

Cover Page



Universiteit Leiden



The handle <http://hdl.handle.net/1887/70036> holds various files of this Leiden University dissertation.

Author: Fillinger, M.J.

Title: Two-Prover Bit-Commitments: Classical, Quantum and Non-Signaling

Issue Date: 2019-03-19

Chapter 3

The Hiding and Binding Properties

3.1 Introduction

Formal security definitions form a crucial part of modern cryptography, where the aim is to mathematically prove the security of cryptographic schemes. Such definitions should capture and refine the informal intuition about the desired security requirements. For example, the informal goal of an encryption scheme is to keep a message secret from an adversary that does not know the appropriate key. The informal requirement guides the development of precise definitions (such as the modern game-based ones), but in formalizing it, one needs to also fill in many details that are left vague in the informal intuition.

Typically, we want security definitions to be as strong as possible while still being satisfiable, in order to offer security guarantees that are as strong as possible. Ideally, they should also be easy to work with. It is also desirable for them to be composable: Informally, this means that the stand-alone security of a scheme implies that security is still satisfied when the scheme is used as a building block in a larger system, and the security of the scheme propagates as expected to the larger system. The security proof for the larger system would not need to concern itself with the internal details of the components if the components satisfy composable security definitions.

In Chapter 1, we have discussed the informal security properties that a bit-commitment scheme should have. They should be hiding, meaning that a dishonest verifier cannot learn the committed value before the opening phase, and binding, meaning that after the commit phase, there is at most one value that can be revealed.

The (information-theoretic) hiding property is straightforward to define formally: even if the verifier is dishonest and arbitrarily deviates from the pro-

tocol, we require that the messages that he sees in the commit phase are statistically independent of the committed value. This definition can be relaxed to allow a limited amount of information by requiring the distributions to have small statistical distance from each other. Defining the information-theoretic binding property for two-prover schemes is more involved. The naive approach of requiring that the value to which the commitment can be opened should be uniquely determined by the verifier's view after the commit phase obviously leads to a contradiction with the information-theoretic binding property (see Section 1.2.1). Thus special care is necessary here.

In this chapter, we study several different and new definitions of the binding property which vary in certain technical aspects, and we analyze how they relate to each other. Our definitions vary in how we formalize the bit or string that the provers supposedly are committed to. One of our definitions, when restricted to bits, turns out to be equivalent to the sum-binding definition, while another one is strictly stronger. Our definitions also vary in how strict we are in not allowing the adversary to open to anything else than the committed value. Schemes that satisfy the less strict definition can quite easily be transformed into schemes that satisfy the stricter one by simply restricting their domain.

Naively, one might think that it suffices to consider the strongest achievable notion. However, some of our weaker definitions play a crucial role in our analysis of multi-round schemes (see Chapter 4).

We also prove that all of the definitions are satisfied by variants of the *CHSH* commitment scheme. This in particular is the first time the *CHSH* commitment scheme is proven secure as a *string* commitment scheme.

3.2 Defining The Binding Property

3.2.1 Possible Strategies

Like the sum-binding property, the binding properties we discuss in this chapter can only hold with respect to some restricted class of strategies. We call these strategies the *possible strategies*. In this chapter, we assume that all possible strategies $(\overline{\text{com}}_{PQ}, \overline{\text{open}}_{PQ})$ are *classical* interactive algorithms with access to joint randomness. We consider strategies that use quantum entanglement in Chapter 5. Our main result holds only in the classical case.

We assume that the set of possible strategies is the *convex hull* of a set of deterministic strategies. That is, if $(\overline{\text{com}}_{PQ}, \overline{\text{open}}_{PQ})$ is a possible randomized strategy, then the deterministic strategies that result from replacing the randomness with fixed values are possible as well. Conversely, if a strategy $(\overline{\text{com}}_{PQ}, \overline{\text{open}}_{PQ})$ instructs the provers to execute a possible strategy selected according to some probability distribution, then $(\overline{\text{com}}_{PQ}, \overline{\text{open}}_{PQ})$ is itself a possible strategy.

We make this assumption because we think of the set of possible strategies not as some arbitrary set, but as the strategies that are permitted by some constraints on the communication between the provers. If a set of strategies is permitted by those constraints, then executing a random strategy from this set should be possible as well. If a randomized strategy is permitted by those constraints, then deterministic strategies that result from replacing the randomness with fixed values should not violate the constraints either.

In the remainder of this chapter, we usually leave the set of possible strategies implicit and take it as understood that when we quantify over strategies, we refer only to possible strategies.

3.2.2 The (Strong) Binding Property

Intuitively, we say that a scheme is binding if after the commit phase there exists a string \hat{s} so that no matter what the provers do in the opening phase, the verifier will output either $s = \hat{s}$ or $s = \perp$ (except with small probability). We consider two definitions of the binding property which interpret this intuitive requirement in two different ways. In the first definition, which we introduce in this section, \hat{s} is a function of the provers' combined view immediately after the commit phase. In the second one, which we introduce in Section 3.2.3, \hat{s} is specified by its distribution only. Both of these definitions admit a composition theorem.

Definition 3.1 (Binding property). *A 2-prover commitment scheme \mathcal{S} is ε -binding if for every commit strategy $\overline{\text{com}}_{PQ}[\bar{\xi}_{PQ}]$ there exists a function $\hat{s}(\bar{\xi}_{PQ}, c)$ of the joint randomness $\bar{\xi}_{PQ}$ and the commitment¹ c such that for every opening strategy $\overline{\text{open}}_{PQ}$ it holds that $p(s \neq \hat{s}(\bar{\xi}_{PQ}, c) \wedge s \neq \perp) \leq \varepsilon$. In short:*

$$\forall \overline{\text{com}}_{PQ} \exists \hat{s}(\bar{\xi}_{PQ}, c) \forall \overline{\text{open}}_{PQ} : p(s \neq \hat{s} \wedge s \neq \perp) \leq \varepsilon. \quad (3.1)$$

The string-commitment scheme \mathcal{CHSH}^q does *not* satisfy this definition (the bit-commitment version does, as we will show): after the commit phase, the provers can still decide to open the commitment to a *fixed* string, chosen before the commit phase, or to a *random* string that is out of their control. We capture this property of \mathcal{CHSH}^q by the following relaxed version of the binding property: we allow V 's output s to be different from \hat{s} and \perp , but in this case the provers should have little control over s ; for any fixed *target string* s_\circ , it should be unlikely that $s = s_\circ$. Formally, this is captured as follows; we will show in Section 3.2.6 that \mathcal{CHSH}^q is fairly-binding in this sense.

Definition 3.2 (Fairly binding property). *A 2-prover commitment scheme \mathcal{S} is ε -fairly-binding if for every commit strategy $\overline{\text{com}}_{PQ}[\bar{\xi}_{PQ}]$ there exists a*

¹Recall that by convention (Remark 2.10), c equals the communication between V and the provers during the commit phase.

function $\hat{s}(\bar{\xi}_{PQ}, c)$ such that for every opening strategy $\overline{\text{open}}_{PQ}[\bar{\eta}_{PQ}]$ and all $s_o \in D$ it holds that $p(s \neq \hat{s}(\bar{\xi}_{PQ}, c) \wedge s = s_o) \leq \varepsilon$. In short:

$$\forall \overline{\text{com}}_{PQ} \exists \hat{s}(\bar{\xi}_{PQ}, c) \forall \overline{\text{open}}_{PQ} \forall s_o(\bar{\xi}_{PQ}, \bar{\eta}_{PQ}) : p(s \neq \hat{s} \wedge s = s_o) \leq \varepsilon. \quad (3.2)$$

If we want to show that a scheme is ε -(fairly-)binding (with respect to all possible strategies), it suffices to show that it is binding with respect to all possible *deterministic* strategies, as the following lemma shows.

Lemma 3.3. *Let \mathcal{S} be a commitment scheme that is ε -(fairly-)binding with respect to all possible deterministic strategies. Then it also is ε -(fairly-)binding with respect to all possible strategies.*

Proof. We prove the lemma for ε -binding schemes. It is easy to see how the proof can be adapted for fairly-binding schemes. Let $\overline{\text{com}}_{PQ}[\bar{\xi}_{PQ}]$ be a possible commit strategy. By our assumptions from Section 3.2.1, it follows that the strategy $\overline{\text{com}}_{PQ}^{r_c}$ where we set the joint randomness $\bar{\xi}_{PQ}$ to the value r_c is also possible for every r_c . By the assumed binding property, for every r_c , there exists a function \hat{s}_{r_c} such that for every deterministic opening strategy $\overline{\text{open}}_{PQ}$, we have $p(s \neq \hat{s}_{r_c} \wedge s \neq \perp) \leq \varepsilon$.

We define $\hat{s}(\bar{\xi}_{PQ}, c) = \hat{s}_{\bar{\xi}_{PQ}}(c)$. If the provers use $\overline{\text{com}}_{PQ}[\bar{\xi}]$ and any possible deterministic opening strategy, we have

$$p(s \neq \hat{s} \wedge s \neq \perp) = \sum_{r_c} p(\bar{\xi}_{PQ} = r_c) p(s \neq \hat{s}_{r_c}(c) \wedge s \neq \perp | \bar{\xi}_{PQ} = r_c) \leq \varepsilon.$$

It is straightforward to extend the above inequality to randomized opening strategies: the above inequality holds when we set the randomness to any particular value, and thus it also holds for the randomized strategy. \square

The next lemma shows that in Definition 3.2, instead of quantifying over strings s_o , we may also quantify over functions of the provers' randomness.

Lemma 3.4. *Let \mathcal{S} be an ε -fairly-binding scheme and $\overline{\text{com}}_{PQ}[\bar{\xi}_{PQ}]$ a commit strategy. There is a function $\hat{s}(\bar{\xi}_{PQ}, c)$ such that for every opening strategy $\overline{\text{open}}_{PQ}[\bar{\eta}_{PQ}]$ and every function $s_o(\bar{\xi}_{PQ}, \bar{\eta}_{PQ})$ with values in D , it holds that $p(s \neq \hat{s}(\bar{\xi}_{PQ}, c) \wedge s = s_o(\bar{\xi}_{PQ}, \bar{\eta}_{PQ})) \leq \varepsilon$.*

Proof. Let $\overline{\text{com}}_{PQ}^{r_c}$ and $\overline{\text{open}}_{PQ}^{r_o}$ be the deterministic strategies that results from fixing the randomness in $\overline{\text{com}}_{PQ}[\bar{\xi}_{PQ}]$ to r_c and the randomness in $\overline{\text{open}}_{PQ}[\bar{\eta}_{PQ}]$ to r_o . Fix an arbitrary function $s_o(\bar{\xi}_{PQ}, \bar{\eta}_{PQ})$. By the binding property, for every r , there is a function $\hat{s}_r(c)$ such that $p(s \neq \hat{s}_r(c) \wedge s = s_o(r_c, r_o) | \bar{\xi}_{PQ} = r_c, \bar{\eta}_{PQ} = r_o) \leq \varepsilon$. Setting $\hat{s}(\bar{\xi}_{PQ}, c) = \hat{s}_{\bar{\xi}_{PQ}}(c)$, we have

$$\begin{aligned} & p(s \neq \hat{s}(\bar{\xi}_{PQ}, c) \wedge s = s_o(\bar{\xi}_{PQ}, \bar{\eta}_{PQ})) \\ &= \sum_{r_c, r_o} p(\bar{\xi}_{PQ} = r_c) p(\bar{\eta}_{PQ} = r_o) p(s \neq \hat{s}(r_c, c) \wedge s = s_o(r_c, r_o) | \bar{\xi}_{PQ} = r_c, \bar{\eta}_{PQ} = r_o) \\ &\leq \varepsilon \end{aligned}$$

which proves our claim. \square

Remark 3.5. *Clearly, the binding property implies the fairly binding property. Furthermore, in the case of bit commitment schemes it obviously holds that $p(b \neq \hat{b} \wedge b \neq \perp) = p(b \neq \hat{b} \wedge b = 0) + p(b \neq \hat{b} \wedge b = 1)$, and thus the fairly-binding property implies the binding property with a factor-2 loss in the parameter. Furthermore, every fairly-binding string commitment scheme gives rise to a binding bit-commitment scheme in a natural way, as shown by the following proposition.*

Proposition 3.6. *Let \mathcal{S} be an ε -fairly-binding string-commitment scheme with domain D . Fix any two distinct strings $s_0, s_1 \in D$ and consider the bit-commitment scheme \mathcal{S}' defined as follows. To commit to $b \in \{0, 1\}$, the provers commit to s_b using \mathcal{S} , and in the opening phase V checks if $s = s_b$ for some bit $b \in \{0, 1\}$ and outputs this bit if it exists and else outputs $b = \perp$. Then, \mathcal{S}' is a 2ε -binding bit-commitment scheme.*

Proof. Fix some commit strategy $\overline{\text{com}}_{PQ}$ for \mathcal{S}' and note that it can also be used to attack \mathcal{S} . Thus, there exists a function $\hat{s}(\bar{\xi}_{PQ}, c)$ as in Definition 3.2. We define

$$\hat{b}(\bar{\xi}_{PQ}, c) = \begin{cases} 0 & \text{if } \hat{s}(\bar{\xi}_{PQ}, c) = s_0 \\ 1 & \text{otherwise} \end{cases}$$

Now fix an opening strategy $\overline{\text{open}}_{PQ}$ for \mathcal{S}' , which again is also a strategy against \mathcal{S} . Thus, we have $p(\hat{s} \neq s = s_\circ) \leq \varepsilon$ for any s_\circ (and in particular $s_\circ = s_0$ or s_1). This gives us

$$\begin{aligned} p(\hat{b} \neq b \neq \perp) &= p(\hat{b} = 1 \wedge b = 0) + p(\hat{b} = 0 \wedge b = 1) \\ &= p(\hat{s} \neq s_0 \wedge s = s_0) + p(\hat{s} = s_0 \wedge s = s_1) \\ &\leq p(\hat{s} \neq s_0 \wedge s = s_0) + p(\hat{s} \neq s_1 \wedge s = s_1), \\ &\leq 2\varepsilon \end{aligned}$$

and thus \mathcal{S}' is a 2ε -binding bit-commitment scheme. \square

Remark 3.7. *The proof of Proposition 3.6 generalizes in a straightforward way: given an ε -fairly-binding commitment scheme \mathcal{S} with domain D , and a subset $D' \subseteq D$, we define a commitment scheme $\mathcal{S}_{D'}$ with domain D' as follows: In the commit phase, the players use \mathcal{S} to produce a commitment to $s \in D'$. In the opening phase, the players run the opening phase of \mathcal{S} . If the result is in D' , V outputs it, and otherwise outputs \perp . Then, $\mathcal{S}_{D'}$ is $|D'| \varepsilon$ -binding.*

When $D' \not\subseteq D$, but $|D'| < |D|$, we can define a similar scheme by fixing an injection from D' to D . In particular, any ε -fairly-binding n -bit string-commitment scheme can be turned into a $2^k \varepsilon$ -binding k -bit string-commitment scheme for any $k < n$.

3.2.3 The Weak Binding Property

Here, we introduce yet another definition for the binding property. It is similar in spirit to Definition 3.1, but weaker. One advantage of this weaker notion is that it is also meaningful when considering quantum attacks, whereas Definition 3.1 is not. Note, however, that in the quantum setting, it does *not* suffice to only consider deterministic attacks. Therefore, results that depend on this property do not automatically carry over to the quantum setting. That includes Theorem 4.13, the composition theorem. In Section 3.2.4, we will see that for *bit*-commitment schemes, this weaker notion of the binding property is equivalent to the sum-binding definition, i.e., Definition 2.14.

Definition 3.8 (Weak binding property). *A 2-prover commitment scheme \mathcal{S} is ε -weak-binding if for all commit strategies $\overline{\text{com}}_{PQ}$ there exists a distribution $p(\hat{s})$ such that for every opening strategy $\overline{\text{open}}_{PQ}$ (which then fixes the distribution $p(s)$ of V 's output s) there is a consistent joint distribution $p(\hat{s}, s)$ such that $p(s \neq \hat{s} \wedge s \neq \perp) \leq \varepsilon$. In short:*

$$\forall \overline{\text{com}}_{PQ} \exists p(\hat{s}) \forall \overline{\text{open}}_{PQ} \exists p(\hat{s}, s) : p(s \neq \hat{s} \wedge s \neq \perp) \leq \varepsilon. \quad (3.3)$$

We also consider a related, i.e., “fairly”, version of this binding property, similar to Definition 3.2.

Definition 3.9 (Fairly weak binding property). *We say that a 2-prover commitment scheme \mathcal{S} is ε -fairly-weak-binding if for all commit strategies $\overline{\text{com}}_{PQ}$ there exists a distribution $p(\hat{s})$ such that for every opening strategy $\overline{\text{open}}_{PQ}$ (which then fixes the distribution $p(s)$ of V 's output s) there is a consistent joint distribution $p(\hat{s}, s)$ so that for all $s_o \in \{0, 1\}^n$, $p(s \neq \hat{s} \wedge s = s_o) \leq \varepsilon$. In short:*

$$\forall \overline{\text{com}}_{PQ} \exists p(\hat{s}) \forall \overline{\text{open}}_{PQ} \exists p(\hat{s}, s) \forall s_o : p(s \neq \hat{s} \wedge s = s_o) \leq \varepsilon. \quad (3.4)$$

Remark 3.10. *Lemma 3.3 and Remark 3.5 also hold for the weak binding properties. Furthermore, it is easy to see that the binding and fairly-binding properties imply their weak counterparts.*

Proposition 3.11. *Let \mathcal{S} be a string-commitment scheme and define \mathcal{S}' as in Proposition 3.6. If \mathcal{S} is ε -fairly-weak-binding, then \mathcal{S}' is a 2ε -weak-binding bit-commitment scheme.*

Proof. The proof of Proposition 3.6 can be easily adapted: Let $p(\hat{s})$ be as required by Definition 3.9. We define $p(\hat{b})$ by taking the marginal of $p(\hat{s}, \hat{b})$ where $\hat{b} = 0$ if $\hat{s} = s_0$, and $\hat{b} = 1$ otherwise. An opening strategy $\overline{\text{open}}_{PQ}$ for \mathcal{S}' can also be viewed as a strategy for \mathcal{S} . As such, there is a joint distribution $p(\hat{s}, s)$ as required by Definition 3.8 which we can extend to $p(\hat{s}, s, b)$ by setting $b = 0$ if $s = s_0$, $b = 1$ if $s = s_1$ and $b = \perp$ otherwise. We define $p(\hat{b}, b) := \sum_{\hat{s}, s} p(\hat{s}, \hat{b}) \cdot p(s, b | \hat{s})$. As in the proof of Proposition 3.6, one can easily check that $p(\hat{b} \neq b \neq \perp) \leq 2\varepsilon$ holds. \square

3.2.4 Relations Between The Definitions

Here, we show that in case of *bit*-commitment schemes, the weak binding property as introduced in Definition 3.8 above is actually *equivalent* to the sum-binding-definition. Even though our focus in this chapter is on classical attacks, the proof immediately carries over to quantum attacks as well.

Theorem 3.12. *A 2-prover bit-commitment scheme is ε -sum-binding if and only if it is ε -weak-binding.*

Proof. First, consider a scheme that is ε -binding according to Definition 2.14. Fix an arbitrary commit strategy $\overline{\text{com}}_{PQ}$. Let $\overline{\text{open}}_{PQ}^0$ and $\overline{\text{open}}_{PQ}^1$ be opening strategies so that $p_0 = p(b_0 = 0)$ and $p_1 = p(b_1 = 1)$ are both maximized, where $b_i \in \{0, 1, \perp\}$ is V 's output when the dishonest provers use the commit strategy $\overline{\text{com}}_{PQ}$ and opening strategy $\overline{\text{open}}_{PQ}^i$. Let ε' be such that $p_0 + p_1 = 1 + 2\varepsilon'$. Since the scheme is ε -binding, we have $\varepsilon' \leq \varepsilon$. We define the distribution $p(\hat{b})$ as $p(\hat{b} = 0) := p_0 - \varepsilon'$ and $p(\hat{b} = 1) := p_1 - \varepsilon'$. To see that this is indeed a probability distribution, note that $p_0, p_1 \geq 2\varepsilon'$ (otherwise, we would have $p_0 > 1$ or $p_1 > 1$) and that $p(\hat{b} = 0) + p(\hat{b} = 1) = p_0 + p_1 - 2\varepsilon' = 1$. Now we consider an arbitrary opening strategy $\overline{\text{open}}_{PQ}$ which fixes a distribution $p(b)$. By definition of p_0 and p_1 , we have $p(b = i) \leq p_i$ and thus $p(b = i) \leq p(\hat{b} = i) + \varepsilon' \leq p(\hat{b} = i) + \varepsilon$. By Lemma 2.2, there exists a consistent joint distribution $p(\hat{b}, b)$ such that $p(\hat{b} = b = i) = \min\{p(b = i), p(\hat{b} = i)\}$. We wish to bound $p(\hat{b} \neq b \wedge b \neq \perp) = p(\hat{b} = 0 \wedge b = 1) + p(\hat{b} = 1 \wedge b = 0)$. For $i \in \{0, 1\}$, it holds that

$$\begin{aligned} p(\hat{b} = 1 - i \wedge b = i) &= p(b = i) - p(\hat{b} = b = i) \\ &= p(b = i) - \min\{p(\hat{b} = i), p(b = i)\} \\ &= \max\{0, p(b = i) - p(\hat{b} = i)\} \\ &\leq \varepsilon \end{aligned}$$

and furthermore, there is at most *one* $i \in \{0, 1\}$ such that $p(b = i) > p(\hat{b} = i)$, for if $p(b = i) > p(\hat{b} = i)$ for both $i = 0$ and $i = 1$, then $p(b = 0) + p(b = 1) > p(\hat{b} = 0) + p(\hat{b} = 1) = 1$ which is a contradiction. Thus, we have $p(\hat{b} \neq b \wedge b \neq \perp) \leq \varepsilon$. This proves one direction of our claim.

For the other direction, consider a scheme that is ε -weak-binding. Fix $\overline{\text{com}}_{PQ}$ and let $p(\hat{b})$ be a distribution such that for every opening strategy $\overline{\text{open}}_{PQ}$, there is a joint distribution $p(\hat{b}, b)$ with $p(\hat{b} \neq b \neq \perp) \leq \varepsilon$. Now consider two opening strategies $\overline{\text{open}}_{PQ}^0$ and $\overline{\text{open}}_{PQ}^1$ which give distributions $p(b_0)$ and $p(b_1)$. We need to bound $p(b_0 = 0) + p(b_1 = 1)$. There is a joint

distribution $p(\hat{b}, b_0)$ such that $p(\hat{b} \neq b_0 \neq \perp) \leq \varepsilon$ and likewise for b_1 . Thus,

$$\begin{aligned} & p(b_0 = 0) + p(b_1 = 1) \\ = & p(\hat{b} = 0, b_0 = 0) + p(\hat{b} = 1, b_0 = 0) + p(\hat{b} = 0, b_1 = 1) + p(\hat{b} = 1, b_1 = 1) \\ \leq & p(\hat{b} = 0) + p(\hat{b} = 1) + p(\hat{b} \neq b_0 \neq \perp) + p(\hat{b} \neq b_1 \neq \perp) \\ \leq & 1 + 2\varepsilon \end{aligned}$$

which proves the other direction. \square

Remark 3.13. *By Remark 3.10, it follows that Definition 3.1 also implies the sum-binding-definition. In fact, Definition 3.1 is strictly stronger (and hence, also strictly stronger than the weak-binding definition). Consider the following (artificial and very non-complete) scheme: In the commit phase, V chooses a uniformly random bit and sends it to the provers, and then accepts anything or rejects anything during the opening phase, depending on that bit. Then, $p_0 + p_1 = 1$, yet a commitment can be opened to $1 - \hat{b}$ (no matter how \hat{b} is defined) with probability $\frac{1}{2}$.*

Since a non-complete separation example may not be fully satisfying, we note that it can be converted into a complete (but even more artificial) scheme. Fix a “good” (i.e., complete, hiding and binding with low parameters) scheme and call our example scheme above the “bad” scheme. We define a combined scheme as follows: At the start, the first prover can request either the “good” or “bad” scheme to be used. The honest prover is instructed to choose the former, guaranteeing completeness. The dishonest prover may choose the latter, so the combined scheme inherits the binding properties of the “bad” scheme: it is binding according to the sum-binding-definition, but not according to Definition 3.1.

3.2.5 Simultaneous Opening

The binding definitions from the previous sections are useful for proving our composition theorem, but it is not clear how to prove in a straightforward way that a commitment scheme satisfies those definitions. In this section, we propose another definition which is easier to check and which implies the binding properties from the previous sections (with some loss in the parameter). We then use this result in Section 3.2.6 to prove that \mathcal{CHSH}^q is fairly-binding.

This binding property is based on the intuition that it should not be possible to open a commitment to two different values *simultaneously* (except with small probability). For this, we observe that when considering a commit strategy $\overline{\text{com}}_{PQ}$, as well as *two* opening strategies $\overline{\text{open}}_{PQ}$ and $\overline{\text{open}}'_{PQ}$, we can run both opening strategies *simultaneously* on the produced commitment with two independent copies of open_V , by applying $\overline{\text{open}}_{PQ}$ and $\overline{\text{open}}'_{PQ}$ to two copies of the respective internal states of P and Q). This gives rise to a

joint distribution $p(s, s')$ of the respective outputs s and s' of the two copies of open_V .

Definition 3.14 and Theorem 3.18 were first considered in [Sca16].

Definition 3.14. A 2-prover commitment scheme \mathcal{S} is ε -binding in the sense of simultaneous opening if for all $\overline{\text{com}}_{PQ}$ and all pairs of opening strategies $\overline{\text{open}}_{PQ}$ and $\overline{\text{open}}'_{PQ}$, it holds that $p(s \neq s' \wedge s \neq \perp \wedge s' \neq \perp) \leq \varepsilon$.

Definition 3.15. A 2-prover commitment scheme \mathcal{S} is ε -fairly-binding in the sense of simultaneous opening if for all $\overline{\text{com}}_{PQ}$, all pairs of opening strategies $\overline{\text{open}}_{PQ}$ and $\overline{\text{open}}'_{PQ}$, and all pairs s_o, s'_o of distinct strings, it holds that $p(s = s_o \wedge s' = s'_o) \leq \varepsilon$.

Remark 3.16. Also for this notion of binding, it is sufficient to consider deterministic strategies, as can easily be seen.

Remark 3.17. It follows directly from Eq. (2.1) that every bit-commitment scheme that is ε -fairly-binding in the sense of simultaneous opening (against classical attacks) is $\varepsilon/2$ -sum-binding (and thus also according to Definition 3.8). The converse is not true though: the schemes from Remark 3.13 again serve as counterexamples.

Theorem 3.18. Let $\mathcal{S} = (\text{com}_{PQV}, \text{open}_{PQV})$ be a 2-prover commitment scheme. If it is ε -binding in the sense of simultaneous opening, and open_V is deterministic, then \mathcal{S} is $2\sqrt{\varepsilon}$ -binding.

Proof. By Lemma 3.3, it suffices to consider deterministic strategies for the provers. We fix some deterministic commit strategy $\overline{\text{com}}_{PQ}$ and an enumeration $\{\overline{\text{open}}^i_{PQ}\}_{i=1}^N$ of all deterministic opening strategies. Since we assume that open_V is deterministic, for any fixed deterministic opening strategy for the provers, the verifier's output s is a *function* of the commitment c . Thus, for each opening strategy $\overline{\text{open}}^i_{PQ}$ there is a function f_i such that the verifier's output is $s = f_i(c)$. We will now define the function $\hat{s}(c)$ that satisfies the properties required by Definition 3.1. We will now define the function $\hat{s}(c)$ that satisfies the properties required by Definition 3.1. Our definition depends on a parameter $\alpha > 0$ which we fix later. To define \hat{s} , we partition the set C of all possible commitments into disjoint sets $C = R \cup \bigcup_i C_i$ that satisfy the following three properties for every i :

- $f_i(c) \neq \perp$ for all $c \in C_i$,
- $p(c \in C_i) \geq \alpha$ or $C_i = \emptyset$,
- and $p(c \in R \wedge f_i(c) \neq \perp) < \alpha$

The second property implies that there are at most α^{-1} non-empty sets C_i . It is easy to see that such a partitioning exists: Start with $R = C$ and while there exists some i with $p(c \in R \wedge f_i(c) \neq \perp) \geq \alpha$, let $C_i = \{c \in R \mid f_i(c) \neq \perp\}$

and remove the elements of C_i from R . For any $c \in C$, we now define $\hat{s}(c)$ as follows. We set $\hat{s}(c) = f_i(c)$ if $c \in C_i$ and $\hat{s}(c) = 0$ for $c \in R$.

Now fix some opening strategy $\overline{\text{open}}_{PQ}^i$ and write s_i for the verifier's output. It follows that

$$\begin{aligned}
& p(s_i \neq \hat{s}(c) \wedge s_i \neq \perp) \\
&= p(f_i(c) \neq \hat{s}(c) \wedge f_i(c) \neq \perp) \\
&\leq p(c \in R \wedge f_i(c) \neq \perp) + \sum_j p(f_i(c) \neq \hat{s}(c) \wedge f_i(c) \neq \perp \wedge c \in C_j) \\
&< \alpha + \sum_{j: C_j \neq \emptyset} P(f_i(c) \neq f_j(c) \wedge f_i(c) \neq \perp \wedge f_j(c) \neq \perp) \\
&\leq \alpha + \alpha^{-1} \cdot \varepsilon
\end{aligned}$$

where the final inequality holds because $p(f_i(c) \neq f_j(c) \wedge f_i(c) \neq \perp \wedge f_j(c) \neq \perp) \leq \varepsilon$ by the assumed binding property. It is easy to see that the upper bound $\alpha + \alpha^{-1} \cdot \varepsilon$ is minimized by setting $\alpha = \sqrt{\varepsilon}$. We conclude that $p(s_i \neq \hat{s}(c) \wedge s_i \neq \perp) < 2\sqrt{\varepsilon}$. \square

Theorem 3.19. *Let $\mathcal{S} = (\text{com}_{PQV}, \text{open}_{PQV})$ be a 2-prover commitment scheme. If \mathcal{S} is ε -fairly-binding in the sense of simultaneous opening and open_V is deterministic, then \mathcal{S} is $2\sqrt{\varepsilon}$ -fairly-binding.*

Proof. It again suffices to consider deterministic strategies for the provers. As in the previous proof, we fix a deterministic commit strategy $\overline{\text{com}}_{PQ}$ and an enumeration $\{\overline{\text{open}}_{PQ}^i\}_{i=1}^N$ of the deterministic opening strategies. The verifier's output when the provers use $\overline{\text{open}}_{PQ}^i$ is a function $f_i(c)$ of the commitment. We now define the function $\hat{s}(c)$ that satisfies the properties required by Definition 3.2. Our definition again depends on a parameter $\alpha > 0$. We partition the set C of all possible commitments into disjoint sets $R \cup \bigcup_{s,i} C_{s,i} = C$ that satisfy the following three properties for every i and every s :

$$C_{s,i} \subseteq f_i^{-1}(\{s\}), \quad p(c \in C_{s,i}) \geq \alpha \text{ or } C_{s,i} = \emptyset, \quad \text{and} \quad p(c \in R \wedge f_i(c) = s) < \alpha.$$

The second property implies that there are at most α^{-1} non-empty sets $C_{s,i}$. Similar to the previous proof, it is easy to see that such a partitioning exists. For any $c \in C$, we now define $\hat{s}(c)$ as follows. We set $\hat{s}(c) = s$ for $c \in C_{s,i}$ and $\hat{s}(c) = 0$ for $c \in R$.

Now fix some opening strategy $\overline{\text{open}}_{PQ}^i$ and a string s_o , and write s_i for the verifier's output. Using $C_{\neq s_o}$ as a shorthand for $\bigcup_{s \neq s_o} \bigcup_j C_{s,j}$, we note

that if $\hat{s}(c) \neq s_o$ then $c \in R \cup C_{\neq s_o}$. Thus, it follows that

$$\begin{aligned}
& p(s_i \neq \hat{s}(c) \wedge s_i = s_o) \\
&= p(\hat{s}(c) \neq s_o \wedge s_i = s_o) \\
&\leq p(c \in (R \cup C_{\neq s_o}) \wedge f_i(c) = s_o) \\
&= p(c \in R \wedge f_i(c) = s_o) + \sum_{s \neq s_o, j} p(c \in C_{s,j} \wedge f_i(c) = s_o) \\
&\leq p(c \in R \wedge f_i(c) = s_o) + \sum_{\substack{s \neq s_o, j \\ \text{s.t. } C_{s,j} \neq \emptyset}} p(f_j(c) = s \wedge f_i(c) = s_o) \\
&< \alpha + \alpha^{-1} \cdot \varepsilon
\end{aligned}$$

where the final inequality holds because $p(f_j(c) = s \wedge f_i(c) = s_o) \leq \varepsilon$ by the assumed binding property. Again, we minimize the upper bound by setting $\alpha = \sqrt{\varepsilon}$ which completes the proof. \square

For the fairly-weak-binding property, we can get better parameters. Also note that we do not require open_V to be deterministic here.

Theorem 3.20. *Every 2-prover commitment scheme \mathcal{S} that is ε -fairly-binding in the sense of simultaneous opening is $\sqrt{2\varepsilon}$ -fairly-weak-binding.*

Proof. Fix a commit strategy $\overline{\text{com}}_{PQ}$ against \mathcal{S} . Enumerate all strings in the domain D of \mathcal{S} as s_o^1, \dots, s_o^d , and for every i , let $\overline{\text{open}}_{PQ}^i$ be an opening strategy maximizing $p_i := p(s_i = s_o^i)$, where s_i is the output of the verifier when P and Q use this strategy. We assume without loss of generality that the p_i are in descending order. We define $p(\hat{s})$ as follows. Let $N \geq 2$ be an integer which we will fix later. By Definition 3.14 and Inequality (2.2), it holds that

$$\sum_{i=1}^N p_i \leq 1 + \binom{N}{2} \cdot \varepsilon = 1 + \frac{N(N-1)}{2} \cdot \varepsilon$$

where we let $p_i = 0$ for $i > d$ in case $N > d$. We would like to define $p(\hat{s})$ as $p(\hat{s} = s_o^i) := p_i - (N-1)\varepsilon/2$ for all $i \leq N, d$; however, this is not always possible because $p_i - (N-1)\varepsilon/2$ may be negative. To deal with this, let N' be the largest integer such that $N' \leq N$ and $p_1, \dots, p_{N'} \geq (N-1)\varepsilon/2$. (We take $N = 0$ if $p_1 < (N-1)\varepsilon/2$.) It follows that

$$\sum_{i=1}^{N'} p_i \leq 1 + \frac{N'(N'-1)}{2} \cdot \varepsilon \leq 1 + \frac{N'(N-1)}{2} \cdot \varepsilon$$

and thus

$$\sum_{i=1}^{N'} p_i = 1 + \frac{N'(N-1)}{2} \cdot \tilde{\varepsilon}$$

for some $\tilde{\varepsilon} \leq \varepsilon$. We now set $p(\hat{s})$ to be $p(\hat{s} = s_i) := p_i - (N-1)\tilde{\varepsilon}/2 \geq p_i - (N-1)\varepsilon/2 \geq 0$ for all $i \leq N'$. Now consider an opening strategy $\overline{\text{open}}_{PQ}$ and let $p(s)$ be the resulting output distribution. By definition of the p_i , it follows that $p(s = s_{\circ}^i) \leq p_i$ for all $i \leq d$, and $p_i \leq p(\hat{s} = s_{\circ}^i) + (N-1)\varepsilon/2$ for all $i \leq N'$. By Lemma 2.2, we can conclude that there exists a consistent joint distribution $p(\hat{s}, s)$ with $p(\hat{s} = s = s_{\circ}^i) = \min\{p(s = s_{\circ}^i), p(\hat{s} = s_{\circ}^i)\} \geq p(s = s_i) - (N-1)\varepsilon/2$ for all $i \leq N'$, and thus $p(\hat{s} \neq s = s_{\circ}^i) = p(s = s_{\circ}^i) - p(\hat{s} = s = s_{\circ}^i) \leq (N-1)\varepsilon/2$ for all $i \leq N'$. Furthermore, when $N' < i \leq N$, we have $p(\hat{s} \neq s = s_{\circ}^i) = p(s = s_{\circ}^i) \leq p_i < (N-1)\varepsilon/2$ by definition of N' . Since the p_i are sorted in descending order, it follows that for all $i > N$

$$p(\hat{s} \neq s = s_{\circ}^i) = p(s = s_{\circ}^i) \leq p_i \leq p_N \leq \frac{1}{N} \sum_{i=1}^N p_i \leq \frac{1}{N} + \frac{N-1}{2} \cdot \varepsilon$$

and thus, we have shown for all $s_{\circ} \in D$ that

$$p(\hat{s} \neq s = s_{\circ}) \leq \frac{1}{N} + \frac{N-1}{2} \cdot \varepsilon.$$

We now select N so that this value is minimized: it is easy to verify that the function $f : \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$, $x \mapsto 1/x + (x-1)\varepsilon/2$ has its global minimum in $\sqrt{2/\varepsilon}$; thus, we pick $N := \lceil \sqrt{2/\varepsilon} \rceil$, which gives us

$$p(\hat{s} \neq s = s_{\circ}) \leq \frac{1}{N} + \frac{N-1}{2} \cdot \varepsilon \leq \frac{1}{\sqrt{2/\varepsilon}} + \frac{\sqrt{2/\varepsilon}}{2} \cdot \varepsilon = \sqrt{2\varepsilon}$$

for any $s_{\circ} \in D$, as claimed. \square

3.2.6 Security of \mathcal{CHSH}^q

Using the results from the previous section, we now show that \mathcal{CHSH}^q is a fairly-binding string-commitment scheme. It is understood that the possible attacks against \mathcal{CHSH}^q are those where the provers do not communicate.

Proposition 3.21. *The string-commitment scheme \mathcal{CHSH}^q is q^{-1} -fairly-binding in the sense of simultaneous opening.*

Proof. By Remark 3.16, it suffices to consider deterministic attack strategies. Fix a deterministic strategy $\overline{\text{com}}_{PQ}$ and two deterministic opening strategies $\overline{\text{open}}_{PQ}$ and $\overline{\text{open}}'_{PQ}$. The strategy $\overline{\text{com}}_{PQ}$ specifies P 's output x as a function $f(a)$ of the verifier's message a . The opening strategies are described by constants y and y' . By definition of \mathcal{CHSH}^q , $s = s_{\circ}$ implies $f(a) + y = a \cdot s_{\circ}$ and likewise, $s' = s'_{\circ}$ implies $f(a) + y' = a \cdot s'_{\circ}$. The condition $s = s_{\circ} \wedge s' = s'_{\circ}$ can hold only if $a = (y - y') / (s_{\circ} - s'_{\circ})$. It follows that $p(s = s_{\circ} \wedge s' = s'_{\circ}) \leq p(a = (y - y') / (s_{\circ} - s'_{\circ})) \leq q^{-1}$, which proves our claim. \square

From Proposition 3.21, Theorem 3.19 and Theorem 3.20, we conclude that the following corollaries hold.

Corollary 3.22. \mathcal{CHSH}^q is $2\sqrt{q^{-1}}$ -fairly-binding.

Corollary 3.23. \mathcal{CHSH}^q is $\sqrt{2q^{-1}}$ -fairly-weak-binding.

Remark 3.24. *It is not too hard to see that Corollary 3.23 above implies an upper bound on the classical value ω of the game CHSH_{2^n} considered in [BS15] of $\omega(\text{CHSH}_{2^n}) \leq 2^{-\frac{n-1}{2}} + 2^{-n}$. As such, Theorem 1.3 in [BS15] implies that the above ε is asymptotically optimal for odd n , i.e., the square root loss to the binding property of the bit-commitment version is unavoidable (for odd n).*

As for security against quantum attacks, we point out that [BS15, RAM16] provide an upper bound on the quantum value $\omega^(\text{CHSH}_q)$ of general finite-field CHSH; however, this does not directly imply security against quantum attacks of \mathcal{CHSH}^q as a (fairly-weak-binding) string-commitment scheme.*

Furthermore, we show that a variant of \mathcal{CHSH}^q is $2\sqrt{q^{-1}}$ -binding. However, this variant requires the opening information to be twice as large as the domain of \mathcal{CHSH}^q , so it is not possible to compose multiple instances of this variant using our composition theorem (see Definition 4.1).

Corollary 3.25. *Let \mathcal{CHSH}_+^q be the scheme defined as follows: The commit phase is the same as in \mathcal{CHSH}^q . In the opening phase, Q sends the opening information and the string s that the provers committed to. Then, V opens the commitment as in \mathcal{CHSH}^q and checks if the result equals the string s he received from Q . If yes, he outputs s and if not, \perp . This scheme is $2\sqrt{q^{-1}}$ -binding and $\sqrt{2q^{-1}}$ -weak-binding.*

Proof. Let $\overline{\text{open}}_{PQ}$ be a dishonest strategy for the opening phase of \mathcal{CHSH}_+^q . Let s_\circ be the string that Q sends along with the opening information. Since V does not send any messages to Q in \mathcal{CHSH}_+^q , the string s_\circ is computed as a function of the provers' randomness. From the strategy $\overline{\text{open}}_{PQ}$, a strategy $\overline{\text{open}}'_{PQ}$ for \mathcal{CHSH}^q can be extracted by simply leaving out s_\circ . By Lemma 3.4 and Corollary 3.22, we conclude that if the provers use $\overline{\text{open}}'_{PQ}$ in \mathcal{CHSH}^q , we have $p(s \neq \hat{s} \wedge s = s_\circ) \leq 2\sqrt{q^{-1}}$. It follows that when they use $\overline{\text{open}}_{PQ}$ in \mathcal{CHSH}_+^q , we have $p(s \neq \hat{s} \wedge s \neq \perp) \leq 2\sqrt{q^{-1}}$. The result for the weak-binding property follows similarly, using Remark 3.10 and Corollary 3.23. \square

While it may seem like a similar idea could be used to transform *any* fairly-binding scheme into a binding scheme at the cost of increasing the size of the opening information, the proof above relies on the assumption that the second prover can not choose the message s_\circ depending on any information sent by the verifier. Otherwise, Lemma 3.4 does not apply.

As a counter-example, consider another variant of \mathcal{CHSH}^q similar to \mathcal{CHSH}_+^q where in the opening phase, P sends the string s and Q sends the opening information. The following strategy breaks the binding property of this variant:

In the commit phase, P sends $x = 1$. The provers can open to 0 by sending 0 and 1 respectively in the opening phase. Since this strategy always opens to 0, $p(\hat{s} = 0)$ needs to be large. On the other hand, if P sends $s_\circ = a^{-1}$ (if it exists) and Q sends 0, it holds that $p(s = s_\circ)$ is large and $p(s_\circ \neq \hat{s})$ is small.