistent with $p(\hat{s})$ fixed above and with $p(\tilde{s})$ determined by $\tilde{s} := \mathsf{Extr}(\hat{y}, c(a))$, and such that $p(\hat{s} \neq \tilde{s} \wedge \tilde{s} = s_\circ) \leq \varepsilon$ for every $s_\circ$. We can now "glue together" $p(\hat{s}, \tilde{s})$ and $p(c, \hat{y}, y, \tilde{s})$, i.e., find a joint distribution that is consistent with both, by setting

$$p(a, \hat{y}, y, \tilde{s}, \hat{s}) := p(a, \hat{y}, y, \tilde{s}) \cdot p(\hat{s}|\tilde{s})\,.$$

We now fix an arbitrary target string $s_\circ$. Furthermore, for any a we let $y_\circ(a)$ be the unique string such that $\mathsf{Extr}(y_\circ(a), c(a)) = s_\circ$ (and to some default string if no such string exists); recall, we assume for the moment that $k(\mathcal{S}) = 1$. With

respect to the above joint distribution, it then holds that

$$\begin{aligned} p(\hat{s} \neq s \wedge s = s_\circ) &= p(\hat{s} \neq s \wedge s = s_\circ \wedge s = \tilde{s}) + p(s = s_\circ \wedge s = s_\circ \wedge s \neq \tilde{s}) \\ &\leq p(\hat{s} \neq \tilde{s} \wedge s = s_\circ \wedge s = \tilde{s}) + p(s \neq \tilde{s} \wedge s = s_\circ) \\ &\leq p(\hat{s} \neq \tilde{s} \wedge \tilde{s} = s_\circ) \\ &\quad + p\big(\mathsf{Extr}(y, c(a)) \neq \mathsf{Extr}(\hat{y}, c(a)) \wedge \mathsf{Extr}(y, c(a)) = s_\circ\big) \\ &\leq p(\hat{s} \neq \tilde{s} \wedge \tilde{s} = s_\circ) + p(y \neq \hat{y} \wedge y = y_\circ(a)) \\ &\leq p(\hat{s} \neq \tilde{s} \wedge \tilde{s} = s_\circ) + \textstyle\sum_a p(a) \cdot p(y \neq \hat{y} \wedge y = y_\circ(a) | a) \\ &\leq \varepsilon + \delta \,. \end{aligned}$$

Thus, the distribution $p(\hat{s}, s)$ is as required.

For the case where $k(\mathcal{S}) > 1$, we can reason similarly, except that we then list the $k \leq k(\mathcal{S})$ possibilities $y_\circ^1(a), \ldots, y_\circ^a(a)$ for $y_\circ(a)$, and conclude that $p(s \neq \tilde{s} \wedge s = s_\circ) \leq \sum_i p\big(y \neq \hat{y} \wedge y = y_\circ^i(a)\big) \leq k(\mathcal{S}) \cdot \delta$, which then results in the claimed bound. □

Remark 4.14. *Analogously to Remark 4.12, we can conclude from Corollary 3.23 and Theorem 4.13 that $\mathcal{CHSH}^q$ is $(m{+}1)\cdot\sqrt{2q^{-1}}$-fairly-weak-binding. It follows from Proposition 3.11 that $\mathcal{CHSH}^q$ is a $(m+1) \cdot 2^{3/2}\sqrt{q^{-1}}$-weak-binding bit-commitment scheme. More generally, we can conclude that for any $k < n$, it is a $(m+1) \cdot 2^{k+1/2}\sqrt{q^{-1}}$-weak-binding k-bit string-commitment scheme. Below, we show how to avoid the factor 2 introduced by invoking Proposition 3.11.*

4.3 Variations

In this section, we show two variants of the composition theorems. The first one says that if we compose a weak-binding with a fairly-weak-binding scheme, we obtain a weak-binding scheme. This allows us to slightly improve the parameter in Remark 4.14. The proof crucially relies on the fact that, in the weak definition, there is some freedom in "gluing together" the distributions $p(s)$ and $p(\hat{s})$. The second variant says that composing two binding (or weak-binding) schemes yields a binding (or weak-binding, respectively) scheme.

We start by proving the following two properties for fairly-weak-binding commitment schemes. The first property shows that one may assume the joint distribution $p(\hat{s}, s)$ to be such that s and $\hat{s}$ are independent conditioned on $s \neq \hat{s}$.

Lemma 4.15. *Let $\mathcal{S}$ be a ε-fairly-weak-binding commitment scheme. Then, for any $\overline{\mathsf{com}}_{PQ}$ and $\overline{\mathsf{open}}_{PQ}$ there exists a joint distribution $p(\hat{s}, s)$ as required by Definition 3.9, but with the additional property that*

$$p(\hat{s}, s | s \neq \hat{s}) = p(\hat{s} | s \neq \hat{s}) \cdot p(s | s \neq \hat{s}) \,.$$

Proof. Since the scheme is ε-fairly-weak-binding, it follows that there exists a consistent joint distribution $p(\hat{s}, s)$ such that $p(s \neq \hat{s} \wedge s = s_\circ) \leq \varepsilon$ for every $s_\circ$. Because of this, we have

$$\begin{aligned} p(s = s_\circ) &= p(s{=}s_\circ \wedge \hat{s}{=}s_\circ) + p(s{=}s_\circ \wedge \hat{s}{\neq}s_\circ) \\ &= p(s{=}s_\circ \wedge \hat{s}{=}s_\circ) + p(s{\neq}\hat{s} \wedge s{=}s_\circ) \\ &\leq p(\hat{s} = s_\circ) + \varepsilon. \end{aligned}$$

We apply Lemma 2.2 to the marginal distributions $p(\hat{s})$ and $p(s)$. The resulting joint distribution $\tilde{p}(\hat{s}, s)$ satisfies $\tilde{p}(\hat{s} = s_\circ \wedge s = s_\circ | s = \hat{s}) = \min\{p(s = s_\circ), p(\hat{s} = s_\circ)\}$ and $\tilde{p}(\hat{s}, s | s \neq \hat{s}) = \tilde{p}(\hat{s} | s \neq \hat{s}) \cdot \tilde{p}(s | s \neq \hat{s})$. It remains to show that $\tilde{p}(s \neq \hat{s} \wedge s = s_\circ) \leq \varepsilon$ for all $s_\circ$. Indeed, we have

$$\begin{aligned} \tilde{p}(s \neq \hat{s} \wedge s = s_\circ) &= \tilde{p}(s = s_\circ) - \tilde{p}(s = \hat{s} \wedge s = s_\circ) \\ &= \tilde{p}(s = s_\circ) - \tilde{p}(\hat{s} = s_\circ \wedge s = s_\circ) \\ &= p(s = s_\circ) - \min\{p(\hat{s} = s_\circ), p(s = s_\circ)\} \\ &\leq p(s = s_\circ) - (p(s = s_\circ) - \varepsilon) \\ &= \varepsilon \end{aligned}$$

as claimed. □

The second property shows that the quantification over all *fixed* $s_\circ$ in Definition 3.9 of the fairly-weak-binding property can be relaxed to $s_\circ$ that may depend on $\hat{s}$, but only on $\hat{s}$. Note that we can obviously not allow $s_\circ$ to depend (arbitrarily) on s, since then one could choose $s_\circ = s$.

Proposition 4.16. *Let $\mathcal{S}$ be a ε-fairly-weak-binding commitment scheme. Then*

$$\forall\, \overline{\mathsf{com}}_{PQ}\ \exists\, p(\hat{s})\ \forall\, \overline{\mathsf{open}}_{PQ}\ \exists\, p(\hat{s}, s)\ \forall\, p(s_\circ | \hat{s}) : p(s \neq \hat{s}\ \wedge\ s = s_\circ) \leq \varepsilon\,,$$

where it is understood that $p(\hat{s}, s, s_\circ) := p(\hat{s}, s) \cdot p(s_\circ | \hat{s})$. Thus, the joint distribution $p(\hat{s}, s)$ is such that $p(s \neq \hat{s}\ \wedge\ s = s_\circ) \leq \varepsilon$ holds in particular for any function $s_\circ = f(\hat{s})$ of $\hat{s}$.

Proof. For given $\overline{\mathsf{com}}_{PQ}$ and $\overline{\mathsf{open}}_{PQ}$, let $p(\hat{s}, s)$ be as guaranteed by the fairly-weak-binding property. By Lemma 4.15, we may assume without loss of generality that $p(\hat{s}, s | s \neq \hat{s}) = p(\hat{s} | s \neq \hat{s})\, p(s | s \neq \hat{s})$. Then, by Lemma 2.7, we also

have that $p(s, s_\circ | s \neq \hat{s}) = p(s|s \neq \hat{s})\, p(s_\circ | s \neq \hat{s})$. It follows that

$$\begin{aligned}
p(s \neq \hat{s} \wedge s = s_\circ) &= p(s \neq \hat{s}) \cdot p(s = s_\circ | s \neq \hat{s}) \\
&= p(s \neq \hat{s}) \sum_{s^*_\circ} p(s = s^*_\circ \wedge s_\circ = s^*_\circ | s \neq \hat{s}) \\
&= p(s \neq \hat{s}) \sum_{s^*_\circ} p(s = s^*_\circ | s \neq \hat{s}) \cdot p(s_\circ = s^*_\circ | s \neq \hat{s}) \\
&= \sum_{s^*_\circ} p(s \neq \hat{s} \wedge s = s^*_\circ) \cdot p(s_\circ = s^*_\circ | s \neq \hat{s}) \\
&\leq \varepsilon \cdot \sum_{s^*_\circ} p(s_\circ = s^*_\circ | s \neq \hat{s}) \\
&= \varepsilon
\end{aligned}$$

where the inequality follows from the fact that $p(s \neq \hat{s} \wedge s = s^*_\circ) \leq \varepsilon$ for every *fixed* $s^*_\circ$. □

For the rest of the section, we take it as understood that we only consider classical attacks.

Theorem 4.17. *Let $(\mathcal{S}, \mathcal{S}')$ be an eligible pair of 2-prover commitment schemes, where $\mathcal{S}$ is ε-weak-binding and $\mathcal{S}'$ is δ-fairly-weak-binding, and let D be the domain of $\mathcal{S}$. Then, the composition $S \star \mathcal{S}'$ is a $(\varepsilon + (|D|-1) \cdot k(\mathcal{S}) \cdot \delta)$-weak-binding commitment scheme.*
In particular, if $\mathcal{S}$ is a bit *commitment scheme then $S \star \mathcal{S}'$ is a $(\varepsilon + k(\mathcal{S}) \cdot \delta)$-weak-binding.*

Proof. We follow the proof of Theorem 4.13, up to when it comes to choosing $y_\circ$. Let us first consider the case $m = 1$, i.e., $\mathcal{S}$ is a *bit*-commitment scheme. In that case, and assuming for the moment that $k(\mathcal{S}) = 1$, we let $y_\circ$ be the unique string that satisfies $\mathsf{Extr}(y_\circ, c) = s_\circ$, but where now $s_\circ := 1 - \tilde{s}$. We emphasize that for a fixed c, this choice of $y_\circ$ is *not* fixed anymore (in contrast to the choice in the proof of Theorem 4.13); namely, it is a function of $\tilde{s} = \mathsf{Extr}(\hat{y}, c)$, which in turn is a function of $\hat{y}$. Therefore, by Proposition 4.16, it still holds that $p(y \neq \hat{y} \wedge y = y_\circ | a) \leq \delta$, and we can conclude that

$$\begin{aligned}
p(\hat{s} \neq s \wedge s \neq \bot) &\leq p(\hat{s} \neq s \wedge s \neq \bot \wedge s = \tilde{s}) + p(s \neq \tilde{s} \wedge s \neq \bot) \\
&= p(\hat{s} \neq \tilde{s} \wedge s \neq \bot \wedge s = \tilde{s}) + p(s \neq \tilde{s} \wedge s = 1 - \tilde{s}) \\
&\leq p(\hat{s} \neq \tilde{s} \wedge \tilde{s} \neq \bot) + p(y \neq \hat{y} \wedge y = y_\circ) \\
&\leq p(\hat{s} \neq \tilde{s} \wedge \tilde{s} \neq \bot) + \textstyle\sum_a p(a)\, p(y \neq \hat{y} \wedge y = y_\circ | a) \\
&\leq \varepsilon + \textstyle\sum_a p(a)\, \delta \\
&= \varepsilon + \delta\,.
\end{aligned}$$

In the case that $k(\mathcal{S}) > 1$, we instead randomly select one of the at most $k(\mathcal{S})$ strings $y_\circ$ that satisfy $\mathsf{Extr}(y_\circ, c) = s_\circ = 1 - \tilde{s}$. Then, conditioned on a, $y_\circ$ is still independent of y given $\hat{y}$, so that Proposition 4.16 still applies, and we can argue as above, except that we get a factor $k(\mathcal{S})$ blow-up from $p(s \neq \tilde{s} \wedge s = 1 - \tilde{s}) \leq k(\mathcal{S}) \cdot p(y \neq \hat{y} \wedge y = y_\circ)$.

Finally, for the case $m > 1$, we first pick a random $s_\circ \in D \setminus \{\tilde{s}\}$, and then choose $y_\circ$ such that $\mathsf{Extr}(y_\circ, c) = s_\circ$, uniquely or at random, depending of $k(\mathcal{S})$. Conditioned on a, $y_\circ$ is still independent of y given $\hat{y}$, and therefore Proposition 4.16 still applies, but now we get an additional factor $(|D| - 1)$ blow-up from $p(s \neq \tilde{s} \wedge s \neq \perp) \leq (|D| - 1)\, p(s \neq \tilde{s} \wedge s = s_\circ)$. □

Remark 4.18. *Theorem 4.17 allows us to slightly improve the bound we obtain in Remark 4.14 on the Lunghi et al. multi-round commitment scheme. By Theorem 4.13, we can compose m instances of $\mathcal{CHSH}^n$ to obtain a $m \cdot 2^{-(n-1)/2}$-fairly-weak-binding string-commitment scheme. Then, we can compose the Crépeau et al.* bit *commitment scheme (i.e., the bit-commitment version of $\mathcal{CHSH}^n$), which is $2^{-(n-1)}$-weak-binding, with this fairly-weak-binding string-commitment scheme; by Theorem 4.17, this composition, which is the Lunghi et al. multi-round bit-commitment scheme, is $\left(m \cdot 2^{-(n-1)/2} + 2^{-(n-1)}\right)$-weak-binding.*

Finally, for completeness, we point out that the composition theorem also applies to two ordinary binding or weak-binding commitment schemes.

Theorem 4.19. *Let $(\mathcal{S}, \mathcal{S}')$ be an eligible pair of 2-prover commitment schemes, where $\mathcal{S}$ is ε-binding and $\mathcal{S}'$ is δ-binding. Then, the composition $S \star \mathcal{S}'$ is $(\varepsilon + \delta)$-binding. The same holds for the weak-binding property.*

Proof. The proof is almost the same as in Theorem 4.11 or Theorem 4.13, respectively, except that now there are no $s_\circ$ and $y_\circ$, and in the end we can simply conclude that

$$\begin{aligned} p(s \neq \hat{s} \wedge s \neq \perp) &\leq p(s \neq \hat{s} \wedge s \neq \perp \wedge s = \tilde{s}) + p(s \neq \tilde{s} \wedge s \neq \perp) \\ &\leq p(\tilde{s} \neq \hat{s} \wedge \tilde{s} \neq \perp) + p(y \neq \hat{y} \wedge y \neq \perp) \\ &\leq \varepsilon + \delta\,, \end{aligned}$$

where the second inequality holds since $y = \perp$ implies that $s = \mathsf{Extr}(y, c) = \perp$. □

4.4 Tightness

We now show that our composition result is nearly tight for $\mathcal{CHSH}^q$. Let $\mathcal{CHSH}^q_m$ be the m-fold composition of $\mathcal{CHSH}^q$ with itself, as defined in Remark 4.5. We show that if $q = p^{2k}$ for some prime p, this composed scheme can be ε-weak-binding as a bit-commitment scheme only if $\varepsilon \gtrsim \frac{1}{4} m \sqrt{q^{-1}}$. A

slightly weaker result was proved in [BC16], which shows that $\varepsilon \gtrsim \frac{1}{6}m2^{-n/2}$ for $q = 2^n$ with n even.[7] Furthermore, we show that, as a string-commitment scheme, $\mathcal{CHSH}^n_m$ can be ε-fairly-weak-binding only if $\varepsilon \gtrsim \frac{1}{2}m\sqrt{q^{-1}}$ (for $q = p^{2k}$).

Lemma 4.20. *Consider functions $X_q, Y_q : \mathbb{F}_q \times R_q \to \mathbb{F}_q$ where R_q is some finite set. Let*

$$\lambda_q = \max_{X_q, Y_q} \; p(X_q(a,r) + Y_q(s,r) = a \cdot s) \tag{4.1}$$

where a, s and r are selected uniformly at random. It holds that:

1. *There are X_q and Y_q such that $p(X_q(a,r) + Y_q(s,r) = a \cdot s) = \lambda_q$ for* all *$a, s \in \mathbb{F}_q$.*
2. *If $q = p^{2k}$ for some prime p, we have $\lambda_q = \Omega\big(\sqrt{q^{-1}}\big)$. Otherwise, we have $\lambda_q = \Omega\big(q^{-2/3}\big)$.*

Proof. Fix X'_q and Y'_q that achieve the maximum in Equation (4.1). We show that there also are functions X_q and Y_q such that for *any* a and s, $p(X_q(a,r) + Y_q(s,r) = a \cdot s) = \lambda_q$: Without loss of generality, X'_q and Y'_q depend only on a and s, not on r. Intuitively, X_q and Y_q do the following: they randomize their inputs a and s by adding uniformly random elements $r_a, r_s \in \mathbb{F}_q$, then apply X'_q and Y'_q, and finally remove the random terms again from the output. Formally, we let

$$\begin{aligned} X_q(a,(r_a,r_s)) &= X'_n(a + r_a) - ar_s - r_a r_s \\ Y_q(a,(r_a,r_s)) &= Y'_n(s + r_s) - r_a s \end{aligned}$$

For r_a and r_s uniformly random, we have $p(X'_q(a+r_a)+Y'_q(s+r_s) = as+ar_s+r_ar_s+sr_a) = \lambda_q$. Thus, it is easy to see that $p(X_q(a,(r_a,r_s))+Y_q(s,(r_a,r_s)) = as) = \lambda_q$.

The functions X_q and Y_q in Equation (4.1) describe classical strategies for the CHSH_q game and λ_q is the maximal winning probability that classical players can achieve in this game. As shown in [BS15], it holds that $\lambda_q = \Omega\big(\sqrt{q^{-1}}\big)$ for $q = p^{2k}$, and $\lambda_q = \Omega\big(q^{-2/3}\big)$ otherwise. □

The following lemma can be seen as a generalization of Theorem 3.12 to string-commitment schemes. Intuitively, it bounds the winning probability of the provers in the following game: First, they have to produce a commitment. Then, they receive a uniformly random string $s_\circ$ and, in order to win, they have to open the commitment to $s_\circ$. The winning probability in this game is at most $\varepsilon + 2^{-n}$, when the scheme is an ε-fairly-weak-binding n-bit string-commitment scheme.

[7] The paper states $\varepsilon \gtrsim \frac{1}{3}m2^{-n/2}$, but their binding definition is $p_0 + p_1 \leq 1 + \varepsilon$; to convert their bound to our definition (equivalent to $p_0 + p_1 \leq 1 + 2\varepsilon$), it must be multiplied by 1/2.

Lemma 4.21. *Let $\mathcal{S}$ be an ε-fairly-weak-binding commitment scheme with domain D. Fix a possible commit strategy $\overline{\mathsf{com}}_{PQ}$ for $\mathcal{S}$ and, for each $s_\circ \in D$, a possible opening strategy $\overline{\mathsf{open}}_{PQ}(s_\circ)$. Let $p(s|s_\circ)$ be the output distribution of $\mathcal{S}$ if the provers use $\overline{\mathsf{com}}_{PQ}$ and $\overline{\mathsf{open}}_{PQ}(s_\circ)$. Let $p(s_\circ)$ be the uniform distribution over D. Then, $p(s = s_\circ) = \sum_{s_\circ \in D} p(s_\circ)p(s = s_\circ|s_\circ) \leq \varepsilon + |D|^{-1}$.*

Proof. Let $p(\hat{s})$ be a distribution that satisfies Equation (3.4) for the commit strategy $\overline{\mathsf{com}}_{PQ}$. Now consider any consistent joint distribution $p(s, \hat{s}|s_\circ)$. Here, consistency also means that $p(\hat{s}|s_\circ) = p(\hat{s})$. Thus, for a uniformly random $s_\circ$, $p(\hat{s} = s_\circ) = |D|^{-1}$. By the ε-fairly-weak-binding property of $\mathcal{S}$, we have

$$\varepsilon \geq p(s \neq \hat{s} \wedge s = s_\circ) \geq p(s = s_\circ) - p(\hat{s} = s_\circ) = p(s = s_\circ) - |D|^{-1}$$

and thus our claim follows. □

With the help of the lemma above, is easy to see that λ_q limits the binding parameter of the one-round scheme $\mathcal{CHSH}^q$: If P sends $X_n(a, r)$ and Q sends $Y_n(s_\circ, r)$ for uniformly random r, then we have $p(s = s_\circ|a \neq 0) = \lambda_q$, and thus $p(s = s_\circ) \geq \lambda_q - q^{-1}$ for every $s_\circ$. Thus, by Lemma 4.21, $\mathcal{CHSH}^q$ can be ε-fairly-weak-binding only if $\varepsilon \geq \lambda_q - 2q^{-1}$. We now show that this bound scales approximately linearly with the number of rounds.

Theorem 4.22. *Let λ_q as in Lemma 4.20. For odd m, the $\mathcal{CHSH}^q_m$ commitment scheme can be ε-fairly-weak-binding as a string-commitment scheme only if*

$$\varepsilon \geq \frac{(m+1)\lambda_q}{2} - \frac{\left(m^2-1\right)\lambda_q^2}{8} - (m+1)q^{-1}.$$

If $m = o\left(\lambda_q^{-1}\right)$, it holds that $\varepsilon \geq \Omega(m\lambda_q)$. If, furthermore, $q = p^{2k}$, we have $\varepsilon \geq \Omega\left(m\sqrt{q^{-1}}\right)$; otherwise, $\varepsilon \geq \Omega\left(mq^{-2/3}\right)$.

Proof. Let $X_q(a, r)$ and $Y_q(b, r)$ be functions as in Lemma 4.20. We define a commit strategy $\overline{\mathsf{com}}_{PQ}$ and an opening strategy $\overline{\mathsf{open}}_{PQ}(s_\circ)$ for every $s_\circ$ which aims to open to $s_\circ$.

We assume that the provers have m uniformly random $r_i \in \mathbb{F}_q$ and $(m+1)/2$ uniformly random inputs r'_i, i odd, for X_q and Y_q as shared randomness. We write $c_i = (a_i, x_i)$ for the communication between the verifier and the active prover in round i, where the x_i are specified below. The dishonest provers exchange their communications as fast as possible, so in round $i + 2$, the active prover knows $c_1, \ldots, c_i$. Let $y_0 = s_\circ$ and for $i > 0$, let y_i such that $\mathsf{Extr}(y_i, c_i) = y_{i-1}$. Such a y_i exists and is unique if $a_i \neq 0$. We only specify our strategy for the case where the verifier's messages a_i are all non-zero and assume that the provers fail to open to $s_\circ$ otherwise. One can compute y_i from $c_1, \ldots, c_i$, so in round $i + 2$, the active prover can compute y_i.

If in any round i, the commitment is $(a_i, r_i + a_i \cdot y_{i-1})$, the provers can open to $s_\circ$ simply by following the honest strategy for $\mathcal{CHSH}^q_m$ from that round on.

The strategy described below is such that the provers have $(m+1)/2$ chances to bring about this situation with probability λ_q.

- Round 1 (commit): P produces a "fake commitment" $x_1 = X_q(a_1, r'_1)$.
- Round i, i even: Q computes $y'_{i-1} = Y_q(y_{i-2}, r'_{i-1})$, hoping that $x_{i-1} + y'_{i-1} = a_{i-1} \cdot y_{i-2}$, i.e., $y'_{i-1} = y_{i-1}$. He honestly commits to y'_{i-1} by computing $x_i = a_i \cdot y'_{i-1} + r_i$.
- Round $i+1$, i even: P checks if $y_{i-1} = y'_{i-1}$. If yes, both provers proceed honestly from this round on, i.e., they follow the honest strategy for $\mathcal{CHSH}^q_m$ in all subsequent rounds.[8] If not, P again produces a "fake commitment" $x_{i+1} = X_q(a_{i+1}, r'_{i+1})$.
- Round $m+1$: Q sends $y'_m = Y_q(y_{m-1}, r'_m)$ to V.

By definition, it holds that $y'_{i-1} = y_{i-1}$ if and only if $X_q(a_{i-1}, r'_{i-1}) + Y_q(y_{i-2}, r'_{i-1}) = a_{i-1} \cdot y_{i-2}$, which happens with probability λ_q. In this case, we have $c_i = (a_i, r_i + a_i \cdot y_{i-1})$, so the provers can indeed open to $s_\circ$ by proceeding honestly (ignoring completeness errors for now).

By definition of X_q, Y_q, and λ_q, if the provers use the strategy $\overline{\mathsf{open}}_{PQ}(s_\circ)$, then for

$$\lambda = 1-(1-\lambda_q)^{(m+1)/2} \geq \frac{(m+1)\lambda_q}{2} - \binom{(m+1)/2}{2}\lambda_q^2 = \frac{(m+1)\lambda_q}{2} - \frac{\big(m^2-1\big)\lambda_q^2}{8}$$

we have $p(s = s_\circ | a_1, \ldots, a_m \neq 0) = \lambda$. Thus, $p(s = s_\circ) \geq \lambda - mq^{-1}$ for all $s_\circ$. Applying Lemma 4.21, we conclude that the scheme can be ε-fairly-weak-binding only if

$$\varepsilon \geq \lambda - (m+1)q^{-1} \geq \frac{(m+1)\lambda_q}{2} - \frac{\big(m^2-1\big)\lambda_q^2}{8} - (m+1)q^{-1}$$

which is in $\Omega(m\lambda_q)$ if $m = o\big(\lambda_q^{-1}\big)$. Finally, we have $\Omega(m\lambda_q) = \Omega\big(m\sqrt{q^{-1}}\big)$ if $q = p^{2k}$ and $\Omega(m\lambda_q) = \Omega\big(mq^{-2/3}\big)$ otherwise, by claim 2 of Lemma 4.20. □

From the analysis in the above proof, we can also derive a version of the theorem for the bit-commitment scheme described in Proposition 3.11.

Corollary 4.23. *For even m, the commitment scheme $\mathcal{CHSH}^q_m$ can be ε-binding as a bit-commitment scheme only if*

$$\varepsilon \geq \frac{m\lambda_q}{4} - \frac{(m^2-2m)\lambda_q^2}{16} - (m+1)q^{-1}.$$

If $m = o\big(\lambda_q^{-1}\big)$, it holds that $\varepsilon \geq \Omega(m\lambda_q)$. If $q = p^{2k}$, we have $\varepsilon \geq \Omega\big(m\sqrt{q^{-1}}\big)$ and if it is odd, $\varepsilon \geq \Omega\big(mq^{-2/3}\big)$.

[8] Q can compute y_{i-1} in round $i+2$ and thus he too knows whether the provers should proceed honestly or not.

Proof. Let $\overline{\mathsf{com}}_P = \mathsf{com}_P(0)$, i.e., P produces an honest commitment to 0. Let $\overline{\mathsf{open}}_{PQ}(0) = \mathsf{open}_{PQ}$, i.e., the honest opening strategy. Since the provers play honestly, they are successful with probability at least $1 - (m+1)q^{-1}$.

For $\overline{\mathsf{open}}_{PQ}(1)$, let $s_\circ$ such that $\mathsf{Extr}(s_\circ, c_1) = 1$. The provers then use the strategy in the proof of Theorem 4.22 to produce a fake commitment c_1 and open it to $s_\circ$. Then, we have

$$p(b = 1 | a_1, \ldots, a_m \neq 0) \geq \frac{m\lambda_q}{2} - \frac{(m^2 - 2m)\lambda_q^2}{8} - q^{-1}$$

and thus,

$$p(b = 1) \geq \frac{m\lambda_q}{2} - \frac{(m^2 - 2m)\lambda_q^2}{8} - (m+1)q^{-1}.$$

It follows that

$$p(b = 0) + p(b = 1) \geq 1 + \frac{m\lambda_q}{2} - \frac{(m^2 - 2m)\lambda_q^2}{8} - (m+1)q^{-1}$$

and, by Theorem 3.12, the scheme can be ε-weak-binding only if

$$\varepsilon \geq \frac{m\lambda_q}{4} - \frac{(m^2 - 2m)\lambda_q^2}{16} - (m+1)q^{-1}.$$

□