Resolving the Round Complexity of Non-Malleable Commitments

Vipul Goyal* Silas Richelson†

Abstract

We present a new non-malleable commitment protocol. Our protocol has the following features:

- The protocol has only *three rounds* of interaction. Pass (TCC 2013) showed an impossibility result for a two-round non-malleable commitment scheme w.r.t. a black-box reduction to any "standard" intractability reduction. Thus, this resolves the round complexity of non-malleable commitments at least w.r.t. black-box security reductions. Our construction is secure as per the standard notion of non-malleability w.r.t. commitment.
- Our protocol is *truly efficient*. In our basic protocol, the entire computation of the committer is dominated by just three invocations of a non-interactive statically binding commitment scheme, while, the receiver computation (in the commitment stage) is limited to just sampling a random string. Unlike many previous works, we directly construct a protocol for large tags and hence avoid any non-malleability amplification steps.
- Our protocol is based on a black-box use of any non-interactive statistically binding commitment scheme. Such schemes, in turn, can be based on any one-way permutations (or one-way functions at the cost of an extra initialization round). Previously, the best known black-box construction of non-malleable commitments required a larger (constant) number of rounds.
- Our construction is public-coin and makes use of only black-box simulation. Prior to our work, no public-coin constant round non-malleable commitment schemes were known based on black-box simulation.

Our techniques depart *significantly* from the techniques used previously to construct non-malleable commitment schemes. As a main technical tool, we rely on non-malleable codes in the split state model. Our proofs of security are purely combinatorial in nature.

*Microsoft Research, India. Email: vipul@microsoft.com.

†MIT. Email: sirichel@csail.mit.edu. Work done in part while visiting Microsoft Research, India.

1 Introduction

Man-in-the-middle (MIM) attacks are one of the most basic attacks in cryptography. The notion of non-malleable commitments was introduced in a seminal work of Dolev, Dwork and Naor [DDN91] as a countermeasure against such attacks. Since their introduction, non-malleable commitments have proven to be capable of preventing MIM attacks in a variety of settings. Non-malleability lies at the heart of secure protocol composition, it allows for round-efficient secure multi-party computation and gives applications to areas as diverse as position based cryptography [CGMO09].

A *commitment scheme* is a useful two party protocol which allows a committer, C, to send a representation of his message v, $\mathrm{Com}(v;r)$ to a receiver, R, in such a way so that 1) R learns nothing about v until C chooses to open his commitment and 2) C is bound to v and cannot open $\mathrm{Com}(v)$ to any value $v' \neq v$. A commitment scheme is *non-malleable* if for every message v, no MIM adversary, intercepting a commitment $\mathrm{Com}(v;r)$ and modifying it at will, is able to efficiently generate a commitment $\mathrm{Com}(\tilde{v};\tilde{r})$ to a related message $\tilde{v}$. Interest in non-malleable commitments is motivated both by the central role that they play in securing protocols under composition (see for example [CLOS02, LPV09]) and by the unfortunate reality that many widely used commitment schemes are actually highly malleable. Indeed, man-in-the-middle (MIM) attacks occur quite naturally when multiple concurrent executions of protocols are allowed, and can be quite devastating.

Since their conceptualization by Dolev, Dwork and Naor [DDN91], non-malleable commitments have been studied extensively, and, with increasing success in terms of characterizing their round complexity. [DDN91] gave a construction of non-malleable commitments which requires O(log λ) rounds, where λ is the security parameter. Barak [Bar02] gave a constant round construction based on non-black-box simulation (which was further improved by Pass and Rosen [PR05b]). More recently, constant round protocols for non-malleable commitment with a black-box proof of security are given by Goyal [Goy11] and Lin and Pass [LP11]. Other constructions include [PR05a, LP09, LPV08, PPV08, PW10, Wee10, GLOV12]. The current state of art is represented by a construction of Goyal, Richelson, Rosen and Vald [GRRV14] whose scheme requires only four rounds of interaction. On the negative side, Pass [Pas13] showed that two-round non-malleable commitments cannot exists w.r.t. black-box proofs of security based on any "standard" intractability assumption. The lower bound of Pass holds even if the construction uses the underlying assumption in a non-black-box way. Thus, the main remaining open problem in the study of round complexity of non-malleable commitments is

Do there exist non-malleable commitments with only three rounds of interaction?

Zero-Knowldge: A Barrier to Three-Round Non-Malleable Commitment. Almost all previous schemes invoke some sort of proof of consistency. Such proofs are usually critical to the proofs of non-malleability, as without consistency one runs into a host of "selective $\perp$ attacks" (where the MIM plays in such a way so that whether or not his commitment is valid, depends on the value inside the commitment he receives) which are difficult to rule out. Generally, zero-knowledge is used for this purpose; for example the recent works of [GRRV14, BGR$^+$15] use the Feige-Shamir paradigm in order to parallelize their proof of consistency down to four rounds. However, zero-knowledge w.r.t. black-box simulation is known to require 4-rounds [GK96], so if one hopes to obtain three-round non-malleable commitment, one must do so without zero-knowledge. Furthermore, zero-knowledge is computationally expensive and any non-malleable commitment which contains zero-knowledge as a subprotocol is libel to be significantly slower when compared to the computation required for an ordinary statically binding commitment. This is indeed true of the recent protocols of [GRRV14, BGR$^+$15].

Our Contributions. We present a new construction of non-malleable commitments which has the following features:

- The protocol has only *three rounds* of interaction. Pass [Pas13] showed that two-round non-malleable commitments unfortunately cannot exists w.r.t. black-box proofs of security based on any "standard" intractability assumption. The lower bound of Pass holds even if the construction uses the underlying assumption in a non-black-box way. Thus, this resolves the round complexity of non-malleable commitments at least w.r.t.

black-box security reductions. Our construction is secure as per the standard notion of non-malleability w.r.t. commitment.

- Our protocol simple and *truly efficient*. In our basic protocol, the entire computation of the committer is dominated by just three invocations of a non-interactive statically binding commitment scheme, while, the receiver computation (in the commitment stage) is limited to just sampling a random string. The decommitment stage is equally basic: the committer would send the openings of these commitments, while, the receiver would be required to check these openings for correctness and perform some simple computations. The protocol is simple to describe, the main complexity lies in the analysis rather than the construction.

  In several previous works (including [GRRV14, BGR$^+$15]), first a non-malleable commitment scheme for "small" tags is constructed. Then, a scheme for large tags is obtained by using non-malleability amplification [DDN91, LP09]. This adds a significant multiplicative overhead to the computation of each party: the multiplicative overhead is typically related to the number of bits in the large tags. Unlike these previous works, our basic protocol works directly with large tags, and hence, we avoid any expensive non-amplification steps. Our basic protocol provides security only against synchronizing adversaries. Extension to non-synchronizing adversaries is addressed later, still with a three round protocol.

- Our protocol is based on a black-box use of any non-interactive statistically binding commitment scheme. Such schemes, in turn, can be based on any one-way permutation, or, at the cost of an extra initialization round, any one-way function. Previously, the best known black-box construction of non-malleable commitments required a larger constant number of rounds [GLOV12, KMO14]. Furthermore, the previous constructions, even though black-box, were significantly less efficient [GLOV12, LP12, KMO14, Kiy14]. For example, the construction of Goyal et al [GLOV12] used "MPC in the head techniques" of Ishai et. al [IKOS07].

- Our construction is public-coin and makes use of only black-box simulation. Prior to our work, no public-coin constant round non-malleable commitment schemes were known based on black-box simulation. The structure of our basic protocol is arguably "as basic as it can be": the sender sends a single commitment to some string, the receiver sends a random challenge, and, in the final round, sender sends another string (but doesn't send any opening).

Key Technical Idea – Using Split-State Non-Malleable Codes. Our key technical tool will be non-malleable codes in the split state model [DPW10, DKO13]. Very informally, non-malleable codes in the split state model allow one to encode a message m into $\mathsf{Enc}(m) = (\mathsf{L}, \mathsf{R})$ and be assured that if an adversary uses functions (f, g) to tamper (L, R) into $(\tilde{\mathsf{L}}, \tilde{\mathsf{R}}) = \big(f(\mathsf{L}), g(\mathsf{R})\big)$, the decoded value $\tilde{m} = \mathsf{Dec}(\tilde{\mathsf{L}}, \tilde{\mathsf{R}})$ will either be equal to m (in, for example, the case when $f = g = \mathbb{I}$ are the identity function), or will be independent from m.

As a first attempt towards constructing non-malleable commitments, what if the sender separately commits to L and R? This is indeed a non-starter since the underlying commitment scheme may have some homomorphic properties allowing the receiver to maul L and R "jointly". Our starting idea is as follows. Let us focus our attention to synchronizing adversaries[1]. The committer encodes the message m as L and R, and, in the first round, the committer sends a commitment $\mathrm{com}(L)$ to L. The receiver responds back with a simple acknowledgement message. Finally, the committer sends R *in clear*. Very roughly, the scheme does seem to have some non-malleability features. In the first round, the MIM could maul $\mathrm{com}(L)$ into $\mathrm{com}(\tilde{L})$ *without the knowledge of* R. In the final round, the MIM receives R and is required to produce $\tilde{R}$. It seems that this mauling must be done independent of L since only $\mathrm{com}(L)$ is available to MIM (rather than L itself). While this is indeed our starting point, this intuition turns out to be not sound and more work will be required. Our basic protocol is quite simple and is given below.

- **Committer's Input:** A value v to commit to.

- **1.** $\mathcal{C} \to \mathcal{R}$**:** $\mathcal{C}$ chooses $(\mathsf{L}, \mathsf{R}) \leftarrow \mathsf{Enc}(v)$ where L is viewed as a field element in $\mathbb{Z}_q$; $\mathcal{C}$ also draws $r \leftarrow \mathbb{Z}_q$ at random and sends $\mathrm{Com}(\mathsf{L}), \mathrm{Com}(r)$ to $\mathcal{R}$ where Com is Blum's non-interactive, perfectly binding commitment scheme.

[1] Roughly, this means that the MIM sends the i-th round message on the right immediately after getting the i-th round message in the left interaction.

- **2.** $\mathcal{R} \rightarrow \mathcal{C}$**:** $\mathcal{R}$ chooses a random $\alpha \leftarrow \mathbb{Z}_q$ and sends it to $\mathcal{C}$.
- **3.** $\mathcal{C} \rightarrow \mathcal{R}$**:** $\mathcal{C}$ sends $a = r\alpha + \mathsf{L}$ and R to $\mathcal{R}$.
- **Decommitment:** To decommit, $\mathcal{C}$ decommits to both commitments in **1.**

To prove security of this protocol, the key challenge would be to reduce any "successful" mauling attack on our protocol to a mauling attack on the underlying non-malleable code. Our adversary for the non-malleable codes would have to work as follows. At a very high level, the adversary would run the left execution (with the MIM) using the given L and R. From the right execution, it would somehow extract $\tilde{L}$ and $\tilde{R}$. If MIM was successful in mauling our non-malleable commitment protocol, this guarantees that the extracted $\tilde{L}$ and $\tilde{R}$ decodes to a message $\tilde{m}$ which is related to the message m represented by L and R. This would presumably contradict the security of the non-malleable code. However one must keep in mind that the adversary we build for non-malleable codes in the split state model does not get to see L and R at once. Hence, it can't simply run our non-malleable commitment protocol, and, extract the tampered $\tilde{L}$ and $\tilde{R}$.

To complete the proof of security, we would need to construct split state functions f and g (which can use the MIM and the distinguisher for our protocol internally). $f(L)$ and $g(R)$ would have to output $\tilde{L}$ and $\tilde{R}$ respectively. However note that neither f nor g can complete the protocol execution with MIM to extract the tampered code (since they will be missing either L or R)!

Thus, the idea of reducing the security of our construction to the security of non-malleable codes (in the split state model) seems like a non-starter. *This would be the key technical challenge we encounter in our proof of non-malleability.*

Our proof strategy, at a very high level would be execute $f(L)$ and $g(R)$ independent as required, and output $\tilde{L}$ and $\tilde{R}$ respectively. This would lead to $\tilde{L}$ and $\tilde{R}$ being extracted from two different protocol transcripts. However, then we show that there would exist a single protocol transcript (from the correct distribution) such that the left execution in that transcript was completed using L and R, and, the MIM completes the right execution using $\tilde{L}$ and $\tilde{R}$. Thus, if MIM was successful in mauling our protocol, f and g were successful in mauling the non-malleable code in the split state model.

Indeed, we are not able to make the above argument go through based on standard split-state non-malleable codes. We need need the following additional properties described below very informally. A more formal description can be found in Section 3:

1. The code must be an *augmented* split state non-malleable code [AAG$^+$16]. Very informally, this means that the distinguisher for the non-malleable code is also given R as input (in addition to the decoded message $\tilde{m}$).
2. We need the non-malleable code to be *conditionally* secure as well. This is a new property we define in our paper. This property in particular implies that the non-malleable code is secure even under leakage of logarithmic many bits from L (but the exact definition is different). The details of this property are given in section 3.
3. The code must satisfy what we call as the *simulatable right state* property. Informally this means the following. For a particular L, sample $R_1, R_2, \ldots, R_n$ s.t. (L, R_i) decodes to the same message m for all i. Then even given $R_1, R_2, \ldots, R_n$ (but not L), an adversary gets no advantage in guessing m.

Our construction is based on the recent split state non-malleable code of Aggarwal et. al [ADL14]. It turns out that the code in [ADL14] already satisfies the first two of the above properties. We then present a modification to add the strong hiding property. Note that the code of Agarwal et. al is purely information theoretic, and, in comparison to cryptographic objects (such as commitments or one-way functions), very efficient. Encoding and decoding simply requires sampling random vectors and taking their inner product, etc. To add the strong hiding property, we add a commitment and a symmetric encryption to the encoding and decoding procedures. Thus, our overall basic protocol has computation which is dominated by about three invocations of a statistically binding commitment scheme (or rather two invocation of a statistically binding commitment and one symmetric encryption to be more precise)[2].

[2]We note that the two commitments which the committer sends to the receiver in our protocol in the first round can, in fact, be combined in a single commitment

Extension to non-synchronizing adversaries. While in some applications, security against synchronizing adversaries is all one needs (e.g., constructing round efficient multi-party computation), in others, non-malleability against arbitrary schedulings is required. Our basic protocol only provides security again synchronizing adversaries, and since our protocol has only three rounds, it is easy to see that the only other potentially problematic scheduling is the sequential one where the left execution finishes entirely even before the right execution starts.

To extend to non-synchronizing adversaries, we make our protocol extractable by running a 3-round extractable (malleable) commitment scheme in parallel to our basic protocol (without adding any "proofs of consistency" of the two parallel executions). Extraction immediately yields non-malleability against a sequential adversary as we may rewind M and extract his commitment without having to rewind the honest committer. The main technical challenge for this portion is proving non-malleability against a synchronizing adversary; *i.e.*, that the extractable commitment doesn't "interfere" with the basic protocol. To achieve this, we will use an extractable commitment scheme such that extracting from this scheme requires a larger number of rewindings compared to our basic protocol. This is inspired by a technique from the constant round protocol of Lin and Pass [LP11]. Our final protocol, however, requires significantly more invocations of the underlying commitment scheme, however we stress that even the extended protocol is significantly more efficient than any of the prior ones known in the literature (besides requiring only 3-rounds of interaction). This protocol is described in more detail in Section E.

2 Preliminaries

We use λ for the security parameter and $\mathsf{negl}(\lambda)$ or negl for a function which tends to zero faster than λ^{-k} for any constant k. We have moved most of our preliminaries to the appendix, so see Appendix 2 for notes on the basic cryptographic building blocks we will use as well as basic definitions relating to non-malleable commitment and non-malleable codes.

3 New Constructions of Non-Malleable Codes

3.1 Conditional Augmented Non-Malleable Codes

Let $(\mathsf{Enc}, \mathsf{Dec})$ be a split-state code with codeword space $\mathcal{L} \times \mathcal{R}$. In proving that our commitment scheme is non-malleable, we will need to choose a random $\mathsf{L} \in \mathcal{L}$ and be ensured that the augmented tampering distribution is independent of m even conditioned on L. We define information theoretic and computational variants of these codes.

Definition 1 (Conditional Augmented Tampering Distribution). *Fix $m \in \mathcal{M}$, $(f, g) \in \mathcal{F}_{\mathsf{split}}$, and $\mathsf{L} \in \mathcal{L}$. The* conditional augmented tampering distribution*, $\mathsf{V}^{\mathsf{L}}_{m,f,g}$ is defined by the following process: draw $\mathsf{R} \leftarrow \mathsf{Enc}(m|\mathsf{L}) = \{\mathsf{R}' \in \mathcal{R} : \mathsf{Dec}(\mathsf{L}, \mathsf{R}') = m\}$, set $(\tilde{\mathsf{L}}, \tilde{\mathsf{R}}) = \big(f(\mathsf{L}), g(\mathsf{R})\big)$ and output $(\mathsf{R}, \tilde{m})$ where $\tilde{m} = \mathsf{Dec}(\tilde{\mathsf{L}}, \tilde{\mathsf{R}})$.*

Definition 2 (Conditional Augmented Simulatable Distribution). *Let $\left\{\mathsf{D}^{\mathsf{L}}_m\right\}_{m,\mathsf{L}}$ be a family of distributions on $\mathcal{R} \times \mathcal{M}$ indexed by $m \in \mathcal{M}$ and $\mathsf{L} \in \mathcal{L}$, where $\mathcal{L}$ and $\mathcal{R}$ are arbitrary sets. We say that $\left\{\mathsf{D}^{\mathsf{L}}_m\right\}_{m,\mathsf{L}}$ is* $\varepsilon-$conditionally augmented simulatable *if there exists a family of distributions $\left\{\mathsf{S}^{\mathsf{L}}\right\}_{\mathsf{L}}$ on $\mathcal{R} \times \big(\mathcal{M} \cup \{\mathsf{same}\}\big)$ which such that*

$$\Pr_{\mathsf{L}}\Big[\Delta\big(\mathsf{D}^{\mathsf{L}}_m, \mathsf{S}^{\mathsf{L}}_m\big) < \varepsilon \;\forall\, m \in \mathcal{M}\Big] \geq 1 - \varepsilon,$$

where the probability is over $\mathsf{L} \leftarrow \mathcal{L}$ drawn uniformly and where $\mathsf{S}^{\mathsf{L}}_m$ draws $(\mathsf{R}, \tilde{m}) \leftarrow \mathsf{S}^{\mathsf{L}}$ and outputs (R, m) if $\tilde{m} = \mathsf{same}$, $(\mathsf{R}, \tilde{m})$ if not. We say $\{\mathsf{D}_m\}_m$ is computationally conditionally augmented simulatable *if for all PPT distinguishers D and non-negligible $\delta > 0$,*

$$\Pr_{\mathsf{L}}\left[\exists\, m \in \mathcal{M} \text{ st } \left|\Pr_{(\mathsf{R},\tilde{m}) \leftarrow \mathsf{D}^{\mathsf{L}}_m}\big(\mathrm{D}(\mathsf{R}, \tilde{m}) = 1\big) - \Pr_{(\mathsf{R},\tilde{m}) \leftarrow \mathsf{S}^{\mathsf{L}}_m}\big(\mathrm{D}(\mathsf{R}, \tilde{m}) = 1\big)\right| > \delta\right] = \mathsf{negl}.$$

Definition 3 (Conditional Augmented Non-Malleable Code). *We say that $(\mathsf{Enc}, \mathsf{Dec})$ is $\varepsilon-$conditionally augmented non-malleable (resp. computationally conditionally augmented non-malleable) against $\mathcal{F}_{\mathsf{split}}$ if for all $(f, g) \in \mathcal{F}_{\mathsf{split}}$, $\left\{\mathsf{V}^{\mathsf{L}}_{m,f,g}\right\}_{m,\mathsf{L}}$ is $\varepsilon-$conditionally augmented simulatable (resp. computationally conditionally augmented simulatable).*

The following claim is implicit in the recent work of [AAG+16, Agg] which builds on [ADL14]. The proof is given in Appendix B.

Claim 1. *The code* $(\mathsf{Enc}, \mathsf{Dec})$ *of [ADL14] is* $\varepsilon'-$*conditionally augmented non-malleable for some negligible quantity* $\varepsilon' = \varepsilon'(\lambda) > 0$.

3.2 Adding the Hiding Property

We will also need our non-malleable code to have a computational hiding property resembling semantic security in order to rule out selective $\perp$ attacks. Essentially we need there to exist a distribution $\mathcal{D}_{\mathsf{hid}}$ on $\mathcal{R}$ such that for almost all $\mathsf{L} \in \mathcal{L}$, $\mathcal{D}_{\mathsf{hid}}$ is indistinguishable from $\mathsf{Enc}(m|\mathsf{L})$ for all $m \in \mathcal{M}$. We formalize this using a game between a challenger $\mathcal{C}$ and a PPT adversary $\mathcal{A}$, parametrized by $N = \mathsf{poly}(\lambda)$, a message $m \in \mathcal{M}$ and a distribution $\mathcal{D}_{\mathsf{hid}}$ on $\mathcal{R}$.

- $\mathcal{C}$ draws $\mathsf{L} \leftarrow \mathcal{L}$ and $b \leftarrow \{0,1\}$. $\mathcal{C}$ sends $\mathsf{R}_1, \ldots, \mathsf{R}_N$ to $\mathcal{A}$ where $\mathsf{R}_i \leftarrow \mathsf{Enc}(m|\mathsf{L})$ if $b = 0$ and $\mathsf{R}_i \leftarrow \mathcal{D}_{\mathsf{hid}}$ if $b = 1$.
- $\mathcal{A}$ outputs b' and wins if $b' = b$.

Definition 4 (Codes with Simulatable State). *We say that a split-state code* $(\mathsf{Enc}, \mathsf{Dec})$ *has a* simulatable right state *if there is a distribution* $\mathcal{D}_{\mathsf{hid}}$ *on* $\mathcal{R}$ *such that for all PPT* $\mathcal{A}$, $N = \mathsf{poly}(\lambda)$, *and* $m \in \mathcal{M}$, *the probability that* $\mathcal{A}$ *wins the above game is at most* $1/2 + \mathsf{negl}$.

Construction. Let $(\mathsf{Enc}_0, \mathsf{Dec}_0)$ be an $\varepsilon-$conditional augmented non-malleable code. Let $(\mathsf{G}, \mathsf{E}, \mathsf{D})$ be a symmetric key encryption scheme, and let $(\mathrm{Com}, \mathrm{Decom})$ be Blum's non-interactive, perfectly binding commitment scheme. The new coding scheme, $(\mathsf{Enc}, \mathsf{Dec})$ is defined as follows.

- $\mathsf{Enc}(m)$**:** Draw $(\mathsf{L}_0, \mathsf{R}_0) \leftarrow \mathsf{Enc}_0(m)$, $k \leftarrow \mathsf{G}(1^\lambda)$, $\sigma \leftarrow \$$, and $c \leftarrow \mathsf{E}_k(\mathsf{R}_0)$. Set $z = \mathrm{Com}(k, \sigma)$ and output (L, R) where $\mathsf{L} = \big(\mathsf{L}_0, (k, \sigma)\big)$ and $\mathsf{R} = (c, z)$.
- $\mathsf{Dec}(\mathsf{L}, \mathsf{R})$**:** If either $\mathsf{L}, \mathsf{R} = \perp_{\mathsf{com}}$ output $\perp_{\mathsf{com}}$. Otherwise, parse $\mathsf{L} = \big(\mathsf{L}_0, (k, \sigma)\big)$ and $\mathsf{R} = (c, z)$, check that $\mathrm{Decom}(z) = (k, \sigma)$. If so set $\mathsf{R}_0 = \mathsf{D}_k(c)$, output $\mathsf{Dec}_0(\mathsf{L}_0, \mathsf{R}_0)$, if not output $\perp_{\mathsf{com}}$.

Claim 2. $(\mathsf{Enc}, \mathsf{Dec})$ *has simulatable right state.*

Proof. Define three challengers $\mathcal{C}_0$, $\mathcal{C}_1$, $\mathcal{C}_2$ for the above hiding game, played with some fixed $m \in \mathcal{M}$. Each challenger draws $\mathsf{L} = \big(\mathsf{L}_0, (k, \sigma)\big) \leftarrow \mathcal{L}$ and for $i = 1, \ldots, N$

- $\mathcal{C}_0$ sets $\mathsf{R}_i = (c_i, z)$ where $z = \mathrm{Com}(k, \sigma)$ and $c_i \leftarrow \mathsf{E}_k(\mathsf{R}_0)$ for some R_0 with $\mathsf{Dec}_0(\mathsf{L}_0, \mathsf{R}_0) = m$.
- $\mathcal{C}_1$ sets $\mathsf{R}_i = (c_i, z)$ where $z = \mathrm{Com}(0)$ and $c_i \leftarrow \mathsf{E}_k(\mathsf{R}_0)$.
- $\mathcal{C}_2$ sets $\mathsf{R}_i = (c_i, z)$ where $z = \mathrm{Com}(0)$ and $c_i \leftarrow \$$ is random.

Note that $\mathcal{C}_0$ draws each R_i from $\mathsf{Enc}(m|\mathsf{L})$, whereas $\mathcal{C}_2$ draws R_i from $\mathcal{D}_{\mathsf{hid}}$. Furthermore, $\mathcal{A}$ cannot distinguish between his interaction with $\mathcal{C}_0$ and $\mathcal{C}_1$ by the hiding of Com. Likewise, he cannot distinguish between his interaction with $\mathcal{C}_1$ and $\mathcal{C}_2$ by the random ciphertext property of $(\mathsf{G}, \mathsf{E}, \mathsf{D})$. □

Hiding Game Variant. We will use another game parametrized by polynomials $N, N' = \mathsf{poly}(\lambda)$. In this game $\mathcal{A}$ sends $\mathcal{C}$ two messages $m, m' \in \mathcal{M}$, $\mathcal{C}$ draws a secret $\mathsf{L} \leftarrow \mathcal{L}$ and sends $\mathcal{A}$ the tuple $\big(\mathsf{R}, \{\mathsf{R}_1\}, \ldots, \{\mathsf{R}_N\}\big)$ where $\mathsf{R} \leftarrow \mathsf{Enc}(m|\mathsf{L})$ and each set has N' elements. Moreover, $\mathsf{R}_i \leftarrow \mathsf{Enc}(m'|\mathsf{L})$ for all $\mathsf{R}_i \in \{\mathsf{R}_i\}$ and all $i = 1, \ldots, N$ except for one random i^* for which $\mathsf{R}_{i^*} \leftarrow \mathsf{Enc}(m|\mathsf{L})$ for all $\mathsf{R}_{i^*} \in \{\mathsf{R}_{i^*}\}$; $\mathcal{A}$ tries to guess i^*. If $(\mathsf{Enc}, \mathsf{Dec})$ has simulatable right state then it is straightforward to show that any PPT adversary $\mathcal{A}$ can guess i^* with probability at most $1/N + \mathsf{negl}$.

Lemma 1. *If* $(\mathsf{G}, \mathsf{E}, \mathsf{D})$ *is a semantically secure private key encryption scheme with the random ciphertext property,* Com *is a perfectly binding non-interactive commitment scheme and* $(\mathsf{Enc}_0, \mathsf{Dec}_0)$ *is* $\varepsilon-$*conditionally augmented non-malleable against* $\mathcal{F}_{\mathsf{split}}$ *for negligible* $\varepsilon > 0$, *then* $(\mathsf{Enc}, \mathsf{Dec})$ *is computationally conditionally augmented non-malleable against* $\mathcal{F}^{\mathsf{poly}}_{\mathsf{split}}$.

Proof Idea. Fix $(f, g) \in \mathcal{F}_{\mathsf{split}}^{\mathsf{poly}}$. We must describe a family of simulators for $\{\mathsf{V}_{m,f,g}^{\mathsf{L}}\}_{m,\mathsf{L}}$. $\mathsf{S}_{f,g}^{\mathsf{L}}$ will be one of two distributions depending on L. The first simply draws $\mathsf{R} \leftarrow \mathcal{D}_{\mathsf{hid}}$ and outputs $(\mathsf{R}, \perp_{\mathsf{com}})$. This will simulate $\{\mathsf{V}_{m,f,g}^{\mathsf{L}}\}_m$ whenever $\mathsf{p}_{\mathsf{L},m}$ is small for all $m \in \mathcal{M}$ where $\mathsf{p}_{\mathsf{L},m}$ is shorthand for $\Pr_{(\mathsf{R},\tilde{m})\leftarrow\mathsf{V}_{m,f,g}^{\mathsf{L}}}\big(\tilde{m} \neq \perp_{\mathsf{com}}\big)$. If there exists $m \in \mathcal{M}$ it can be shown that $\mathsf{p}_{\mathsf{L},m}$ for all $m \in \mathcal{M}$ (Claim 8). This means that $\mathsf{V}_{m,f,g}^{\mathsf{L}}$ can be related to a conditional augmented tampering distribution $\mathsf{V}_{m,f_0,g_0}^{\mathsf{L}_0}$ on $(\mathsf{Enc}_0, \mathsf{Dec}_0)$ for tampering functions $(f_0, g_0) \in \mathcal{F}_{\mathsf{split}}$ related to (f, g). We use $\mathsf{S}_{f_0,g_0}^{\mathsf{L}_0}$ to construct $\mathsf{S}_{f,g}^{\mathsf{L}}$. The formal proof is in Appendix C, we describe the functions (f_0, g_0) here.

- **Random Choices:** Draw $k \leftarrow \mathsf{G}(1^\lambda)$, $\sigma \leftarrow \$$, set $z = \mathrm{Com}(k, \sigma)$, and draw $c^\$ \leftarrow \mathsf{E}_k(0)$. Set $(\cdot, \tilde{z}^\$) = g(c^\$, z)$. Save the values $(k, \sigma, z, \tilde{z}^\$)$.
- $f_0(\mathsf{L}_0)$**:** Compute $\big(\tilde{\mathsf{L}}_0, (\tilde{k}, \tilde{\sigma})\big) = f\big(\mathsf{L}_0, (k, \sigma)\big)$. Output $\tilde{\mathsf{L}}_0$.
- $g_0(\mathsf{R}_0)$**:** Draw $c \leftarrow \mathsf{E}_k(\mathsf{R}_0)$ and set $(\tilde{c}, \tilde{z}) = g(c, z)$. If $\tilde{z} \neq \tilde{z}^\$$, output $\perp_{\mathsf{com}}$. Otherwise, use superpolynomial time to break open $\tilde{z}$ and recover the pair $(\tilde{k}', \tilde{\sigma}')$. Output $\tilde{\mathsf{R}}_0 = \mathsf{D}_{\tilde{k}'}(\tilde{c})$.

The main observation is that whenever $(\mathsf{L}, c^\$)$ are such that

$$\mathrm{Decom}(\tilde{z}^\$) = (\tilde{k}, \tilde{\sigma}),$$

where $\big(\cdot, (\tilde{k}, \tilde{\sigma})\big) = f(\mathsf{L})$ and $(\cdot, \tilde{z}^\$) = g(\mathsf{R})$, the conditional augmented tampering distributions $\mathsf{V}_{m,f,g}^{\mathsf{L}}$ and $\mathsf{V}_{m,f_0,g_0}^{\mathsf{L}_0}$ are the same (up to encrypting the right state output by $\mathsf{V}_{m,f_0,g_0}^{\mathsf{L}_0}$). In this case, we can simulate $\mathsf{V}_{m,f,g}^{\mathsf{L}}$ using $\mathsf{S}_{f_0,g_0}^{\mathsf{L}_0}$, so it suffices to show that $c^\$$ exists such that $\mathrm{Decom}(\tilde{z}^\$) = (\tilde{k}, \tilde{\sigma})$. This follows from the semantic security of $(\mathsf{G}, \mathsf{E}, \mathsf{D})$ as $(c^\$, z)$ and $\mathsf{R} \leftarrow \mathsf{Enc}(m|\mathsf{L})$ differ only in their encrypted value, and $\mathsf{p}_{\mathsf{L},m}$ is not too small.

4 The Basic Protocol

The protocol is shown in Figure 1.

Setup: Let Com be a non-interactive, perfectly binding commitment scheme. Let $(\mathsf{Enc}, \mathsf{Dec})$ be a conditional non-malleable code with indistinguishable views. Fix a large prime q. Let $\mathsf{id} \in \{0,1\}^\lambda$ be C's identity.

Commiter's Private Input: $v \in \mathcal{M}_{\langle \mathrm{C},\mathrm{R}\rangle}$ to be committed to.

Commit Phase:

1. $\mathrm{C} \to \mathrm{R}$**:** Set $m = v \circ \mathsf{id}$ and draw $(\mathsf{L}, \mathsf{R}) \leftarrow \mathsf{Enc}(m)$, where $\mathsf{L} \in \mathcal{L} \subset \mathbb{Z}_q$. Choose random $r \in \mathbb{Z}_q$ and send $\mathrm{Com}(\mathsf{L})$, $\mathrm{Com}(r)$ to R.
2. $\mathrm{R} \to \mathrm{C}$**:** Send random challenge $\alpha \in \mathbb{Z}_q$.
3. $\mathrm{C} \to \mathrm{R}$**:** Send response $a = r\alpha + \mathsf{L} \in \mathbb{Z}_q$ and also send R.

Decommit Phase:

1. $\mathrm{C} \to \mathrm{R}$**:** Open the commitments sent in step 1. Let $\mathsf{L}', r' \in \mathbb{Z}_q$ be the decommitted values.

Receiver's Output: If L' and r' do not satisfy $r'\alpha + \mathsf{L}' = a$ then output the special symbol $\perp_{\mathsf{inc}}$. Otherwise, compute $m' = \mathsf{Dec}(\mathsf{L}', \mathsf{R})$ and parse $m' = v' \circ \mathsf{id}'$. Output v' if $\mathsf{id}' = \mathsf{id}$, $\perp_{\mathsf{id}}$ if not.

Figure 1: Non-malleable commitment scheme $\langle \mathrm{C}, \mathrm{R}\rangle$.

Claim 3. $\langle \mathrm{C}, \mathrm{R} \rangle$ *is a perfectly binding commitment scheme.*

Proof Sketch. Perfect binding follows immediately from the perfect binding of Com. Computational hiding follows in a straightforward fashion from the hiding of Com and the well known fact that any split-state non-malleable code is also a 2-out-of-2 secret sharing scheme. □

Theorem 1. $\langle \mathrm{C}, \mathrm{R} \rangle$ *is non-malleable against a synchronizing adversary.*

5 Proof of Non-Malleability (Theorem 1)

5.1 Notation and Proof Overview

Transcripts. Suppose a PPT man-in-the-middle, M, participates in two protocol executions. We denote the transcript of M's view with the letter $\mathbb{T}$. So

$$\mathbb{T} = \big(\mathsf{id}, \tilde{\mathsf{id}}, \mathrm{Com}(\mathsf{L}), \mathrm{Com}(r), \mathrm{Com}(\tilde{\mathsf{L}}), \mathrm{Com}(\tilde{r}), \tilde{\alpha}, \alpha, a, \mathsf{R}, \tilde{a}, \tilde{\mathsf{R}}\big).$$

We write $\mathbf{Com}_{\mathsf{L}}$ as shorthand for $\mathrm{Com}(\mathsf{L})$ and $\mathrm{Com}(r)$. Note that $\mathbf{Com}_{\mathsf{L}}$ specifies a linear polynomial $\varphi(x) = rx + \mathsf{L}$ which C uses to answer M's query α. We will usually write a transcript $\mathbb{T}$ more consicely as $\mathbb{T} = \big(\mathbf{Com}_{\mathsf{L}}, \tilde{\alpha}, a, \mathsf{R}\big)$, surpressing C's identity id and the quantities which are outputs of M. Since without loss of generality M is deterministic, these values uniquely define a full transcript.

Partial Transcripts. We also will find it useful to speak of partial transcripts, as this will let us isolate certain random choices made during the execution of $\mathbb{T}$. We use τ to denote the partial transcript where C's value R is unspecified. We write τ consicely as $\tau = (\mathbf{Com}_{\mathsf{L}}, \tilde{\alpha}, a)$. Note that M's third message is not specified given τ, however τ extends to a full transcript $\mathbb{T}$ once R is chosen. We write this full transcript $\mathbb{T}(\tau, \mathsf{R})$.

Proof Overview. Before continuing, we go though the proof at a high level. As mentioned in the intro, our main idea is to use an M who mauls $\langle \mathrm{C}, \mathrm{R} \rangle$ to construct split state tampering functions (f, g) which maul the code $(\mathsf{Enc}, \mathsf{Dec})$. Our f and g are given a partial transcript $\tau = \big(\mathbf{Com}_{\mathsf{L}}, \tilde{\alpha}, a\big)$ and additionally they share a random $\mathsf{R}^{\$} \leftarrow \mathsf{Enc}(m^* | \mathsf{L})$ for some arbitrary fixed $m^* \in \mathcal{M}$. This defines a full transcript $\mathbb{T} = \mathbb{T}(\tau, \mathsf{R}_{\$})$. If M's third message $\tilde{a}_{\$}$ in $\mathbb{T}$ is correct, then f will be able to extract $\tilde{\mathsf{L}}$ by rewinding M and asking a new challenge $\tilde{\beta}$, answering honestly on the left using L and the point (α, a) from τ and $\mathsf{R}_{\$}$. $g(\mathsf{R})$ simply outputs $\tilde{\mathsf{R}}$ from M's third message in the transcript $\mathbb{T}(\tau, \mathsf{R})$.

Two main issues with the above must be overcome. The first is that we must argue that M's answer to $\tilde{\alpha}$ is the same when he is given $(a, \mathsf{R}^{\$})$ on the right and when he is given (a, R). There is no reason necessarily for this to be the case. However, by the simulatability of the right state of $(\mathsf{Enc}, \mathsf{Dec})$, we are able to show that M cannot decide whether to answer $\tilde{\alpha}$ correctly or not in a way which depends on C's commitment on the left. This aspect of the proof is a series of reductions to the hiding game of $(\mathsf{Enc}, \mathsf{Dec})$ which we call the hiding machine. We used such arguments also in order to construct $(\mathsf{Enc}, \mathsf{Dec})$. For space reasons we have put this portion of our proof in Appendix D. We are left with the absolute cases where M is either committing to $\perp_{\mathsf{inc}}$ on the right with high probability regardless of the left commitment, in which case he is certainly not mauling, or he will answer correctly with non-negligible probability even if he receives $\mathsf{R}_{\$}$ on the left instead of R, in which case $f(\mathsf{L})$ will succeed in extracting $\tilde{\mathsf{L}}$.

The second issue is the dependence of τ on L. Certainly (f, g) are not split-state as g depends on τ which contains information about L. However, τ contains only a commitment to L and so does not depend computationally on L. Since f and g are polytime, we are able to argue that their output is indistinguishable from truly split-state functions which share a bogus partial transcript which contains no information about L. One subtlety here is that originally, L was generated along with (f, g). This is in contrast to the usual situation in non-malleable codes where the tampering functions are fixed and then (L, R) are sampled. This is why we need our non-malleable code to have the conditional property where non-malleability is certain to hold even conditioned on an L which was chosen before. We now continue with the notation and proceed to the proof.

The Distribution M_m^τ and Distinguisher D^τ. Our goal is to use a MIM who breaks the non-malleability of $\langle \mathrm{C}, \mathrm{R} \rangle$ to violate the security of the code $(\mathsf{Enc}, \mathsf{Dec})$. In order to do this we make some notational changes which make it easier to relate M to the non-malleability game of the code. By definition, if M breaks the non-malleability of $\langle \mathrm{C}, \mathrm{R} \rangle$ then there exists $v \in \mathcal{M}_{\langle \mathrm{C},\mathrm{R} \rangle}$, a PPT distinguisher D, and non-negligible $\delta = \delta(\lambda) > 0$ such that

$$\left| \Pr_{(\mathbb{T},\tilde{v}) \leftarrow \mathbf{MIM}_v} \big(\mathrm{D}(\mathbb{T}, \tilde{v}) = 1 \big) - \Pr_{(\mathbb{T},\tilde{v}) \leftarrow \mathbf{MIM}_0} \big(\mathrm{D}(\mathbb{T}, \tilde{v}) = 1 \big) \right| = \delta. \tag{1}$$

Let $m = v \circ \mathsf{id}$, $m' = 0 \circ \mathsf{id}$ and $\tilde{m} = \tilde{v} \circ \tilde{\mathsf{id}}$. So $m, m' \in \mathcal{M}$ are the messages encoded during the left executions of $\langle \mathrm{C}, \mathrm{R} \rangle$ in the real/ideal world, and $\tilde{m} \in \mathcal{M}$ is the message encoded on the right. For a given partial transcript $\tau = \big(\mathbf{Com}_{\mathsf{L}}, \tilde{\alpha}, a \big)$, let M_m^τ be the distribution which draws $\mathsf{R} \leftarrow \mathsf{Enc}(m | \mathsf{L})$ and outputs $(\mathsf{R}, \tilde{m})$, where $\tilde{m}$ is M's encoded message in $\mathbb{T}(\tau, \mathsf{R})$. Let D^τ be the PPT distinguisher which on input $(\mathsf{R}, \tilde{m})$, sets $\mathbb{T} = \mathbb{T}(\tau, \mathsf{R})$, parses $\tilde{m} = \tilde{v} \circ \tilde{\mathsf{id}}'$ and outputs $\mathrm{D}(\mathbb{T}, \tilde{v})$ if $\tilde{\mathsf{id}}' = \tilde{\mathsf{id}}$ (recall $\tilde{\mathsf{id}}$ is part of τ), $\mathrm{D}(\mathbb{T}, \perp_{\mathsf{id}})$ otherwise. With these notational changes in place, (1) gives

$$\Pr_\tau \left[\left| \Pr\Big(\mathrm{D}^\tau\big(\mathsf{M}_m^\tau \big) = 1 \Big) - \Pr\Big(\mathrm{D}^\tau\big(\mathsf{M}_{m'}^\tau \big) = 1 \Big) \right| \geq \frac{\delta}{2} \right] \geq \frac{\delta}{2}. \tag{2}$$

Notice also that since M is required to produce a commitment using a tag $\tilde{\mathsf{id}} \neq \mathsf{id}$, when $(\mathsf{R}, \tilde{m})$ is drawn from M_m^τ or $\mathsf{M}_{m'}^\tau$, we will always have $\tilde{m} \notin \{m, m'\}$.

Definition 5 (Mauling-Friendly Partial Transcripts). *For $m, m' \in \mathcal{M}$, write $\tau \in \mathsf{MAUL}_{m,m'}$ if $\left| \Pr\Big(\mathrm{D}^\tau\big(\mathsf{M}_m^\tau \big) = 1 \Big) - \Pr\Big(\mathrm{D}^\tau\big(\mathsf{M}_{m'}^\tau \big) = 1 \Big) \right| \geq \delta/2$. In this case we say τ is* mauling-friendly.

So, moving forward, if M breaks the non-malleability of $\langle \mathrm{C}, \mathrm{R} \rangle$ then there exist $m, m' \in \mathcal{M}$ such that $\Pr_\tau \big[\tau \in \mathsf{MAUL}_{m,m'} \big] \geq \delta/2$.

5.2 Ruling out Selective $\perp_{\mathsf{inc}}$ Attacks

Recall that $\tilde{v} = \perp_{\mathsf{inc}}$ when M's response $\tilde{a}$ is incorrect. In this section we will use the shorthand $p(\tau, m) = \Pr_{(\mathsf{R},\tilde{m}) \leftarrow \mathsf{M}_m^\tau} \big(\tilde{v} \neq \perp_{\mathsf{inc}} \big)$ and we will prove that if M is mauling then he is doing so by answering correctly. The main lemma of this section is the following.

Lemma 2. *If M breaks the non-malleability of $\langle \mathrm{C}, \mathrm{R} \rangle$ then there exist $m, m' \in \mathcal{M}$ and non-negligible $\delta, \delta' > 0$ such that for all $m^* \in \mathcal{M}$*

$$\Pr_\tau \left[\tau \in \mathsf{MAUL}_{m,m'} \;\&\; p(\tau, m^*) \geq \delta' \right] \geq \frac{\delta}{4} - \mathsf{negl}.$$

Lemma 2 follows immediately from (2) and Claims 4 and 5 which rule out two separate types of mauling behavior. We prove these claims in Appendix D.

Claim 4. *For all $m, m' \in \mathcal{M}$ and non-negligible $\xi > 0$ we have:*

$$\Pr_\tau \left[\Big| p(\tau, m) - p(\tau, m') \Big| > \xi \right] = \mathsf{negl}.$$

Claim 5. *Let $\delta = \delta(\lambda) > 0$ be as the statement of Lemma 3. For all $m, m' \in \mathcal{M}$, we have:*

$$\Pr_\tau \left[\tau \in \mathsf{MAUL}_{m,m'} \;\&\; p(\tau, m) < \lambda^{-2} \delta^3 \right] \leq \frac{\delta}{4}.$$

5.3 The Distribution $\mathcal{D}_{\mathrm{M}}$

We now use M to define a polynomial time sampleable distribution, $\mathcal{D}_{\mathrm{M}}$, which outputs a tampering function pair (f, g), as follows.

- **Random Choices:** Instantiate M and play the first two rounds of $\langle \mathrm{C}, \mathrm{R} \rangle$, obtaining a partial transcript $\tau = \big(\mathbf{Com}_{\mathsf{L}}, \tilde{\alpha}, a\big)$ where $a = \varphi(\alpha)$ and $\varphi(x)$ is the linear map specified by $\mathbf{Com}_{\mathsf{L}}$. Also, draw $\mathsf{R}_{\$} \leftarrow \mathsf{Enc}(m^*|\mathsf{L})$ for some arbitrary fixed $m^* \in \mathcal{M}$ and let $\tilde{a}_{\$}$ be M's response in the full transcript $\mathbb{T}(\tau, \mathsf{R}_{\$})$. Finally draw $\tilde{\beta} \leftarrow \mathbb{Z}_q$.
- $f_{\tau,\mathsf{R}_{\$},\tilde{\beta}}(\mathsf{L})$**:** Let $\psi(x)$ be the unique linear function with constant term L and $\psi(\alpha) = a$.
  - rewind M back to the second message of the right interaction and ask $\tilde{\beta}$, receive β on the left;
  - send $(b, \mathsf{R}_{\$})$ where $b = \psi(\beta)$ and receive $(\tilde{b}, \cdot)$ on the right;
  - output $\tilde{\mathsf{L}}$, the constant term of the line spanned by $\big\{(\tilde{\alpha}, \tilde{a}_{\$}), (\tilde{\beta}, \tilde{b})\big\}$.
- $g_{\tau,\mathsf{R}_{\$}}(\mathsf{R})$**:** Let $(\tilde{a}, \tilde{\mathsf{R}})$ be M's final message in $\mathbb{T}(\tau, \mathsf{R})$. If $\tilde{a} = \tilde{a}_{\$}$ output $\tilde{\mathsf{R}}$, otherwise $\perp_{\mathsf{inc}}$.
- **Output:** $(f, g) = (f_{\tau,\mathsf{R}_{\$},\tilde{\beta}}, g_{\tau,\mathsf{R}_{\$}})$.

Notice (f, g) output by $\mathcal{D}_{\mathrm{M}}$ are not split-state as the randomness $(\tau, \mathsf{R}_{\$})$ shared by both f and g depends on L. Nonetheless, we show in Section 5.4 below that $\mathcal{D}_{\mathrm{M}} \approx_{\mathsf{c}} \mathcal{D}_{\mathsf{split}}$, a distribution which outputs polytime split-state tampering functions. Combined with the next lemma, this shows that $\langle \mathrm{C}, \mathrm{R} \rangle$ is non-malleable: an M which breaks the non-malleability of $\langle \mathrm{C}, \mathrm{R} \rangle$ can be used to construct a distribution $\mathcal{D}_{\mathsf{split}}$ on $\mathcal{F}^{\mathsf{poly}}_{\mathsf{split}}$ which breaks the security of the code $(\mathsf{Enc}, \mathsf{Dec})$.

Lemma 3. *Let $\delta, \delta' > 0$ be as in the statement of Lemma 2. If M breaks the non-malleability of $\langle \mathrm{C}, \mathrm{R} \rangle$ then there exist $m, m' \in \mathcal{M}$ such that*

$$\Pr_{\mathsf{L},(f,g)}\left[\left|\Pr\big(\mathrm{D}^{\tau}(\mathsf{V}^{\mathsf{L}}_{m,f,g}) = 1\big) - \Pr\big(\mathrm{D}^{\tau}(\mathsf{V}^{\mathsf{L}}_{m',f,g}) = 1\big)\right| > \frac{\delta}{2}\right] > \frac{(\delta\delta')^3}{256} - \mathsf{negl},$$

where the outer probability is over $\mathsf{L} \leftarrow \mathcal{L}$ and $(f, g) = (f_{\tau,\mathsf{R}_{\$},\tilde{\beta}}, g_{\tau,\mathsf{R}_{\$}}) \leftarrow \mathcal{D}_{\mathrm{M}}$, where $\tau = (\mathbf{Com}_{\mathsf{L}}, \tilde{\alpha}, a)$

Proof Idea. The randomness needed in order to draw $(f, g) \leftarrow \mathcal{D}_{\mathrm{M}}$ consists of a random partial transcript $\tau = \big(\mathbf{Com}_{\mathsf{L}}, \tilde{\alpha}, a\big)$, $\mathsf{R}_{\$} \leftarrow \mathsf{Enc}(m^*|\mathsf{L})$ and $\tilde{\beta} \leftarrow \mathbb{Z}_q$. Given these choices, the distributions $\mathsf{V}^{\mathsf{L}}_{m,f,g}$ and M^{τ}_{m} are very similar: both draw $\mathsf{R} \leftarrow \mathsf{Enc}(m|\mathsf{L})$ and output $(\mathsf{R}, \tilde{m})$. We prove Lemma 3 by showing that whenver $\big(\tau, \mathsf{R}_{\$}, \tilde{\beta}\big)$ is such that M's response $\tilde{a}_{\$}$ is correct, $\mathsf{V}^{\mathsf{L}}_{m,f,g}$ and M^{τ}_{m} are actually identical for all $m \in \mathcal{M}$. The proof follows almost immediately since either M gives correct $\tilde{a}_{\$}$ with non-negligible probability, or he is always committing to $\perp_{\mathsf{inc}}$ given a commitment to m^* on the left. In the latter case, he cannot be mauling as we ruled out selective $\perp_{\mathsf{inc}}$ attacks in the previous section.

Proof of Lemma 3. For randomness $(\tau, \mathsf{R}_{\$}) = \big(\mathbf{Com}_{\mathsf{L}}, \tilde{\alpha}, r\alpha + \mathsf{L}, \mathsf{R}_{\$}\big)$, say the "extraction event", denoted EXT, occurs whenever $\tilde{a}_{\$}$ is correct in $\mathbb{T}(\tau, \mathsf{R}_{\$})$, and

$$\Pr_{\tilde{\beta}}\Big(\tilde{b} \text{ correct in } \mathbb{T}\big(\mathbf{Com}_{\mathsf{L}}, \tilde{\beta}, r\beta + \mathsf{L}, \mathsf{R}_{\$}\big)\Big) \geq \frac{(\delta\delta')^2}{32}.$$

It follows from Lemma 2 that if M mauls $\langle \mathrm{C}, \mathrm{R} \rangle$ then there exist messages $m, m' \in \mathcal{M}$ such that $\Pr_{\tau,\mathsf{R}_{\$}}\big(\tau \in \mathsf{MAUL}_{m,m'} \;\&\; \tilde{a}_{\$} \text{ correct in } \mathbb{T}(\tau, \mathsf{R}_{\$})\big) \geq \delta\delta'/4 - \mathsf{negl}$, and so using Bayes' theorem,

$$\Pr_{\tau,\mathsf{R}_{\$}}\big(\tau \in \mathsf{MAUL}_{m,m'} \;\&\; \mathsf{EXT}\big) \geq \frac{\delta\delta'}{4} - \frac{\delta\delta'}{8} - \mathsf{negl} = \frac{\delta\delta'}{8} - \mathsf{negl}.$$

If $\tilde{a}_\$$ is correct then $g(\mathsf{R})$ identifies when M is committing to $\perp_{\mathsf{inc}}$ on the right: $\tilde{a}$ is correct if and only if $\tilde{a} = \tilde{a}_\$$. Moreover, if EXT occurs then $f(\mathsf{L})$ outputs the correct value of $\tilde{\mathsf{L}}$ with probability at least $(\delta\delta')^2/32$ (over the choice of $\tilde{\beta} \leftarrow \mathbb{Z}_q$). Indeed, $f(\mathsf{L})$ outputs the correct $\tilde{\mathsf{L}}$ whenever $\tilde{b}$ is correct.

So we have seen that $\Pr_{\tau,\mathsf{R}_\$}\big(\tau \in \mathsf{MAUL}_{m,m'} \,\&\, \mathsf{EXT}\big) \geq \delta\delta'/8 - \mathsf{negl}$, and moreover conditioned on EXT occuring, $\tilde{a}_\$$ is correct and so $\mathsf{V}^{\mathsf{L}}_{m,f,g} \equiv \mathsf{M}^{\tau}_{m}$ for all $m \in \mathcal{M}$ with probability at least $(\delta\delta')^2/32$. But if $\tau \in \mathsf{MAUL}_{m,m'}$ and EXT occurs and $\tilde{a}_\$$ is correct then

$$\left|\Pr\big(\mathrm{D}^{\tau}(\mathsf{V}^{\mathsf{L}}_{m,f,g}) = 1\big) - \Pr\big(\mathrm{D}^{\tau}(\mathsf{V}^{\mathsf{L}}_{m',f,g}) = 1\big)\right| > \frac{\delta}{2},$$

and so Lemma 3 follows. □

5.4 A Hybrid Argument to Prove $\mathcal{D}_{\mathrm{M}} \approx_{\mathsf{c}} \mathcal{D}_{\mathsf{split}}$

Note that the distribution $\mathcal{D}_{\mathrm{M}}$ from the previous section does not output $(f, g) \in \mathcal{F}^{\mathsf{poly}}_{\mathsf{split}}$ as the randomness $(\tau, \mathsf{R}_\$) = \big(\mathrm{Com}(\mathsf{L}), \mathrm{Com}(r), \tilde{\alpha}, a, \mathsf{R}_\$\big)$ shared between f and g depends in three ways on L: 1) τ contains a commitment to L, 2) $a = r\alpha + \mathsf{L}$, and 3) $\mathsf{R}_\$ \leftarrow \mathsf{Enc}(m^*|\mathsf{L})$. We show in this section, however, that $\mathcal{D}_{\mathrm{M}}$ is computationally indistinguishable from a polynomial time sampleable distribution $\mathcal{D}_{\mathsf{split}}$ on $\mathcal{F}^{\mathsf{poly}}_{\mathsf{split}}$. This, together with Lemma 3, completes the proof that $\langle \mathrm{C}, \mathrm{R}\rangle$ is non-malleable since we will have used an M which breaks non-malleability to construct split-state, polytime tampering functions which break the non-malleability of the code $(\mathsf{Enc}, \mathsf{Dec})$.

$\mathcal{D}_0 = \mathcal{D}_{\mathrm{M}} -$ This is the distribution defined above. It draws $\mathsf{L} \leftarrow \mathcal{L}$, a random partial transcript $\tau = \big(\mathrm{Com}(\mathsf{L}), \mathrm{Com}(r), \tilde{\alpha}, a\big)$ where $a = r\alpha + \mathsf{L}$, $\mathsf{R}_\$ \leftarrow \mathsf{Enc}(m^*|\mathsf{L})$ and outputs $(f_0, g_0) = (f_{\tau,\mathsf{R}_\$}, g_{\tau,\mathsf{R}_\$})$.

$\mathcal{D}_1 -$ This distribution outputs functions (f_1, g_1) which behave exactly like (f_0, g_0) except that they are seeded with $(\tau, \mathsf{R}_\$) = \big(\mathrm{Com}(\mathsf{L}), \mathrm{Com}(r), \tilde{\alpha}, a, \mathsf{R}_\$\big)$, where $a \leftarrow \mathbb{Z}_q$ is random instead of equal to $r\alpha + \mathsf{L}$.

$\mathcal{D}_2 -$ This outputs (f_2, g_2) which, again, differ from (f_1, g_1) only in their shared randomness. This time $(\tau, \mathsf{R}_\$) = \big(\mathrm{Com}(0), \mathrm{Com}(r), \tilde{\alpha}, a, \mathsf{R}_\$\big)$, where $a \in \mathbb{Z}_q$ is random. Now the only dependence on L is that $\mathsf{R}_\$ \leftarrow \mathsf{Enc}(m^*|\mathsf{L})$.

$\mathcal{D}_3 = \mathcal{D}_{\mathsf{split}} -$ This outputs (f_3, g_3) which are the same as (f_2, g_2) except that $\mathsf{R}_\$ \leftarrow \mathcal{D}_{\mathsf{hid}}$ instead of $\big\{\mathsf{R} : \mathsf{L} \leftarrow \mathcal{L}, \mathsf{R} \leftarrow \mathsf{Enc}(m^*|\mathsf{L})\big\}$, where $\mathcal{D}_{\mathsf{hid}}$ is the distribution on $\mathcal{R}$ whose existence is guaranteed by the hiding property. Since $(\tau, \mathsf{R}_\$)$ no longer depends on L, $(f_3, g_3) \in \mathcal{F}^{\mathsf{poly}}_{\mathsf{split}}$.

Claim 6. $\mathcal{D}_0 \approx_{\mathsf{c}} \mathcal{D}_1 \approx_{\mathsf{c}} \mathcal{D}_2 \approx_{\mathsf{c}} \mathcal{D}_3$.

Proof. The first two indistinguishabilities follow from the hiding of Com. For the first, consider an adversary $\mathcal{A}$ who interacts with a challenger $\mathcal{C}$ in the hiding game by choosing $r_0, r_1 \in \mathbb{Z}_q$ at random and sending (r_0, r_1) to $\mathcal{C}$, receiving a commitment $z = \mathrm{Com}(r_b)$ for a random $b \in \{0, 1\}$. Then $\mathcal{A}$ draws $\mathsf{L} \leftarrow \mathcal{L}$ and $\tilde{\alpha} \leftarrow \mathbb{Z}_q$ at random and sends $\big(\mathrm{Com}(\mathsf{L}), z\big)$ and $\tilde{\alpha}$ to M (corresponding to the first message of the left interaction and the second message of the right interaction), receiving α as the second message on the left. $\mathcal{A}$ sets $a = r_0\alpha + \mathsf{L}$, draws $\mathsf{R}_\$ \leftarrow \mathsf{Enc}(m_\$|\mathsf{L})$ and outputs $\big(\mathrm{Com}(\mathsf{L}), z, \tilde{\alpha}, a, \mathsf{R}_\$\big)$. If $b = 0$ then $\mathcal{A}$'s output is distributed according to $\mathcal{D}_0$, while if $b = 1$, $\mathcal{A}$'s output is distributed like $\mathcal{D}_1$. This proves that $\mathcal{D}_0 \approx_{\mathsf{c}} \mathcal{D}_1$; $\mathcal{D}_1 \approx_{\mathsf{c}} \mathcal{D}_2$ follows even more readily. Finally, $\mathcal{D}_2 \approx_{\mathsf{c}} \mathcal{D}_3$ follows from the hiding property of $(\mathsf{Enc}, \mathsf{Dec})$. □

It follows from Lemma 3 and Claim 6 that if M breaks the non-malleability of $\langle \mathrm{C}, \mathrm{R}\rangle$ then there exist $m, m' \in \mathcal{M}$ and non-negligible $\delta, \delta' > 0$ such that

$$\Pr_{\mathsf{L},(f,g)}\left[\left|\Pr\big(\mathrm{D}^{\tau}_{m,m'}(\mathsf{V}^{\mathsf{L}}_{m,f,g}) = 1\big) - \Pr\big(\mathrm{D}^{\tau}_{m,m'}(\mathsf{V}^{\mathsf{L}}_{m',f,g}) = 1\big)\right| > \frac{\delta}{2}\right] > \frac{(\delta\delta')^3}{256} - \mathsf{negl},$$

where the outer probability is over $\mathsf{L} \leftarrow \mathcal{L}$ and $(f, g) \leftarrow \mathcal{D}_{\mathsf{split}}$, drawn independently. This breaks the computational conditional augmented non-malleability of $(\mathsf{Enc}, \mathsf{Dec})$, thus completing our proof that $\langle \mathrm{C}, \mathrm{R}\rangle$ is non-malleable.

References

[AAG+16] Divesh Aggarwal, Shashank Agrawal, Divya Gupta, Hemanta Maji, Omkant Pandey, and Manoj Prabhakaran. Optimal computational split-state non-malleable codes. In *TCC*, 2016.

[ADL14] Divesh Aggarwal, Yevgeniy Dodis, and Shachar Lovett. Non-malleable codes from additive combinatorics. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 774–783, 2014.

[Agg] Divesh Aggarwal. Personal communication, 10/30/2015.

[Bar02] Boaz Barak. Constant-Round Coin-Tossing with a Man in the Middle or Realizing the Shared Random String Model. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science*, FOCS '02, pages 345–355, 2002.

[BGR+15] Hai Brenner, Vipul Goyal, Silas Richelson, Alon Rosen, and Margarita Vald. Fast non-malleable commitments. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, Denver, CO, USA, October 12-6, 2015*, pages 1048–1057, 2015.

[CGMO09] Nishanth Chandran, Vipul Goyal, Ryan Moriarty, and Rafail Ostrovsky. Position based cryptography. In Shai Halevi, editor, *CRYPTO*, volume 5677 of *Lecture Notes in Computer Science*, pages 391–407. Springer, 2009.

[CLOS02] Ran Canetti, Yehuda Lindell, Rafail Ostrovsky, and Amit Sahai. Universally composable two-party and multi-party secure computation. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, STOC '02, pages 494–503, 2002.

[DDN91] Danny Dolev, Cynthia Dwork, and Moni Naor. Non-Malleable Cryptography (Extended Abstract). In *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing*, STOC '91, pages 542–552, 1991.

[DKO13] Stefan Dziembowski, Tomasz Kazana, and Maciej Obremski. Non-malleable codes from two-source extractors. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part II*, volume 8043 of *Lecture Notes in Computer Science*, pages 239–257. Springer, 2013.

[DPW10] Stefan Dziembowski, Krzysztof Pietrzak, and Daniel Wichs. Non-malleable codes. In *Innovations in Computer Science - ICS 2010, Tsinghua University, Beijing, China, January 5-7, 2010. Proceedings*, pages 434–452, 2010.

[GK96] Oded Goldreich and Hugo Krawczyk. On the composition of zero-knowledge proof systems. *SIAM J. Comput.*, 25(1):169–192, 1996.

[GLOV12] Vipul Goyal, Chen-Kuei Lee, Rafail Ostrovsky, and Ivan Visconti. Constructing non-malleable commitments: A black-box approach. In *FOCS*, pages 51–60. IEEE Computer Society, 2012.

[Goy11] Vipul Goyal. Constant Round Non-malleable Protocols Using One-way Functions. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*, STOC '11, pages 695–704. ACM, 2011.

[GRRV14] Vipul Goyal, Silas Richelson, Alon Rosen, and Margarita Vald. An algebraic approach to non-malleability. In *FOCS*, 2014.

[IKOS07] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Zero-knowledge from Secure Multiparty Computation. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, STOC '07, pages 21–30, 2007.

[Kiy14] Susumu Kiyoshima. Round-efficient black-box construction of composable multi-party computation. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part II*, volume 8617 of *Lecture Notes in Computer Science*, pages 351–368. Springer, 2014.

[KMO14] Susumu Kiyoshima, Yoshifumi Manabe, and Tatsuaki Okamoto. Constant-round black-box construction of composable multi-party computation protocol. In Yehuda Lindell, editor, *Theory of Cryptography - 11th Theory of Cryptography Conference, TCC 2014, San Diego, CA, USA, February 24-26, 2014. Proceedings*, volume 8349 of *Lecture Notes in Computer Science*, pages 343–367. Springer, 2014.

[LP09] Huijia Lin and Rafael Pass. Non-malleability Amplification. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*, STOC '09, pages 189–198, 2009.

[LP11] Huijia Lin and Rafael Pass. Constant-round Non-malleable Commitments from Any One-way Function. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*, STOC '11, pages 705–714, 2011.

[LP12] Huijia Lin and Rafael Pass. Black-box constructions of composable protocols without set-up. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology - CRYPTO 2012 - 32nd Annual Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2012. Proceedings*, volume 7417 of *Lecture Notes in Computer Science*, pages 461–478. Springer, 2012.

[LPV08] Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkitasubramaniam. Concurrent Non-malleable Commitments from Any One-Way Function. In *Theory of Cryptography, 5th Theory of Cryptography Conference, TCC 2008*, pages 571–588, 2008.

[LPV09] Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkitasubramaniam. A Unified Framework for Concurrent Security: Universal Composability from Stand-alone Non-malleability. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*, STOC '09, pages 179–188, 2009.

[Pas13] Rafael Pass. Unprovable security of perfect NIZK and non-interactive non-malleable commitments. In *TCC*, pages 334–354, 2013.

[PPV08] Omkant Pandey, Rafael Pass, and Vinod Vaikuntanathan. Adaptive One-Way Functions and Applications. In *Advances in Cryptology — CRYPTO '08*, pages 57–74, 2008.

[PR05a] Rafael Pass and Alon Rosen. Concurrent Non-Malleable Commitments. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of ComputerScience*, FOCS '05, pages 563–572, 2005.

[PR05b] Rafael Pass and Alon Rosen. New and improved constructions of non-malleable cryptographic protocols. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, STOC '05, pages 533–542, 2005.

[PW10] Rafael Pass and Hoeteck Wee. Constant-Round Non-malleable Commitments from Sub-exponential One-Way Functions. In *Advances in Cryptology — EUROCRYPT '10*, pages 638–655, 2010.

[Wee10] Hoeteck Wee. Black-Box, Round-Efficient Secure Computation via Non-malleability Amplification. In *Proceedings of the 51th Annual IEEE Symposium on Foundations of Computer Science*, pages 531–540, 2010.

A Preliminaries

A.1 Commitment schemes

Commitment schemes are used to enable a party, known as the sender, to commit itself to a value while keeping it secret from the receiver (hiding). Furthermore, the commitment should be such that when, in a later (decommitment)

stage, the commitment is opened, there is a single value which can result (binding). In this work, we consider commitment schemes that are statistically (or perfectly) binding, namely while the hiding property only holds against computationally bounded (non-uniform) adversaries, the binding property is required to hold against unbounded adversaries. We denote a commitment scheme $\langle \mathrm{C}(m), \mathsf{R} \rangle$.

Definition 6 (Statistically Binding Commitment Scheme). *Let $\langle \mathrm{C}(m), \mathsf{R} \rangle$ be a two phase protocol between C and R where m is C's secret input. Let $z = \mathrm{Com}(m; r)$ denote R's view after the first phase. Let $(m, w) = \mathrm{Decom}(m, r, z)$ be R's view after the second phase. We say that $\langle \mathrm{C}(m), \mathsf{R} \rangle$ is a* statistically binding commitment scheme *if the following properties hold:*

Correctness: If parties follow the protocol, then $\mathsf{R}(z, m, w) = 1$;

Binding: With high probability over R's randomness, there does not exist a (m', w') with $m' \neq m$ such that $\mathsf{R}(z, m', w') = 1$;

Hiding: For all $m_0 \neq m_1$, $\{\mathrm{Com}(m_0; r)\}_r \approx_c \{\mathrm{Com}(m_1; r)\}_r$.

Tag-based Commitment Scheme. Following [PR05b, DDN91], we consider tag-based commitment schemes where, in addition to the security parameter, the committer and the receiver also receive a "tag" a.k.a. the identity id as common input.

A.2 Non-malleable commitments

We follow the definition of non-malleable commitments of Lin et al [LPV08]. In the real interaction, there is a man-in-the-middle adversary M interacting with a committer, C, in the left session a receiver R in the right. We denote the various quantities associated with the right interaction as "tilde'd" versions of their left counterparts. So for example, C commits to m in the left interaction while M commits to $\tilde{m}$ in the right. Let $\mathbf{MIM}_{\langle \mathrm{C}, \mathsf{R} \rangle}(m, z)$ denote a random variable that describes $(\tilde{m}, v)$, the value M commits to and M's view in the full experiment. In the simulated experiment, a simulator $\mathcal{S}$ directly interacts with M. Let $\mathbf{SIM}^{\mathcal{S}}_{\langle \mathrm{C}, \mathsf{R} \rangle}(1^\lambda, z)$ denote the random variable describing $(\tilde{m}, v)$ in this simulated interaction. If the tag tag for the left interaction is equal to the tag $\tilde{tag}$ for the right interaction, the value $\tilde{m}$ committed to in the right interaction is defined to be $\perp$ in both experiments. This is analogous to the uninteresting case when C is committing to himself.

Definition 7 (Non-Malleable Commitments). *A commitment scheme $\langle \mathrm{C}(m), \mathsf{R} \rangle$ is said to be non-malleable if for every PPT man-in-the-middle adversary M, there exists a PPT simulator $\mathcal{S}$ such that the following ensembles are computationally indistinguishable:*

$$\{\mathbf{MIM}_{\langle \mathrm{C}, \mathsf{R} \rangle}(m, z)\}_{m \in \{0,1\}^\lambda, z \in \{0,1\}^\star}, \text{ and } \{\mathbf{SIM}^{\mathcal{S}}_{\langle \mathrm{C}, R \rangle}(1^\lambda, z)\}_{z \in \{0,1\}^\star}.$$

A.3 Non-Malleable Codes

A *coding scheme* is a pair of functions $(\mathsf{Enc}, \mathsf{Dec})$ where $\mathsf{Enc} : \mathcal{M} \to \mathcal{C}$ and $\mathsf{Dec} : \mathcal{C} \to \mathcal{M}$ for a message space $\mathcal{M}$ and codeword space $\mathcal{C}$. It should be the case that $\mathsf{Dec} \circ \mathsf{Enc}(m) = m$ for all $m \in \mathcal{M}$ with high probability over the randomness of Enc (it needn't be the case that Enc is randomized at all, in which case correctness requires $\mathsf{Dec} \circ \mathsf{Enc}(m) = m$ with probability 1). Historically, coding schemes are usually designed in order to be resilient to some form of tampering. In their important 2010 paper Dziembowski, Pietrzak and Wichs [DPW10] introduced non-malleable codes which are codes with strong security in the presence of tampering. Informally, for a family $\mathcal{F} \subset \{f : \mathcal{C} \to \mathcal{C}\}$, we say that $(\mathsf{Enc}, \mathsf{Dec})$ is non-malleable with respect to $\mathcal{F}$ if for all $f \in \mathcal{F}$, the tamper distribution $\big(\mathsf{Dec} \circ f \circ \mathsf{Enc}\big)(m)$ (over the randomness of Enc) outputs $\tilde{m}$ which is either equal to m if f copying or else is independent of m. In this work we are interested in split-state non-malleable codes.

Let $(\mathsf{Enc}, \mathsf{Dec})$ be a split state coding scheme so $\mathsf{Enc} : \mathcal{M} \to \mathcal{L} \times \mathcal{R}$ and let

$$\mathcal{F}_{\mathsf{split}} = \big\{(f, g) \big| f : \mathcal{L} \to \mathcal{L}, g : \mathcal{R} \to \mathcal{R}\big\}$$

be the set of split-state tampering functions.

Definition 8 (Tampering Distribution). *Fix $m \in \mathcal{M}$, $(f,g) \in \mathcal{F}_{\mathsf{split}}$. The* tampering distribution*, denoted $\mathsf{T}_{m,f,g}$ is: draw $(\mathsf{L},\mathsf{R}) \leftarrow \mathsf{Enc}(m)$, set $(\tilde{\mathsf{L}},\tilde{\mathsf{R}}) = \big(f(\mathsf{L}), g(\mathsf{R})\big)$ and output $\tilde{m} = \mathsf{Dec}(\tilde{\mathsf{L}},\tilde{\mathsf{R}})$.*

Definition 9 (Simulatable Distribution). *Let $\{\mathsf{D}_m\}_{m\in\mathcal{M}}$ be a family of distributions on $\mathcal{M}$ indexed by m. We say that $\{\mathsf{D}_m\}$ is* $\varepsilon-$simulatable *if there exists a distribution S on $\mathcal{M} \cup \{\mathsf{same}\}$ such that $\Delta(\mathsf{D}_m, \mathsf{S}_m) < \varepsilon$ for all m, where S_m is the distribution on $\mathcal{M}$ induced by drawing $\tilde{m} \leftarrow \mathsf{S}$ and outputting m if $\tilde{m} = \mathsf{same}$, $\tilde{m}$ if not.*

Definition 10 (Split-State Non-Malleable Code). *We say that $(\mathsf{Enc},\mathsf{Dec})$ is* $\varepsilon-$non-malleable against $\mathcal{F}_{\mathsf{split}}$ *if for all $(f,g) \in \mathcal{F}_{\mathsf{split}}$, $\{\mathsf{T}_{m,f,g}\}_m$ is $\varepsilon-$simulatable.*

The Non-Malleable Code of [ADL14]. The construction of [ADL14] encodes $m \in \mathcal{M}$ into (L,R) where $\mathsf{L},\mathsf{R} \in \mathbb{Z}_p^n$ are random subject to the condition that $\langle \mathsf{L},\mathsf{R}\rangle \in \mathsf{H}_m \subset \mathbb{Z}_p$ (p is a prime much larger than $|\mathcal{M}|$, and the $\{\mathsf{H}_m\}_{m\in\mathcal{M}}$ are carefully chosen disjoint subsets of $\mathbb{Z}_p$). Non-malleability follows from an extensive analysis of the inner product function which makes heavy use of its properties as a randomness extractor. For any $(f,g) \in \mathcal{F}_{\mathsf{split}}$ and $x \in \mathbb{Z}_p$, the following random process is considered: choose $\mathsf{L},\mathsf{R} \in \mathbb{Z}_p^n$ randomly such that $\langle \mathsf{L},\mathsf{R}\rangle = x$, set $(\tilde{\mathsf{L}},\tilde{\mathsf{R}}) = \big(f(\mathsf{L}), g(\mathsf{R})\big)$, and output $\tilde{x} = \langle \tilde{\mathsf{L}},\tilde{\mathsf{R}}\rangle$. The main lemma of [ADL14] says that $\tilde{x}$ is either 1) independent of x, or 2) of the form $\tilde{x} = ax + b$ for some $a,b \in \mathbb{Z}_p$ which depend only on (f,g). Non-malleability in [ADL14] then follows from the design of *affine evasive sets* as the $\{\mathsf{H}_m\}_m$. This aspect of their construction is very elegant but as we will not need to change their $\{\mathsf{H}_m\}_m$, we do not discuss this portion further. The interested reader should see [ADL14] for more information. We note that the earlier work of [DKO13] used essentially the same outline in order to give a non-malleable code for one bit messages. Their construction is also very elegant and is much simpler: they use $\mathsf{H}_0 = \{0\}$ and $\mathsf{H}_1 = \mathbb{Z}_p - \{0\}$.

A.4 Augmented Non-Malleable Codes

Very recently Aggarwal et al. [AAG+16] proved that the [ADL14] construction is non-malleable even when the tamper distribution outputs $\tilde{m}$ along with one of the states. Their proof looks at the randomized process: choose $\mathsf{L},\mathsf{R} \in \mathbb{Z}_p^n$ randomly such that $\langle \mathsf{L},\mathsf{R}\rangle = x$, set $(\tilde{\mathsf{L}},\tilde{\mathsf{R}}) = \big(f(\mathsf{L}), g(\mathsf{R})\big)$ and output $(\mathsf{R},\tilde{x})$ where $\tilde{x} = \langle \tilde{\mathsf{L}},\tilde{\mathsf{R}}\rangle$. The same randomness extraction properties of the inner product function used in [ADL14] show that even conditioned on R, $\tilde{x}$ is either independent of x or else $\tilde{x} = ax + b$ for $a,b \in \mathbb{Z}_p$ which depend only on (f,g). They call this stronger notion augmented non-malleability.

Definition 11 (Augmented Tampering Distribution). *Fix $m \in \mathcal{M}$ and $(f,g) \in \mathcal{F}_{\mathsf{split}}$. The* augmented tampering distribution*, denoted $\mathsf{V}_{m,f,g}$ is: draw $(\mathsf{L},\mathsf{R}) \leftarrow \mathsf{Enc}(m)$, set $(\tilde{\mathsf{L}},\tilde{\mathsf{R}}) = \big(f(\mathsf{L}), g(\mathsf{R})\big)$ and output $(\mathsf{R},\tilde{m})$ where $\tilde{m} = \mathsf{Dec}(\tilde{\mathsf{L}},\tilde{\mathsf{R}})$.*

We use the letter V for "view": the output of the augmented tampering distribution will be basically what the MIM sees during a mauling attack on our non-malleable commitment scheme.

Definition 12 (Augmented Simulatable Distribution). *Let $\{\mathsf{D}_m\}_{m\in\mathcal{M}}$ be a family of distributions on $\mathcal{R}\times\mathcal{M}$ indexed by m, where $\mathcal{R}$ is an arbitrary set. We say that $\{\mathsf{D}_m\}$ is* $\varepsilon-$augmented simulatable *if there exists a distribution S on $\mathcal{R}\times(\mathcal{M}\cup\{\mathsf{same}\})$ such that $\Delta(\mathsf{D}_m,\mathsf{S}_m) < \varepsilon$ for all m, where S_m is the distribution on $\mathcal{M}$ induced by drawing $(\mathsf{R},\tilde{m}) \leftarrow \mathsf{S}$ and outputting (R,m) if $\tilde{m} = \mathsf{same}$, $(\mathsf{R},\tilde{m})$ if not.*

Definition 13 (Augmented Non-Malleable Code). *We say that $(\mathsf{Enc},\mathsf{Dec})$ is* $\varepsilon-$augmented non-malleable *against $\mathcal{F}_{\mathsf{split}}$ if for all $(f,g) \in \mathcal{F}_{\mathsf{split}}$, $\{\mathsf{V}_{m,f,g}\}_m$ is $\varepsilon-$augmented simulatable.*

Before moving on, we remark that the proof in [AAG+16] actually shows something slightly stronger. Let us think of the randomized process above as drawing $\mathsf{L} \leftarrow \mathbb{Z}_p^n$ at random and then drawing $\mathsf{R} \leftarrow \{\mathbf{v} \in \mathbb{Z}_p^n : \langle \mathsf{L},\mathbf{v}\rangle = x\}$, computing $(\tilde{\mathsf{L}},\tilde{\mathsf{R}})$ and outputting $(\mathsf{R},\tilde{x})$. The analysis of [AAG+16] does not actually require L to be uniform in $\mathbb{Z}_p^n$ and works whenever L has sufficient min-entropy. In particular, given $n,p,\mathcal{M},\{\mathsf{H}_m\}_m$ as in [ADL14] and any sufficiently large $\mathcal{L} \subset \mathbb{Z}_p^n$, define the coding scheme:

- $\underline{\mathsf{Enc}^{\mathcal{L}}(m)}$: choose $\mathsf{L} \leftarrow \mathcal{L}$ and $\mathsf{R} \leftarrow \mathbb{Z}_p^n$ randomly such that $\langle \mathsf{L},\mathsf{R}\rangle \in \mathsf{H}_m$.

- $\underline{\mathsf{Dec}(\mathsf{L},\mathsf{R})}$: if $\langle \mathsf{L},\mathsf{R}\rangle \in \mathsf{H}_m$, output m, otherwise output $\perp$.

Claim 7. *There exists an absolute constant $c > 0$ such that for all $\mathcal{M}$, there exist $n, p = \mathsf{poly}(|\mathcal{M}|, \lambda)$ such that for all $\mathcal{L} \subset \mathbb{Z}_p^n$ of size at least $|\mathcal{L}| \geq p^{cn}$, $(\mathsf{Enc}^{\mathcal{L}}, \mathsf{Dec})$ is $2^{-\Omega(\lambda)}-$augmented non-malleable.*

Moreover, the simulator for $(\mathsf{Enc}^{\mathcal{L}}, \mathsf{Dec})$ is identical to the simulator for $(\mathsf{Enc}, \mathsf{Dec})$ except that it draws $\mathsf{L} \leftarrow \mathcal{L}$ instead of $\mathsf{L} \leftarrow \mathbb{Z}_p^n$. Claim 7 follows from the proof of the main theorem in [AAG+16].

B Proof of Claim 1

The fact that the code from [ADL14] is conditionally augmented non-malleable is implicit in the recent work of [AAG+16, Agg] (which in turn builds on [ADL14]). Below we sketch an independent proof using a simple probability argument. This is reminiscent of the way in which one argues that a sufficiently good two-source extractor is also a strong two-source extractor.

Claim 1 (Restated). *The code $(\mathsf{Enc}, \mathsf{Dec})$ of [ADL14] is $\varepsilon'-$conditionally augmented non-malleable for some negligible quantity $\varepsilon' = \varepsilon'(\lambda) > 0$.*

Proof. Given $(f, g) \in \mathcal{F}_{\mathsf{split}}$, the simulator $\mathsf{S}_{f,g}$ guaranteed by the augmented non-malleability of $(\mathsf{Enc}, \mathsf{Dec})$ behaves as follows: draw $\mathsf{L}, \mathsf{R} \leftarrow \mathbb{Z}_p^n$ at random, set $(\tilde{\mathsf{L}}, \tilde{\mathsf{R}}) = \big(f(\mathsf{L}), g(\mathsf{R})\big)$ and output $\big(\mathsf{R}, \langle \tilde{\mathsf{L}}, \tilde{\mathsf{R}}\rangle\big)$ unless $\langle \tilde{\mathsf{L}}, \tilde{\mathsf{R}}\rangle = \langle \mathsf{L}, \mathsf{R}\rangle$, in which case output $(\mathsf{R}, \mathsf{same})$. We define the family $\{\mathsf{S}_{f,g}^{\mathsf{L}}\}_{\mathsf{L}}$ of simulators similarly: $\mathsf{S}_{f,g}^{\mathsf{L}}$ draws $\mathsf{R} \leftarrow \mathbb{Z}_p^n$ at random and outputs $\big(\mathsf{R}, \langle \tilde{\mathsf{L}}, \tilde{\mathsf{R}}\rangle\big)$ or $(\mathsf{R}, \mathsf{same})$ according to whether $\langle \tilde{\mathsf{L}}, \tilde{\mathsf{R}}\rangle$ is distinct from or equal to $\langle \mathsf{L}, \mathsf{R}\rangle$. Define

$$\mathcal{L}_{\mathsf{bad}} = \Big\{ \mathsf{L} \in \mathbb{Z}_p^n : \exists\, m \in \mathcal{M} \text{ st } \Delta\big(\mathsf{V}_{m,f,g}^{\mathsf{L}}, \mathsf{S}_{m,f,g}^{\mathsf{L}}\big) > \varepsilon' \Big\}.$$

If $|\mathcal{L}_{\mathsf{bad}}| < \varepsilon' p^n$ then we are done so assume $|\mathcal{L}_{\mathsf{bad}}| \geq \varepsilon' p^n$. By Claim 7, the restricted code $(\mathsf{Enc}^{\mathcal{L}_{\mathsf{bad}}}, \mathsf{Dec})$ is $2^{-\Omega(\lambda)}-$augmented non-malleable with simulator $\mathsf{S}_{f,g}^{\mathcal{L}_{\mathsf{bad}}}$ identical to that for $(\mathsf{Enc}, \mathsf{Dec})$ except that the initial choices of L, R are $\mathsf{L} \leftarrow \mathcal{L}_{\mathsf{bad}}$, $\mathsf{R} \leftarrow \mathbb{Z}_p^n$. However, by definition of $\mathcal{L}_{\mathsf{bad}}$, $\mathsf{S}_{f,g}^{\mathsf{L}}$ does not simulate $\big\{\mathsf{V}_{m,f,g}^{\mathsf{L}}\big\}_m$ when $\mathsf{L} \in \mathcal{L}_{\mathsf{bad}}$. Therefore, it must be that $\big|\mathcal{L}_{\mathsf{bad}}\big| < \varepsilon' p^n$ and so $(\mathsf{Enc}, \mathsf{Dec})$ is $\varepsilon'-$conditionally augmented non-malleable. □

C Proof of Lemma 1

Lemma 1 (Restated). *If $(\mathsf{G}, \mathsf{E}, \mathsf{D})$ is a semantically secure private key encryption scheme with the random ciphertext property, Com is a perfectly binding non-interactive commitment scheme and $(\mathsf{Enc}_0, \mathsf{Dec}_0)$ is $\varepsilon-$conditionally augmented non-malleable against $\mathcal{F}_{\mathsf{split}}$ for negligible $\varepsilon > 0$, then $(\mathsf{Enc}, \mathsf{Dec})$ is computationally conditionally augmented non-malleable against $\mathcal{F}_{\mathsf{split}}^{\mathsf{poly}}$.*

Claim 8. *For any $m, m' \in \mathcal{M}$, $(f, g) \in \mathcal{F}_{\mathsf{split}}^{\mathsf{poly}}$ and non-negligible $\xi > 0$ we have*

$$\Pr_{\mathsf{L}}\left[\left|\Pr_{(\mathsf{R},\tilde{m})\leftarrow \mathsf{V}_{m,f,g}^{\mathsf{L}}}\big(\tilde{m} \neq \perp_{\mathsf{com}}\big) - \Pr_{(\mathsf{R},\tilde{m})\leftarrow \mathsf{V}_{m',f,g}^{\mathsf{L}}}\big(\tilde{m} \neq \perp_{\mathsf{com}}\big)\right| > \xi\right] = \mathsf{negl}. \tag{3}$$

Proof. Fix $m, m' \in \mathcal{M}$, and $(f, g) \in \mathcal{F}_{\mathsf{split}}^{\mathsf{poly}}$. Let BAD be the set of $\mathsf{L} \in \mathcal{L}$ such that the inequality $\mathsf{p}_{\mathsf{L},m} > \xi + \mathsf{p}_{\mathsf{L},m'}$ holds (using shorthand $\mathsf{p}_{\mathsf{L},m}$ instead of $\Pr_{(\mathsf{R},\tilde{m})\leftarrow \mathsf{V}_{m,f,g}^{\mathsf{L}}}\big(\tilde{m} \neq \perp_{\mathsf{com}}\big)$ as above), and suppose for contradiction that there is a non-negligible $\xi' = \xi'(\lambda) > 0$ such that $\Pr_{\mathsf{L}}\big[\mathsf{L} \in \mathsf{BAD}\big] \geq \xi'$. Fix $N = 3/(\xi\xi')$, $N' = \Omega(\lambda\xi^{-2})$ and consider the PPT adversary $\mathcal{A}$ who interacts with a challenger $\mathcal{C}$ as follows.

- $\mathcal{A}$ sends m, m' to $\mathcal{C}$ and receives $\big(\mathsf{R}, \{\mathsf{R}_1\}, \ldots, \{\mathsf{R}_N\}\big)$.
- $\mathcal{A}$ computes $(\tilde{c}, \tilde{z}) = g(\mathsf{R})$ and $(\tilde{c}_i, \tilde{z}_i) = g(\mathsf{R}_i)$ for each $i = 1, \ldots, N$ and $\mathsf{R}_i \in \{\mathsf{R}_i\}$.
- For $i = 1, \ldots, N$, $\mathcal{A}$ sets $p_i = \Pr_{\mathsf{R}_i \in \{\mathsf{R}_i\}}\big(\tilde{z}_i = \tilde{z}\big)$ and outputs i^* such that p_{i^*} is maximal.

Note that if the random secret $\mathsf{L} \in \mathcal{L}$ chosen by $\mathcal{C}$ is in BAD then $\tilde{z}$ is correct (*i.e.*, $\mathrm{Decom}(\tilde{z}) = (\tilde{k}, \tilde{\sigma})$ where $f(\mathsf{L}) = \big(\tilde{\mathsf{L}}_0, (\tilde{k}, \tilde{\sigma})\big)$) with probability at least ξ. If $\tilde{z}$ is correct then for all i and $\mathsf{R}_i \in \{\mathsf{R}_i\}$, $\mathsf{Dec}(\tilde{\mathsf{L}}, \tilde{\mathsf{R}}_i) \neq \perp_{\mathsf{com}}$ if and only if $\tilde{z}_i = \tilde{z}$. Therefore, p_i approximates $\mathsf{p}_{\mathsf{L},m_i}$ where $m_i = m'$ if $i \neq i^*$ and $m_{i^*} = m$. Therefore,

$$p_{i^*} \geq \mathsf{p}_{\mathsf{L},m} - \frac{\xi}{3} > \mathsf{p}_{\mathsf{L},m'} + \xi - \frac{\xi}{3} \geq p_i + \frac{\xi}{3},$$

for all $i \neq i^*$. We have used the Chernoff-Hoeffding bound, facilitated by our choice of large N'. So we see that

$$\Pr\big(\mathcal{A} \text{ wins}\big) \geq \Pr\big(\mathsf{L} \in \mathsf{BAD}\big)\Pr\big(\tilde{z} \text{ correct} \big| \mathsf{L} \in \mathsf{BAD}\big)(1 - \mathsf{negl}) = \xi\xi' - \mathsf{negl} > \frac{2}{N},$$

and so $\mathcal{A}$ breaks the right state simulatability of $(\mathsf{Enc}, \mathsf{Dec})$. □

Fix any $m^* \in \mathcal{M}$ and non-negligible $\delta = \delta(\lambda) > 0$ and define

$$\mathcal{L}_{\mathsf{BOT}} = \Big\{\mathsf{L} \in \mathcal{L} : \Pr_{(\mathsf{R},\tilde{m}) \leftarrow \mathsf{V}^{\mathsf{L}}_{m^*,f,g}}\big(\tilde{m} \neq \perp_{\mathsf{com}}\big) < \lambda^{-2}\delta^3\Big\}.$$

If $\mathsf{L} \in \mathcal{L}_{\mathsf{BOT}}$ then $\mathsf{S}^{\mathsf{L}}_{f,g}$ samples $\mathsf{R} \leftarrow \mathsf{Enc}(m^*|\mathsf{L})$ and outputs $(\mathsf{R}, \perp_{\mathsf{com}})$.

Claim 9. *For any PPT distinguisher* D*,*

$$\Pr_{\mathsf{L}}\Big[\mathsf{L} \in \mathcal{L}_{\mathsf{BOT}} \;\&\; \exists\, m \; st \; \Big|\Pr_{(\mathsf{R},\tilde{m}) \leftarrow \mathsf{V}^{\mathsf{L}}_{m,f,g}}\big(\mathrm{D}(\mathsf{R},\tilde{m}) = 1\big) - \Pr_{(\mathsf{R},\tilde{m}) \leftarrow \mathsf{S}^{\mathsf{L}}_{m,f,g}}\big(\mathrm{D}(\mathsf{R},\tilde{m}) = 1\big)\Big| > \delta\Big] < \delta.$$

Proof. We use the hiding machine again. Set BAD' to be the set of $\mathsf{L} \in \mathcal{L}_{\mathsf{BOT}}$ such that there exists $m \in \mathcal{M}$ such that

$$\Pr_{(\mathsf{R},\tilde{m}) \leftarrow \mathsf{V}^{\mathsf{L}}_{m,f,g}}\big(\mathrm{D}(\mathsf{R},\tilde{m}) = 1\big) > \delta + \Pr_{(\mathsf{R},\tilde{m}) \leftarrow \mathsf{S}^{\mathsf{L}}_{m,f,g}}\big(\mathrm{D}(\mathsf{R},\tilde{m}) = 1\big).$$

Assume for contradiction that $\Pr_{\mathsf{L}}\big[\mathsf{L} \in \mathsf{BAD}'\big] \geq \delta$. Fix $N = 5/\delta$, $N' = \Omega(\lambda\delta^{-2})$ and consider the PPT adversary $\mathcal{A}$ who interacts with a challenger $\mathcal{C}$ as follows.

- $\mathcal{A}$ sends m, m^* to $\mathcal{C}$ and receives $\big(\mathsf{R}, \{\mathsf{R}_1\}, \ldots, \{\mathsf{R}_N\}\big)$.
- For $i = 1, \ldots, N$, $\mathcal{A}$ sets $p_i = \Pr_{\mathsf{R}_i \in \{\mathsf{R}_i\}}\big(\mathrm{D}(\mathsf{R}, \perp_{\mathsf{com}}) = 1\big)$ and outputs i^* such that p_{i^*} is maximal.

Note that if the random secret $\mathsf{L} \in \mathcal{L}$ chosen by $\mathcal{C}$ is in BAD' and furthermore is not in the negligible fraction of $\mathcal{L}$ for which (3) does not hold, then $\mathsf{p}_{\mathsf{L},m}, \mathsf{p}_{\mathsf{L},m^*} < 2\lambda^{-2}\delta^3$ (this uses the definition of $\mathcal{L}_{\mathsf{BOT}}$ and Claim 8 with $\xi = \lambda^{-2}\delta^3$). Therefore, the expected number of (i, R_i) such that $i \in \{1, \ldots, N\}$, $\mathsf{R}_i \in \{\mathsf{R}_i\}$, and $\tilde{m}_i \neq \perp_{\mathsf{com}}$ is at most $2\lambda^{-2}\delta^3 NN' < 1/2$, where $\tilde{m}_i = \mathsf{Dec}(\tilde{\mathsf{L}}, \tilde{\mathsf{R}}_i)$; so with probability at least $1/2$ there exist no such (i, R_i). In this case, p_i approximates $\Pr_{(\mathsf{R},\tilde{m}) \leftarrow \mathsf{X}}\big(\mathrm{D}(\mathsf{R},\tilde{m}) = 1\big)$, where $\mathsf{X} = \mathsf{S}^{\mathsf{L}}_{f,g}$ when $i \neq i^*$ and $\mathsf{X} = \mathsf{V}^{\mathsf{L}}_{m,f,g}$ when $i = i^*$. Therefore,

$$p_{i^*} \geq \Pr_{(\mathsf{R},\tilde{m}) \leftarrow \mathsf{V}^{\mathsf{L}}_{m,f,g}}\big(\mathrm{D}(\mathsf{R},\tilde{m})\big) - \frac{\delta}{3} > \Pr_{(\mathsf{R},\tilde{m}) \leftarrow \mathsf{S}^{\mathsf{L}}_{f,g}}\big(\mathrm{D}(\mathsf{R},\tilde{m}) = 1\big) + \delta - \frac{\delta}{3} \geq p_i + \frac{\delta}{3},$$

for all $i \neq i^*$. We have used the Chernoff-Hoeffding bound with sufficiently large N'. So we see that

$$\Pr\big(\mathcal{A} \text{ wins}\big) \geq \Pr\big(\mathsf{L} \in \mathsf{BAD}'\big)\Pr\big(\tilde{m}_i = \perp_{\mathsf{com}} \forall\, (i, \mathsf{R}_i) \big| \mathsf{L} \in \mathsf{BAD}'\big) - \mathsf{negl} \geq \frac{\delta}{2} - \mathsf{negl} > \frac{2}{N},$$

and so $\mathcal{A}$ breaks the right state simulatability of $(\mathsf{Enc}, \mathsf{Dec})$. □

Let $\mathcal{L}_{\mathsf{VALID}} = \mathcal{L} \setminus \mathcal{L}_{\mathsf{BOT}}$ be the set of L for which $\mathsf{p}_{\mathsf{L},m^*} \geq \lambda^2\delta^{-3}$ and furthermore, assume $\mathcal{L}_{\mathsf{VALID}}$ comprises at least a $\delta-$fraction of $\mathcal{L}$; if not, we are done as $\{\mathsf{S}^{\mathsf{L}}_{f,g}\}_{\mathsf{L} \in \mathcal{L}_{\mathsf{BOT}}}$ are sufficient to conditionally simulate $\big\{\mathsf{V}^{\mathsf{L}}_{m,f,g}\big\}_{m,\mathsf{L}}$. Given (f, g), we construct $(f_0, g_0) \in \mathcal{F}_{\mathsf{split}}$ which tampers $(\mathsf{Enc}_0, \mathsf{Dec}_0)$. We will then use $\mathsf{S}^{\mathsf{L}_0}_{f_0,g_0}$ whose existence is guaranteed by the conditional non-malleability of $(\mathsf{Enc}_0, \mathsf{Dec}_0)$, to construct $\mathsf{S}^{\mathsf{L}}_{f,g}$. The construction of (f_0, g_0) is as follows.

- **Random Choices:** Draw $k \leftarrow \mathsf{G}(1^\lambda)$, $\sigma \leftarrow \$$, set $z = \mathrm{Com}(k,\sigma)$, and draw $c^\$ \leftarrow \mathsf{E}_k(0)$. Set $(\cdot, \tilde{z}^\$) = g(c^\$, z)$. Save the values $(k, \sigma, z, \tilde{z}^\$)$.
- $f_0(\mathsf{L}_0)$**:** Compute $\big(\tilde{\mathsf{L}}_0, (\tilde{k}, \tilde{\sigma})\big) = f\big(\mathsf{L}_0, (k,\sigma)\big)$. Output $\tilde{\mathsf{L}}_0$.
- $g_0(\mathsf{R}_0)$**:** Draw $c \leftarrow \mathsf{E}_k(\mathsf{R}_0)$ and set $(\tilde{c}, \tilde{z}) = g(c, z)$. If $\tilde{z} \neq \tilde{z}^\$$, output $\perp_{\mathsf{com}}$. Otherwise, use superpolynomial time to break open $\tilde{z}$ and recover the pair $(\tilde{k}', \tilde{\sigma}')$. Output $\tilde{\mathsf{R}}_0 = \mathsf{D}_{\tilde{k}'}(\tilde{c})$.

Remark. Note that g_0 above does not run in polynomial time. It is possible to change the construction to use polynomial time (f_0, g_0) instead, however the proof would be longer and more difficult.

Definition 14. *Let* $\mathsf{L} = \big(\mathsf{L}_0, (k,\sigma)\big)$ *be a left state and let* $(f_0, g_0) \in \mathcal{F}^{\mathsf{poly}}_{\mathsf{split}}$ *be the tampering function pair obtained from* (f, g) *using randomness* (k, σ) *and* $c^\$$*. We say that* $\big(\mathsf{L}, (f_0, g_0)\big)$ *is* good *and write either* $\big(\mathsf{L}, (f_0, g_0)\big) \in \mathsf{GOOD}$ *or* $(\mathsf{L}, c^\$) \in \mathsf{GOOD}$ *if*

$$\mathrm{Decom}(\tilde{z}^\$) = (\tilde{k}, \tilde{\sigma}),$$

where $(\cdot, \tilde{z}^\$) = g\big(c^\$, \mathrm{Com}(k,\sigma)\big)$*, and* $\big(\cdot, (\tilde{k}, \tilde{\sigma})\big) = f(\mathsf{L})$.

If $\big(\mathsf{L}, (f_0, g_0)\big) \in \mathsf{GOOD}$ then tampering with (f, g) is the same as tampering with (f_0, g_0). To see this, note that $(\mathsf{R}, \tilde{m}) \leftarrow \mathsf{V}^{\mathsf{L}}_{m,f,g}$ has $\tilde{m} = \perp_{\mathsf{com}}$ if and only if $\mathrm{Decom}(\tilde{z}) \neq (\tilde{k}, \tilde{\sigma})$, or equivalently, since $(\mathsf{L}, c^\$) \in \mathsf{GOOD}$ and Com is perfectly binding, if $\tilde{z} \neq \tilde{z}^\$$. Moreover, if $\tilde{m} \neq \perp_{\mathsf{com}}$ then $g(\mathsf{R}) = \big(\mathsf{E}_{\tilde{k}}(g(\mathsf{R}_0)), \tilde{z}\big)$. It follows that if $\big(\mathsf{L}, (f_0, g_0)\big) \in \mathsf{GOOD}$, then $\mathsf{V}^{\mathsf{L}}_{m,f,g}$ is identical to the distribution: draw $(\mathsf{R}_0, \tilde{m}) \leftarrow \mathsf{V}^{\mathsf{L}_0}_{m,f_0,g_0}$, set $\mathsf{R} = \big(\mathsf{E}_k(\mathsf{R}_0), \mathrm{Com}(k,\sigma)\big)$ and output $(\mathsf{R}, \tilde{m})$.

We now define our simulators $\big\{\mathsf{S}^{\mathsf{L}}_{f,g}\big\}_{\mathsf{L}}$ for $\mathsf{L} \in \mathcal{L}_{\mathsf{VALID}}$: $\mathsf{S}^{\mathsf{L}}_{f,g}$ chooses $c^\$$ st $\big(\mathsf{L}, (f_0, g_0)\big) \in \mathsf{GOOD}$, then draws $(\mathsf{R}_0, \tilde{m}) \leftarrow \mathsf{S}^{\mathsf{L}_0}_{f_0,g_0}$, sets $c \leftarrow \mathsf{E}_k(\mathsf{R}_0)$, $z = \mathrm{Com}(k,\sigma)$ and outputs $(\mathsf{R}, \tilde{m})$ where $\mathsf{R} = (c, z)$. Defining $\mathsf{S}^{\mathsf{L}}_{f,g}$ this way is possible whenever there exists $c^\$$ such that $(\mathsf{L}, c^\$) \in \mathsf{GOOD}$.

Claim 10. *If* $\Pr_{\mathsf{L}}\big[\mathsf{L} \in \mathcal{L}_{\mathsf{VALID}}\big] \geq \delta$ *then* $\Pr_{\mathsf{L},c^\$}\big[(\mathsf{L}, c^\$) \in \mathsf{GOOD}\big] \geq \lambda^{-2}\delta^4$.

Proof Sketch. If $\mathsf{L} \in \mathcal{L}_{\mathsf{VALID}}$ then $\Pr_{\mathsf{R} \leftarrow \mathsf{Enc}(m^*|\mathsf{L})}\big(\mathrm{Decom}(\tilde{z}) = (\tilde{k}, \tilde{\sigma})\big) \geq \lambda^{-2}\delta^3$. It follows from the security of $(\mathsf{G}, \mathsf{E}, \mathsf{D})$ that $\Pr_{c^\$}\big(\mathrm{Decom}(\tilde{z}^\$) = (\tilde{k}, \tilde{\sigma})\big) \geq \lambda^{-2}\delta^3 - \mathsf{negl}$, and so a non-negligible fraction of $c^\$$ are so that $(\mathsf{L}, c^\$) \in \mathsf{GOOD}$. This can be formalized using an easy application of the hiding machine. □

Proof of Lemma 1. Fix non-negligible $\delta = \delta(\lambda) > 0$ and $(f, g) \in \mathcal{F}^{\mathsf{poly}}_{\mathsf{split}}$, let $\big\{\mathsf{S}^{\mathsf{L}}_{f,g}\big\}_{\mathsf{L}}$ be the family of simulators described above and let $\mathsf{E}_{f,g}$ be the event

$$\exists\, m \in \mathcal{M} \text{ st } \Big|\Pr_{(\mathsf{R},\tilde{m}) \leftarrow \mathsf{V}^{\mathsf{L}}_{m,f,g}}\big(\mathrm{D}(\mathsf{R}, \tilde{m}) = 1\big) - \Pr_{(\mathsf{R},\tilde{m}) \leftarrow \mathsf{S}^{\mathsf{L}}_{m,f,g}}\big(\mathrm{D}(\mathsf{R}, \tilde{m}) = 1\big)\Big| > \delta.$$

We must show that $\Pr_{\mathsf{L}}\big[\mathsf{E}_{f,g}\big] < 2\delta + \lambda^{-2}\delta^4 + 2\varepsilon$. We have

$$\begin{aligned}
\Pr_{\mathsf{L}}\big[\mathsf{E}_{f,g}\big] &= \Pr_{\mathsf{L}}\big[\mathsf{L} \in \mathcal{L}_{\mathsf{BOT}} \,\&\, \mathsf{E}_{f,g}\big] + \Pr_{\mathsf{L}}\big[\mathsf{L} \in \mathcal{L}_{\mathsf{VALID}} \,\&\, \mathsf{E}_{f,g}\big] \\
&< 2\delta + \Pr_{\mathsf{L}}\big[\mathsf{E}_{f,g} \big| |\mathcal{L}_{\mathsf{VALID}}| \geq \delta|\mathcal{L}|\big] \\
&\leq 2\delta + \lambda^{-2}\delta^4 + \Pr_{\mathsf{L}}\big[\mathsf{E}_{f,g} \big| \exists\, c^\$ \ (\mathsf{L}, c^\$) \in \mathsf{GOOD}\big] \\
&\leq 2\delta + \lambda^{-2}\delta^4 + \varepsilon + \Pr_{\mathsf{L}}\big[\exists\, m \in \mathcal{M} \text{ st } \Delta\big(\mathsf{V}^{\mathsf{L}_0}_{m,f_0,g_0}, \mathsf{S}^{\mathsf{L}_0}_{m,f_0,g_0}\big) > \varepsilon\big] \\
&< 2\delta + \lambda^{-2}\delta^4 + 2\varepsilon,
\end{aligned}$$

where the second line is from Claim 9, the fourth uses the above discussion about the ramifications of $(\mathsf{L}, c^\$) \in \mathsf{GOOD}$ and the last uses the $\varepsilon-$conditional augmented non-malleability of $(\mathsf{Enc}_0, \mathsf{Dec}_0)$. □

D Extras from Section 4

A Hiding Game For $\langle \mathrm{C}, \mathrm{R} \rangle$. Before proving Theorem 1, we specify a hiding game for $\langle \mathrm{C}, \mathrm{R} \rangle$, analogous to the hiding game for $(\mathsf{Enc}, \mathsf{Dec})$. We will use that this game is hard for a PPT adversary to win in the proof of non-malleability. Consider the following hiding game between a challenger $\mathcal{C}$ and a PPT adversary $\mathcal{A}$.

- **Partial Transcript:** $\mathcal{C}$ chooses $\mathsf{L} \leftarrow \mathcal{L}$, $r \leftarrow \mathbb{Z}_q$ and sends $\mathbf{Com}_{\mathsf{L}} = \big(\mathrm{Com}(\mathsf{L}), \mathrm{Com}(r)\big)$ to $\mathcal{A}$, $\mathcal{A}$ returns $\alpha \in \mathbb{Z}_q$ and receives $a = r\alpha + \mathsf{L}$ from $\mathcal{C}$.
- **Message Choice:** $\mathcal{A}$ chooses $m, m' \in \mathcal{M}$ and sends them to $\mathcal{C}$.
- **Challenge Message:** $\mathcal{C}$ chooses $\mathsf{R} \leftarrow \mathsf{Enc}(m|\mathsf{L})$, $b \leftarrow \{0,1\}$ and sends $(\mathsf{R}, \mathsf{R}_0, \mathsf{R}_1)$ to $\mathcal{A}$ where $\mathsf{R}_b \leftarrow \mathsf{Enc}(m|\mathsf{L})$ and $\mathsf{R}_{1-b} \leftarrow \mathsf{Enc}(m'|\mathsf{L})$.
- **Guess:** $\mathcal{A}$ outputs a guess $b' \in \{0,1\}$ and wins if $b' = b$.

Just as for the hiding game of the code $(\mathsf{Enc}, \mathsf{Dec})$, we will usually use an $(N, N')-$way variant of the above game, where the challenge message is $\big(\mathsf{R}, \{\mathsf{R}_1\}, \ldots, \{\mathsf{R}_N\}\big)$ where R is as in the basic game and each $^{\#}\{\mathsf{R}_i\} = N'$, and also $\mathsf{R}_i \leftarrow \mathsf{Enc}(m'|\mathsf{L})$ for all $\mathsf{R}_i \in \{\mathsf{R}_i\}$ and all i except for a random i^*, for which $\mathsf{R}_{i^*} \leftarrow \mathsf{Enc}(m|\mathsf{L})$ for all $\mathsf{R}_{i^*} \in \{\mathsf{R}_{i^*}\}$. In the $(N, N')-$way variant, $\mathcal{A}$ wins if he guesses i^*.

Claim 11. *If* Com *is computationally hiding and* $(\mathsf{Enc}, \mathsf{Dec})$ *has the hiding property then for all PPT adversaries* $\mathcal{A}$*, the probability that* $\mathcal{A}$ *wins the above game (resp. its* $(N, N')-$*way variant) is at most* $1/2 + \mathsf{negl}$ *(resp.* $1/N + \mathsf{negl}$*).*

Proof Sketch. The main difference between the above game and the hiding game for $(\mathsf{Enc}, \mathsf{Dec})$ is the partial transcript generation phase. Consider the version of the above with the bogus partial transcript phase: $\mathcal{C}$ chooses $\mathsf{L} \leftarrow \mathcal{L}$ but sends $\big(\mathrm{Com}(0), \mathrm{Com}(0)\big)$, and upon receiving $\alpha \in \mathbb{Z}_q$, returns a random $a \in \mathbb{Z}_q$. The bogus game is indistinguishable from the original by the hiding of Com. But any $\mathcal{A}$ who wins the bogus game can be used to construct an $\mathcal{A}'$ who wins the hiding game for $(\mathsf{Enc}, \mathsf{Dec})$ since $\mathcal{A}'$ can generate the bogus partial transcript himself and forward the other messages from his challenger to $\mathcal{A}$. □

D.1 Proof of Claims 4 and 5

Claim 4 (Restated). *For all* $m, m' \in \mathcal{M}$ *and non-negligible* $\xi > 0$ *we have:*

$$\mathrm{Pr}_{\tau}\left[\Big|p(\tau, m) - p(\tau, m')\Big| > \xi\right] = \mathsf{negl}.$$

Proof. We utilize the hiding machine. We prove the first formally and we will set up the machine and give the proof outline for the second. Fix $m, m' \in \mathcal{M}$ and non-negligible $\xi > 0$. Define $\mathsf{BAD} = \big\{\tau : p(\tau, m) > \xi + p(\tau, m')\big\}$ and suppose for contradiction that there is some non-negligible $\xi' > 0$ for which $\mathrm{Pr}\big(\tau \in \mathsf{BAD}\big) \geq \xi'$. Set $N = 4/(\xi\xi')$, $N' = \Omega\big(\lambda\xi^{-2}\big)$ and consider the PPT adversary $\mathcal{A}$ who interacts with $\mathcal{C}$ in the $(N, N')-$way hiding game for $\langle \mathrm{C}, \mathrm{R} \rangle$ as follows.

- $\mathcal{A}$ instantiates M. Upon receiving $\mathbf{Com}_{\mathsf{L}}$ from $\mathcal{C}$, it plays the first two rounds of $\langle \mathrm{C}, \mathrm{R} \rangle$ with M, giving input $\mathbf{Com}_{\mathsf{L}}$ and uniform $\tilde{\alpha}$ and receiving α in the second round of the left interaction. $\mathcal{A}$ sends α to $\mathcal{C}$ and receives a. This defines a partial transcript $\tau = \big(\mathbf{Com}_{\mathsf{L}}, \tilde{\alpha}, a\big)$.
- $\mathcal{A}$ sends (m, m') to $\mathcal{C}$ and receives challenge $\big(\mathsf{R}, \{\mathsf{R}_1\}, \ldots, \{\mathsf{R}_N\}\big)$. $\mathcal{A}$ forwards (a, R) to M and receives $(\tilde{a}, \tilde{\mathsf{R}})$. This defines a full transcript $\mathbb{T} = \mathbb{T}(\tau, \mathsf{R})$. Moreover, for all $i = 1, \ldots, N$ and $\mathsf{R}_i \in \{\mathsf{R}_i\}$, $\mathcal{A}$ sends (a, R_i) to M and receives $(\tilde{a}_i, \tilde{\mathsf{R}}_i)$, defining transcripts $\{\mathbb{T}_i\}$ for $i = 1, \ldots, N$, where $\mathbb{T}_i = \mathbb{T}(\tau, \mathsf{R}_i)$.
- $\mathcal{A}$ computes $p_i = \mathrm{Pr}_{\mathsf{R}_i \in \{\mathsf{R}_i\}}\big(\tilde{a}_i = \tilde{a}\big)$, and outputs i^* such that p_{i^*} is maximal.

Note that if $\tau \in \mathsf{BAD}$, then M's response $\tilde{a}$ in $\mathbb{T}$ is correct with probability at least ξ. Moreover, if $\tilde{a}$ is correct then $\tilde{v} \neq \perp_{\mathsf{inc}}$ in $\mathbb{T}_i$ if and only if $\tilde{a}_i = \tilde{a}$. Therefore, conditioned on $\tau \in \mathsf{BAD}$ and $\tilde{a}$ being correct, we see that

$$p_{i^*} \geq p(\tau, m) - \frac{\xi}{3} > p(\tau, m') + \xi - \frac{\xi}{3} \geq p_i + \xi - \frac{\xi}{3} - \frac{\xi}{3} = p_i + \frac{\xi}{3},$$

for all $i \neq i^*$ with probability at least $1 - 2^{-\Omega(\lambda)}$. We have used the Chernoff-Hoeffding bound, made possible by our choice of large enough N'. We conclude that

$$\Pr\big(\mathcal{A} \text{ wins}\big) \geq \xi\xi'(1 - \mathsf{negl}) > \frac{2}{N},$$

which violates the security of the hiding game for $\langle \mathrm{C}, \mathrm{R} \rangle$. □

Claim 5 (Restated). *Let $\delta = \delta(\lambda) > 0$ be as the statement of Lemma 3. For all $m, m' \in \mathcal{M}$, we have:*

$$\Pr_\tau\left[\tau \in \mathsf{MAUL}_{m,m'} \;\&\; p(\tau, m) < \lambda^{-2}\delta^3\right] \leq \frac{\delta}{4}.$$

Proof. This is another invocation of the hiding machine. Define BAD' to be $\tau \in \mathsf{MAUL}_{m,m'}$ and $p(\tau, m) < \lambda^{-2}\delta^3$, and suppose for contradiction that $\Pr_\tau\big[\tau \in \mathsf{BAD}'\big] > \delta/4$. By part one, whp over τ, if $\tau \in \mathsf{BAD}'$ then $p(\tau, m') < 2\lambda^{-2}\delta^3$ as well. Set $N = 17/\delta$, $N' = \Omega(\lambda\delta^{-2})$ and consider the $\mathcal{A}$ who instantiates M and plays with $\mathcal{C}$ as follows.

- $\mathcal{A}$ instantiates M and plays until he has partial transcript τ. Then $\mathcal{A}$ sends m, m' and receives $\big(\mathsf{R}, \{\mathsf{R}_1\}, \ldots, \{\mathsf{R}_N\}\big)$, and sets $\mathbb{T}_i = \mathbb{T}(\tau, \mathsf{R}_i)$, for all $\mathsf{R}_i \in \{\mathsf{R}_i\}$.
- $\mathcal{A}$ computes $p_i = \Pr_{\mathsf{R}_i \in \{\mathsf{R}_i\}}\big(\mathrm{D}^\tau(\mathsf{R}_i, \perp_{\mathsf{inc}}) = 1\big)$ and outputs i^* such that p_{i^*} is maximal.

If $\tau \in \mathsf{BAD}'$ then with probability at least 1/2, M will answer incorrectly in every $\mathbb{T}_i \in \{\mathbb{T}_i\}$ for all $i = 1, \ldots, N$, and so M's commitment in every $\mathbb{T}_i$ is to $\perp_{\mathsf{inc}}$. If this happens then p_i approximates $\Pr_{\mathsf{R} \leftarrow \mathsf{Enc}(m_i|\mathsf{L})}\big(\mathrm{D}^\tau(\mathsf{R}, \tilde{m}) = 1\big)$ where $m_{i^*} = m$ and $m_i = m'$ when $i \neq i^*$. In this case, we have by Chernoff-Hoeffding,

$$p_{i^*} \geq \Pr_{\mathsf{R} \leftarrow \mathsf{Enc}(m|\mathsf{L})}\big(\mathrm{D}^\tau(\mathsf{R}, \tilde{m})\big) - \frac{\delta}{6} > \Pr_{\mathsf{R} \leftarrow \mathsf{Enc}(m'|\mathsf{L})}\big(\mathrm{D}^\tau(\mathsf{R}, \tilde{m})\big) + \frac{\delta}{2} - \frac{\delta}{6} \geq p_i + \frac{\delta}{6},$$

for all $i \neq i^*$ with probability $1 - 2^{-\Omega(\lambda)}$. And so,

$$\Pr\big(\mathcal{A} \text{ wins}\big) \geq \frac{\delta}{8} - \mathsf{negl} > \frac{2}{N},$$

which violates the security of the hiding game for $\langle \mathrm{C}, \mathrm{R} \rangle$. □

E The Extended Protocol

In this section we modify the protocol of Section 4 so it remains non-malleable against a non-synchronizing adversary. The only non-synchronizing scheduling available to the adversary which is not trivially dealt with is the sequential scheduling where he lets the left interaction complete before beginning the right. Note this scheduling could not yeild a mauling attack against an extractable commitment, since this scheduling allows one to extract M's commitment without rewinding C. However, our protocol in Section 4 is not extractable. In this section we make it extractable while, and thus non-malleable against a sequential adversary, while still maintaining its non-malleability against a synchronizing adversary.

Let Π_{NM} be the commitment of Section 4 and let Π_{ext} be a (malleable) three round extractable commitment scheme. Our commitment scheme in this section commits to v by using Π_{NM} to commit to v while, in parallel, using Π_{ext} to commit to the decommitment information of the first part. We prove that this composition enjoys the best of both of its building blocks: it is extractable (and so non-malleable against a sequential adversary) while still being non-malleable against a synchronizing adversary. One technical point is that in the proof of synchronizing non-malleability, we need to rewind the protocol one time, therefore to make our proof go through, we need extraction from Π_{ext} to require two rewinds. The protocol is shown in Figure 2.

Setup: Let Com be a non-interactive, perfectly binding commitment scheme. Let $(\mathsf{Enc}, \mathsf{Dec})$ be a conditional non-malleable code with indistinguishable views. Fix a large prime q. Let $\mathsf{id} \in \{0,1\}^\lambda$ be C's identity.

Commiter's Private Input: $v \in \mathcal{M}_{\langle \mathrm{C},\mathrm{R}\rangle}$ to be committed to.

Commit Phase:

1. $\mathrm{C} \to \mathrm{R}$**:**
   - Set $m = v \circ \mathsf{id}$ and draw $(\mathsf{L}, \mathsf{R}) \leftarrow \mathsf{Enc}(m)$, where $\mathsf{L} \in \mathcal{L} \subset \mathbb{Z}_q$. Choose random $r \leftarrow \mathbb{Z}_q$ and $\omega, \omega' \leftarrow \$$ send $\mathrm{Com}(\mathsf{L}; \omega)$, $\mathrm{Com}(r; \omega')$ to R.
   - Set $\mathsf{X} = (\mathsf{L}, r, \omega, \omega') \in \{0,1\}^{\mathsf{poly}(\lambda)}$, viewed as a string and for $i = 1, \ldots, \lambda$ choose $\mathsf{X}_i^0, \mathsf{X}_i^1, \mathsf{X}_i^2 \in \{0,1\}^{\mathsf{poly}(\lambda)}$ randomly such that $\mathsf{X}_i^0 \oplus \mathsf{X}_i^1 \oplus \mathsf{X}_i^2 = \mathsf{X}$. Send $\mathsf{Y}_i^b = \mathrm{Com}(\mathsf{X}_i^b)$, to R for $b \in \{0,1,2\}$.
2. $\mathrm{R} \to \mathrm{C}$**:** Send random challenge $\alpha \in \mathbb{Z}_q$ and $\mathsf{c} \in \{0,1,2\}^\lambda$ to C.
3. $\mathrm{C} \to \mathrm{R}$**:** Send $a = r\alpha + \mathsf{L}$, R and $\mathrm{Decom}(\mathsf{Y}_i^{\mathsf{c}_i})$ for $i = 1, \ldots, \lambda$ to C.

Decommit Phase:

1. $\mathrm{C} \to \mathrm{R}$**:** Open the commitments sent in step 1. Let $\mathsf{L}', r' \in \mathbb{Z}_q$ and $\mathsf{X}_i^b \in \{0,1\}^{\mathsf{poly}(\lambda)}$ be the decommitted values.

Receiver's Output: If the strings $\mathsf{X}_i^0 \oplus \mathsf{X}_i^1 \oplus \mathsf{X}_i^2$ are not equal for all i or if they are all equal to $\mathsf{X} \in \{0,1\}^{\mathsf{poly}(\lambda)}$ but X is not a valid decommitment to $\mathrm{Com}(\mathsf{L}; \omega)$, $\mathrm{Com}(r; \omega')$ sent in step one output $\perp_{\mathsf{fail}}$. If L' and r' do not satisfy $r'\alpha + \mathsf{L}' = a$ then output $\perp_{\mathsf{inc}}$. Otherwise, compute $m' = \mathsf{Dec}(\mathsf{L}', \mathsf{R})$ and parse $m' = v' \circ \mathsf{id}'$. Output v' if $\mathsf{id}' = \mathsf{id}$, $\perp_{\mathsf{id}}$ if not.

Figure 2: Non-malleable commitment scheme $\langle \mathrm{C}, \mathrm{R}\rangle$.

Claim 12. $\langle \mathrm{C}, \mathrm{R}\rangle$ *is a perfectly binding extractable commitment scheme.*

Theorem 2. $\langle \mathrm{C}, \mathrm{R}\rangle$ *is non-malleable.*

Proof. It suffices to prove that it is non-malleable against a synchronizing adversary as the only other non-trivial scheduling is the sequential one and non-malleability against a sequential adversary follows from extractability. We follow the same outline and use the same notation as in the proof of Theorem 1. Recall that if M breaks the non-malleability of $\langle \mathrm{C}, \mathrm{R}\rangle$ then there exist $m, m' \in \mathcal{M}$ and non-negligible $\delta > 0$ such that $\Pr_\tau\left[\tau \in \mathsf{MAUL}_{m,m'}\right] \geq \delta/2$, where $\tau \in \mathsf{MAUL}_{m,m'}$ if

$$\left|\Pr\left(D_{m,m'}^\tau\left(\mathsf{M}_m^\tau\right) = 1\right) - \Pr\left(D_{m,m'}^\tau\left(\mathsf{M}_{m'}^\tau\right) = 1\right)\right| \geq \frac{\delta}{2}.$$

The partial transcript τ includes the first two rounds of the right execution and the entire left execution except for R. As before, τ is completed to a full transcript once R is specified. Note that if M's commitment $\tilde{v}$ is to $\perp_{\mathsf{fail}}$, then M's first message is bad, either because the strings $\tilde{\mathsf{X}}_i^0 \oplus \tilde{\mathsf{X}}_i^1 \oplus \tilde{\mathsf{X}}_i^2$ are not all equal, or because they are not all equal to a valid decommitment to $\mathrm{Com}(\tilde{\mathsf{L}}; \tilde{\omega})$, $\mathrm{Com}(\tilde{r}; \tilde{\omega}')$. In any case, M cannot be mauling if $\tilde{v} = \perp_{\mathsf{fail}}$ as C's commitment on the left is not even defined after the first message. Also, the proof of Lemma 2 goes through unchanged for this new protocol and so it follows that if M is mauling with non-negligible probability then he is doing so while also sending the correct value for $\tilde{a}$ with non-negligible probability.

This allows us to build a distribution $\mathcal{D}_{\mathrm{M}}$ on tampering functions which we will use to break the security of $(\mathsf{Enc}, \mathsf{Dec})$. As before $(f, g) \leftarrow \mathcal{D}_{\mathrm{M}}$ share a random partial transcript τ and a random $\mathsf{R}_{\$}$, let $\tilde{a}_{\$}$ be M's response on the right in the transcript $\mathbb{T}(\tau, \mathsf{R}_{\$})$. $f(\mathsf{L})$ extracts $\tilde{\mathsf{L}}$ by rewinding M and asking a new challenge $\tilde{\beta}$ on the right, using L and τ to answer on the left. We provide f with the decommitments $\mathrm{Decom}(\mathsf{Y}_i^b)$ so he can answer this part of M's query on the left honestly. $g(\mathsf{R})$ completes τ to $\mathbb{T}(\tau, \mathsf{R})$ and checks whether M's answer $\tilde{a}$ is equal to $\tilde{a}_{\$}$ or not. If so $g(\mathsf{R}) = \tilde{\mathsf{R}}$, where $\tilde{\mathsf{R}}$ is from M's final message of $\mathbb{T}(\tau, \mathsf{R})$. If not $g(\mathsf{R}) = \perp_{\mathsf{inc}}$. The same proof of Lemma 3 shows that if M breaks the non-malleability of $\langle \mathrm{C}, \mathrm{R} \rangle$ then there exist $m, m' \in \mathcal{M}$ and non-negligible $\delta, \delta' > 0$ such that

$$\mathrm{Pr}_{\mathsf{L},(f,g)}\left[\left|\Pr\left(\mathrm{D}^{\tau}(\mathsf{V}^{\mathsf{L}}_{m,f,g}) = 1\right) - \Pr\left(\mathrm{D}^{\tau}(\mathsf{V}^{\mathsf{L}}_{m',f,g}) = 1\right)\right| > \frac{\delta}{2}\right] > \delta', \tag{4}$$

where the outer probability is over $\mathsf{L} \leftarrow \mathcal{L}$ and $(f, g) \leftarrow \mathcal{D}_{\mathrm{M}}$ using randomness which depends on L. In order to complete the reduction to the security of $(\mathsf{Enc}, \mathsf{Dec})$ we need to exhibbit a distribution $\mathcal{D}_{\mathsf{split}}$ on split state tampering function pairs such that

$$\mathrm{Pr}_{\mathsf{L},(f,g)}\left[\left|\Pr\left(\mathrm{D}^{\tau}(\mathsf{V}^{\mathsf{L}}_{m,f,g}) = 1\right) - \Pr\left(\mathrm{D}^{\tau}(\mathsf{V}^{\mathsf{L}}_{m',f,g}) = 1\right)\right| > \frac{\delta}{2}\right] > \delta', \tag{5}$$

where the outer probability is over $\mathsf{L} \leftarrow \mathcal{L}$ and $(f, g) \leftarrow \mathcal{D}_{\mathsf{split}}$ drawn independently. In the proof of Theorem 1, we argued that $\mathcal{D}_{\mathrm{M}} \approx_{\mathsf{c}} \mathcal{D}_{\mathsf{split}}$ and so (5) followed straight from (4). Here, this won't quite work because f gets every $\mathrm{Decom}(\mathsf{Y}_i^b)$ so it can answer M's query in the rewind. This means f knows L and so we will not be able to hybrid away f's dependence on L without f noticing. However, the following observation lets us deduce (5) anyway: as f only rewinds once its output contains no information about $\mathsf{X} = \mathsf{X}_i^0 \oplus \mathsf{X}_i^1 \oplus \mathsf{X}_i^2$. $\square$