Electronic Colloquium on Computational Complexity, Revision 1 of Report No. 117 (2019)

Locally Testable Non-Malleable Codes*

Silas Richelson† Sourya Roy‡

Abstract

In this work we adapt the notion of non-malleability for codes of Dziembowski, Pietrzak and Wichs (ICS 2010) to locally testable codes. Roughly speaking, a locally testable code is non-malleable if any tampered codeword which passes the local test with good probability is close to a valid codeword which either encodes the original, or an unrelated message.

We instantiate our definition by proving that a Reed-Muller-type code is non-malleable in the following sense: any adversary who independently tampers the coordinates of the code so that the tampered code passes the test with good probability, is tampering the underlying polynomial according to an affine transformation.

To the best of our knowledge, prior to this work, polynomial codes were not known to possess any non-malleability guarantees. Our analysis builds on the sampler-based decoding techniques common to several recent works.

As an additional contribution, we describe a new (standard) non-malleable code against affine tampering which is much simpler than prior constructions, and achieves better parameters. Finally, we prove a composition theorem for locally testable non-malleable codes which allows for obtaining codes via concatenation.

1 Introduction

A *coding scheme* is a pair $(\mathsf{Enc}, \mathsf{Dec})$ of functions $\mathsf{Enc} : \Gamma^k \to \Gamma^n$ (possibly randomized) and $\mathsf{Dec} : \Gamma^n \to \Gamma^k \cup \{\perp\}$ such that $\mathsf{Dec}\big(\mathsf{Enc}(m)\big) = m$ holds with probability 1 for all $m \in \Gamma^k$. We say $\mathbf{x} \in \Gamma^n$ is a *valid codeword* if $\mathbf{x} = \mathsf{Enc}(m)$ for some $m \in \Gamma^k$ (and some choice of randomness for Enc). The quantity k/n is called the *rate* of the code. Given $\mathbf{x}, \mathbf{y} \in \Gamma^n$, the *distance* between $\mathbf{x}$ and $\mathbf{y}$ is $\Pr_{i\sim[n]}\big[\mathbf{x}_i \neq \mathbf{y}_i\big]$. The *distance of the code* is the minimum distance between any two distinct valid codewords. When a code's distance is bounded away from zero, one can design decoding algorithms with error-correction capabilities. We say $(\mathsf{Enc}, \mathsf{Dec})$ is an *error-correcting code* [Sha49, Ham50] if there exists $\delta > 0$ such that $\mathsf{Dec}(\mathbf{y}) = m$ for all $\mathbf{y}$ which are within distance δ of some valid codeword $\mathbf{x} = \mathsf{Enc}(m)$. Error-correcting codes are incredibly useful objects in both theory and practice. Thanks to an extensive research effort over the past 70 years, the theory of error-correcting codes has advanced to the point that many of the first order questions

*Supported by UC Lab Fees grant LFR-18-548554. All opinions expressed are those of the authors.

†UC Riverside. Email: silas@cs.ucr.edu.

‡UC Riverside. Email: sourya.roy@email.ucr.edu.

ISSN 1433-8092

have been answered. For example, at this point it is known how to construct codes with constant rate and which can decode from a constant fraction of errors [RS60, Jus72].

Most of the time, $\mathsf{Dec}(\mathbf{y})$ – in addition to outputting the message underlying the nearest valid codeword to $\mathbf{y}$ – will detect whether $\mathbf{y}$ itself is a valid codeword, and if not, identify the incorrect symbols. *Locally-testable codes* (LTCs) [FS95, GS06] support a very efficient, randomized test of this type which reads only a few (usually a constant number of) symbols from the code and outputs a bit indicating whether or not it thinks the codeword is valid. Roughly speaking, the requirement is that for all $\mathbf{y} \in \Gamma^n$, the probability that $\mathsf{Test}(\mathbf{y}) = 0$ should be proportional to the distance between $\mathbf{y}$ and the nearest valid codeword. So $\mathsf{Test}(\mathbf{x}) = 1$ with probability 1 for all valid codewords $\mathbf{x} \in \Gamma^n$; and if $\mathbf{y}$ is very far from being valid, then $\mathsf{Test}(\mathbf{y}) = 0$ should occur with high probability. Interest in LTCs is derived from their connection to probabilistically checkable proofs [ALM+98, AS98] and to property testing [BLR93, RS96].

Non-malleable codes (NMCs) [DPW18] provide security against a channel which, rather than being honest but noisy, actively tampers codewords using a function $f : \Gamma^n \to \Gamma^n$. This model was initially motivated by applications to leakage and tamper resilient cryptography [DPW18, AGM+15, CDTV16]. However, since their introduction, NMCs have found numerous other applications, for example to secure protocol design [GPR16, GR19], complexity theory [DJMW12], and pseudorandomness [CGL16, CZ16]. Given a message $m \in \Gamma^k$ and $f : \Gamma^n \to \Gamma^n$, the *tampering distribution* outputs $\big(\mathsf{Dec} \circ f \circ \mathsf{Enc}\big)(m) \in \Gamma^k \cup \{\perp\}$. Roughly speaking, we say that $(\mathsf{Enc}, \mathsf{Dec})$ is *non-malleable* against a function family $\mathcal{F} \subset \big\{f : \Gamma^n \to \Gamma^n\big\}$ if for all $f \in \mathcal{F}$ and $m \in \Gamma^k$, the tampering distribution either outputs m (such is the case when f is the identity) or is statistically independent of m (such is the case when f is a constant function).

1.1 Our Contributions

LTCs and NMCs represent two generalizations of error-correcting codes along different axes. In this work, we combine the notions and define *locally testable, non-malleable codes* (LTNMCs). Roughly speaking, a LTNMC is a LTC which has the following non-malleability guarantee: any tampered codeword which passes the test with good probability is close to a valid codeword which either encodes m, or else encodes an unrelated message. We then instantiate our notion by proving that a Reed-Muller-type code is non-malleable against the family of coordinate-wise tampering functions. Our construction has three parts.

1. We prove that when the Reed-Muller-type LTC of Raz and Safra [RS97] (*i.e.*, the "planes table") is tampered by a coordinate-wise tampering function then either the tampered codeword is far from a valid codeword (and so fails the local test with high probability) or else is close to a valid codeword which encodes an affine function of the original message. In NMC terminology, we show that the planes table is a *non-malleable reduction* from coordinate-wise tampering to affine tampering.

2. We describe an elementary construction of a (standard) NMC against the family of affine tampering functions. Such codes were previously known [ADL14, Agg15, CL17], but our construction is much simpler those in prior work. When the message space is large, our construction is more efficient than the one in [ADL14, Agg15] as our encoding algorithm

does not require drawing large random primes. Our code achieves a better rate/error tradeoff than the construction of [CL17].

3. We combine the codes from the above points into a concatenated code, obtaining a LTNMC against coordinate-wise tampering via a composition theorem. The local test of our composed code works by decoding a symbol of the outer code and checking validity using the inner code. This idea has been used previously to analyze the composition of LTCs and PCPs [AS98].

1.2 Motivation

LTCs and NMCs are both well studied objects from coding theory. As mentioned, both are generalizations of error-correcting codes (ECCs) − LTCs extend ECCs by adding efficient testing capabilities, while NMCs extend ECCs by protecting against more sophisticated forms of tampering. So LTCs are ECCs with *more functionality* while NMCs are ECCs with *better security*. For this reason, it is natural to ask whether these two extensions can be achieved simultaneously. Rather than constructing a special purpose code which is locally testable and non-malleable, we show that one of the most widely-used LTCs already possesses non-malleability properties. Thus, our result might have applications in the areas of theoretical computer science where these codes are commonly used.

For example, although non-malleable proofs have been studied in cryptography since [DDN91], non-malleable PCPs have not to our knowledge been defined or constructed. LTNMCs might be the "combinatorial analogues" of non-malleable PCPs, just like LTCs are analogues of standard PCPs. Thus our work could be a stepping stone towards new PCP constructions, and more optimistically, new hardness of approximation results. We elaborate below on a specific potential application of our result which we believe is a promising research direction. In fact, it was by pursuing this direction that this paper originated.

Split-State NMCs with Optimal Parameters. In the split-state model, the encoding algorithm splits the message into a codeword with two parts,[1] $\mathsf{Enc} : \Gamma^k \to \Gamma^n \times \Gamma^n$, and the tampering function acts on the parts independently: $\mathcal{F}_{\mathsf{split}} := \big\{(f,g)\big| f,g : \Gamma^n \to \Gamma^n\big\}$. So the tampering distribution for $(f,g) \in \mathcal{F}_{\mathsf{split}}$ and $m \in \Gamma^k$, outputs:

$$\big(\mathsf{Dec} \circ (f,g) \circ \mathsf{Enc}\big)(m) = \mathsf{Dec}\big(f(\mathsf{L}), g(\mathsf{R})\big) = \mathsf{Dec}(\tilde{\mathsf{L}}, \tilde{\mathsf{R}}) = \tilde{m}, \tag{1}$$

where $\mathsf{L}, \mathsf{R} \in \Gamma^n$ and $(\tilde{\mathsf{L}}, \tilde{\mathsf{R}}) = \big(f(\mathsf{L}), g(\mathsf{R})\big)$. Let us say that $(\mathsf{Enc}, \mathsf{Dec})$ is an $[n,k,\varepsilon]^{\Gamma}_{\mathsf{NM}}$−code if it has $\mathsf{Enc} : \Gamma^k \to \Gamma^n \times \Gamma^n$, and is ε−non-malleable against $\mathcal{F}_{\mathsf{split}}$ (omit Γ if $\Gamma = \{0,1\}$).

Split-state codes have a tremendous number of applications. The works [ADKO15, AKO17, BDG+18, ADN+19, BGW19] (and more) use split-state non-malleable codes to construct codes which are secure against more sophisticated classes of adversarial tampering. Additionally, several applications of split-state codes have been found outside of coding theory [GPR16, CGL16, CZ16, GR19] (and more). For most of these examples, an optimal construction of split-state non-malleable codes would yield improvements in the applications.

[1] In general $\mathsf{Enc} : \Gamma^k \to \mathcal{L} \times \mathcal{R}$ with $\mathcal{L}$ and $\mathcal{R}$ different. We enforce $\mathcal{L} = \mathcal{R} = \Gamma^n$ for simplicity; note the rate is $k/2n$.

The split-state model for non-malleable codes was conceived in [DPW18], where it was proved (non-constructively) that $[n, k, \varepsilon]_{\mathsf{NM}}-$codes exist with $n = \mathcal{O}(k)$ and $\varepsilon = 2^{-\Omega(k)}$. Numerous explicit constructions followed [DKO13, ADL14, ADKO15, CGL16, Li17] (and more). The state of the art today is represented by two works [Li19, AO19]. Li constructs an $[n, k, \varepsilon]_{\mathsf{NM}}-$code with $n = \mathcal{O}\big(k \cdot \frac{\log k}{\log\log k}\big)$, $\varepsilon = 2^{-\Omega(k)}$; Aggarwal and Obremski get $n = \mathcal{O}(k)$ and $\varepsilon = 2^{-k^{\alpha}}$ for a constant $0 < \alpha < 1$. Both constructions use the "alternating extraction" technique of [DP07], which is both a very powerful method for proving non-malleability, and the source of the sub-optimalities. For this reason, it seems as though new ideas will be required in order to obtain optimal split-state non-malleable codes.

One such idea is to build a split-state code using a LTC with extra properties. Specifically, suppose that $(\mathsf{Enc}, \mathsf{Dec})$ is a LTC with $\mathsf{Enc} : \Gamma^k \to \Gamma^n$. Additionally, assume that the local testing algorithm has a decoding feature where it takes an additional input $j \in [k]$ and $\mathsf{Test}(\mathbf{x}, j)$ reads two symbols, say $(\mathbf{x}_{i_{\mathsf{L}}}, \mathbf{x}_{i_{\mathsf{R}}})$ and outputs either $\perp$ if the test fails or $m_j \in \Gamma$, the $j-$th symbol of the encoded message, if the test passes. Given such a code, define the split state code $(\mathsf{Enc}', \mathsf{Dec}')$ as follows.

- $\mathsf{Enc}'(m)$**:** Given $m \in \Gamma$, compute $\mathbf{x} \sim \mathsf{Enc}(m, 0, \ldots, 0)$ where $(m, 0, \ldots, 0) \in \Gamma^k$, draw $(i_{\mathsf{L}}, i_{\mathsf{R}}) \in [n]^2$ according to $\mathsf{Test}(\mathbf{x}, 1)$, and output $(\mathsf{L}, \mathsf{R}) = \big((i_{\mathsf{L}}, \mathbf{x}_{i_{\mathsf{L}}}), (i_{\mathsf{R}}, \mathbf{x}_{i_{\mathsf{R}}})\big)$.
- $\mathsf{Dec}'(\mathsf{L}, \mathsf{R})$**:** Parse $(\mathsf{L}, \mathsf{R}) = \big((i_{\mathsf{L}}, \mathbf{x}_{i_{\mathsf{L}}}), (i_{\mathsf{R}}, \mathbf{x}_{i_{\mathsf{R}}})\big)$, and output what the test outputs (either $\perp$ or m).

Note the rate of $(\mathsf{Enc}', \mathsf{Dec}')$ is $\frac{1}{2} \cdot \log|\Gamma| / \big(\log n + \log|\Gamma|\big)$ which is constant whenever $n = |\Gamma|^{\mathcal{O}(1)}$. The list of requirements on the LTC in order for $(\mathsf{Enc}', \mathsf{Dec}')$ to be non-malleable is extensive though plausible. Our result handles the special tampering case when f tampers $(i_{\mathsf{L}}, \mathbf{x}_{i_{\mathsf{L}}})$ to $(i_{\mathsf{L}}, \tilde{\mathbf{x}}_{i_{\mathsf{L}}})$, keeping i_{L} fixed, and similarly when g fixes i_{R}. Indeed, such tampering of $(\mathsf{Enc}', \mathsf{Dec}')$ corresponds to coordinate-wise tampering of the LTC. We obtain the optimal non-malleability error of $|\Gamma|^{-\Omega(k)}$ in this case.

1.3 Other Prior Work

Sampler-Based Decoding. Our work fits into a recent line of work on sampler-based decoding [IKW12, Mos17, BDN17, DHK+19, DHKR19] (and more). In these works, sampling properties of a code's index set are exploited in order to give non-trivial decoding algorithms. Our work builds on techniques developed in these papers in order to "decode" a coordinate-wise tampering function which respects codeword proximity, to a small list of affine functions.

Locally Decodable Non-Malleable Codes. A few works combine the notions of local decodability with non-malleability [DLSZ15, CKR15, DLSZ20]. These works give constructions of non-malleable codes which admit local decode/update subroutines. Our work differs in several ways from these. First, the codes in these works achieve high rate with super-constant locality, whereas our main construction achieves optimal locality with very poor rate. Moreover, our techniques differ significantly. The constructions in prior work achieve local decodeability and updateability by separately encoding each element of the message, they do not support a local test of proximity to a valid codeword. Our techniques on the other hand, are similar to those used in the LTC literature.

1.4 Technical Overview

In this technical overview we focus on our first contribution − the proof that any coordinate-wise tampering of the planes-table code results either in a codeword which fails the local test with high probability, or else in a codeword which encodes an affine function of the original message. Let $k, d \in \mathbb{N}$ be integers with $k \geq 5$ and $d \geq 2$, and let $\mathbb{F}$ be a finite field. The code we analyze encodes a field element $m \in \mathbb{F}$ by choosing a random $k-$variate, degree d polynomial Φ such that $\Phi(\mathbf{0}) = m$, then for each $3-$plane $\mathsf{a} \subset \mathbb{F}^k$, the $\mathsf{a}-$th codeword symbol is the restriction of Φ to a, $\alpha = \Phi|_{\mathsf{a}}$. Note α is a $3-$variate polynomial of degree at most d. This code was shown to be a LTC in [RS97] where the testing algorithm chooses a random point $\mathsf{c} \in \mathbb{F}^k$ and two random planes $\mathsf{a}, \mathsf{a}' \subset \mathbb{F}^k$ containing c, reads the $\mathsf{a}-$th and $\mathsf{a}'-$th symbols α and α' and outputs 1 if $\alpha|_{\mathsf{c}} = \alpha'|_{\mathsf{c}}$ and 0 otherwise. Here, $\alpha|_{\mathsf{c}}$ denotes the evaluation of the $3-$variate polynomial α at the point c. The tampering function family we consider is the family of coordinate-wise tampering functions, $\mathcal{F} = \left\{\{\mathsf{f}_{\mathsf{a}}\}_{\mathsf{a}\subset\mathbb{F}^k}\right\}$ where for each $3-$plane $\mathsf{a} \subset \mathbb{F}^k$, f_{a} is an arbitrary function mapping $3-$variate, degree d polynomials to $3-$variate, degree d polynomials. We write $\tilde{\alpha} = \mathsf{f}_{\mathsf{a}}(\alpha)$.

Given $m \in \mathbb{F}$ and $\{\mathsf{f}_{\mathsf{a}}\} \in \mathcal{F}$, the tampering distribution chooses a random $k-$variate Φ of degree d such that $\Phi(\mathbf{0}) = m$, and for each $3-$plane $\mathsf{a} \subset \mathbb{F}^k$, tampers to obtain $\tilde{\alpha} = \mathsf{f}_{\mathsf{a}}(\alpha)$ where $\alpha = \Phi|_{\mathsf{a}}$; the distribution outputs $\{(\mathsf{a}, \tilde{\alpha})\}_{\mathsf{a}\subset\mathbb{F}^k}$. We must show that either $\{(\mathsf{a}, \tilde{\alpha})\}$ fails the local test with high probability, or else is close to an encoding of an affine function of m. Since the code is a LTC, if $\{(\mathsf{a}, \tilde{\alpha})\}$ passes the local test with good probability, then there exists a $k-$variate polynomial $\tilde{\Phi}$ of degree at most d such that $\tilde{\alpha} = \tilde{\Phi}|_{\mathsf{a}}$ holds for a good fraction of the $\mathsf{a} \subset \mathbb{F}^t$. Our main theorem says, intuitively, that for all $\{\mathsf{f}_{\mathsf{a}}\} \in \mathcal{F}$ there exists an affine map T on $k-$variate, degree d polynomials such that $\tilde{\Phi} = \mathsf{T}(\Phi)$. This means that a good fraction of the coordinate-wise functions are restrictions of a global affine function. The result of such tampering on the message is that $\tilde{m}$ is an affine function of m.

So we now zoom in on the proof of the main theorem: that for all $\{\mathsf{f}_{\mathsf{a}}\} \in \mathcal{F}$ such that the output of the tampering experiment passes the local test with good probability, there exists an affine map T such that a good fraction of the f_{a} agree with T. One can think of the data $\{\mathsf{f}_{\mathsf{a}}\}$ as assigning $\tilde{\alpha} = \mathsf{f}_{\mathsf{a}}(\alpha)$, a $3-$variate low degree polynomial, to the plane/polynomial pair (a, α). Thus $\{\mathsf{f}_{\mathsf{a}}\}$ is a planes/polynomials table, *i.e.*, it is like the planes table from [RS97] which assigns a polynomial to each plane, except that the index set now consists of all plane/polynomial pairs. Thus our main theorem involves analyzing a low-degree test. We do this in two steps.

The first part is similar to the low-degree theorem of [BDN17]. We show that if $\{\mathsf{f}_{\mathsf{a}}\}$ is such that the tampering experiment passes the local test with non-negligible probability, then the index set of the planes table splits into a few "almost cliques". These are subsets of the index set with non-negligible weight where agreement within the set holds with high probability. Typically this argument uses the pseudorandomness properties of the index set. For example, [BDN17] makes heavy use of the fact that the planes/points "incidence graph"[2] is a good sampler. In our context, we will need that the "incidence×agreement" graph is a good sampler. This is the bipartite graph $G = (A \cup B, E)$ where A is the set of pairs (a, α) where $\mathsf{a} \subset \mathbb{F}^k$ is a $3-$plane and α a $3-$variate, degree d polynomial; $B = \mathbb{F}^k \times \mathbb{F}$ and $\big((\mathsf{a}, \alpha), (\mathsf{c}, \gamma)\big) \in$ iff $\mathsf{c} \in \mathsf{a}$ and $\alpha|_{\mathsf{c}} = \gamma$. We establish the sampling properties we need in Section 5.

Once we know that the set of plane/polynomial pairs separates into cliques we focus in on

[2]This is the bipartite graph $G = (A \cup B, E)$ where A is the set of planes in $\mathbb{F}^k$, $B = \mathbb{F}^k$ and $(\mathsf{a}, \mathsf{c}) \in E$ iff $\mathsf{c} \in \mathsf{a}$.

one of these cliques and we use a linearity test of [RS96] to show that there is a global affine map which many of the f_a agree with. It is key that we work within a clique because this puts us in the "low error" regime rather than the "high error" one where we originally started.

2 Preliminaries

2.1 Locally Testable Codes and Non-Malleable Codes

Definition 1 (Locally Testable Code). *Fix $q \in \mathbb{N}$ and $\varepsilon > 0$. We say that a code $(\mathsf{Enc}, \mathsf{Dec})$, is a $(q, \varepsilon)-$*locally testable code *(LTC) if there exists a randomized algorithm Test which reads q symbols of a supposed codeword $\mathbf{y} \in \Gamma^n$ (the symbols are indexed by $I \subset [n]$ of size $|I| = q$) and outputs a bit such that 1) $\mathsf{Test}(\mathbf{x}) = 1$ with probability 1 for all valid codewords $\mathbf{x} \in \Gamma^n$, and 2) there exists a constant $c > 0$ such that for all $\mathbf{y} \in \Gamma^n$ with $\mathsf{dist}(\mathbf{y}) \geq \varepsilon$,*

$$\Pr_I\Big[\mathsf{Test}(\mathbf{y}; I) = 0\Big] \geq c \cdot \mathsf{dist}(\mathbf{y}),$$

where $\mathsf{dist}(\mathbf{y})$ denotes the distance between $\mathbf{y}$ and the nearest valid codeword.

Intuitively, the second point says that if $\mathbf{y} \in \Gamma^n$ is such that $\mathsf{Test}(\mathbf{y}) = 1$ with non-negligible probability, then $\mathbf{y}$ has non-negligible agreement with a valid codeword. List decoding for LTCs refers to the stronger guarantee: for any $\mathbf{y} \in \Gamma^n$, there is a short list of valid codewords which explain nearly all of $\mathsf{Test}(\mathbf{y})$'s acceptance probability.

Definition 2 (List-Decoding for LTCs). *Fix $\ell \in \mathbb{N}$ and $\varepsilon > 0$. A locally testable code is said to be $(\ell, \varepsilon)-$*list-decodable *if for all $\mathbf{y} \in \Gamma^n$ there exists a set $\mathsf{L}_\mathbf{y} \subset \Gamma^n$ of valid codewords such that $|\mathsf{L}_\mathbf{y}| \leq \ell$ and*

$$\Pr_I\Big[\mathsf{Test}(\mathbf{y}; I) = 1 \;\&\; \mathbf{y}_I \neq \mathbf{x}_I \;\forall\, \mathbf{x} \in \mathsf{L}_\mathbf{y}\Big] < \varepsilon.$$

Non-malleable codes [DPW18] (NMCs) provide meaningful security guarantees even in situations where error correction is impossible. Non-malleable reductions [ADKO15] are useful relaxations which allow constructing non-malleable codes via concatenation. Intuitively, a non-malleable reduction from $\mathcal{F}$ to $\mathcal{G}$ guarantees that the tampering of codewords by functions in $\mathcal{F}$ is captured by tampering messages by functions in $\mathcal{G}$. The key feature of non-malleable reductions is that they compose well. For example, if $(\mathsf{Enc}_\mathcal{F}, \mathsf{Dec}_\mathcal{F})$ is a non-malleable reduction from $\mathcal{F}$ to $\mathcal{G}$ and $(\mathsf{Enc}_\mathcal{G}, \mathsf{Dec}_\mathcal{G})$ is a non-malleable code against $\mathcal{G}$, then $\big(\mathsf{Enc}_\mathcal{F} \circ \mathsf{Enc}_\mathcal{G}, \mathsf{Dec}_\mathcal{G} \circ \mathsf{Dec}_\mathcal{F}\big)$ is a non-malleable code against $\mathcal{F}$.

Definition 3 (Non-Malleable Reductions). *Fix $\varepsilon > 0$ and tampering function families*

$$\mathcal{F} \subset \big\{f : \Gamma^n \to \Gamma^n\big\} \text{ and } \mathcal{G} \subset \big\{g : \Gamma^k \to \Gamma^k \cup \{\perp\}\big\}.$$

*We say that a coding scheme $(\mathsf{Enc}, \mathsf{Dec})$ is an $\varepsilon-$*non-malleable reduction from $\mathcal{F}$ to $\mathcal{G}$ *if for all $f \in \mathcal{F}$ there exists a distribution G_f on $\mathcal{G}$ such that $\Delta\big((\mathsf{Dec} \circ f \circ \mathsf{Enc})(m), \mathsf{G}_f(m)\big) \leq \varepsilon$ for all $m \in \Gamma^k$, where $\mathsf{G}_f(m)$ is the distribution which draws $g \sim \mathsf{G}_f$ and outputs $g(m)$ (Δ denotes statistical distance). A* non-malleable code *is a non-malleable reduction to the family of "trivial" tampering functions, containing only the identity and constants.*

Remark. We have chosen to define NM codes via NM reductions because it is syntactically simpler and much more intuitive for readers who are not already experts in non-malleability. It was one of the main contributions of [ADKO15] that non-malleable codes could be defined this way. We refer readers to the original definition of [DPW18] and the discussion in [ADKO15] for more details.

2.2 Sampler Graphs

Notations. For a finite set S, $s \sim S$ indicates that s is drawn uniformly from S. For a bipartite graph $(A \cup B, E)$ and $a \in A$, $B(a)$ denotes the uniform distribution on the neighborhood of a in B: $\{b \in B : (a,b) \in E\}$. The neighborhood distribution $A(b)$ for $b \in B$ is defined analogously. For all bipartite graphs used in this work, the edge relations are natural. For example, A might be the set of lines in $\mathbb{F}^k$ ($\mathbb{F}$ a finite field), B the set of points in $\mathbb{F}^k$, and the edge relation captures incidence: $(a,b) \in E$ iff $b \in a$. For this reason, we simplify notations by supressing E and denoting bipartite graphs as A/B instead of $(A \cup B, E)$, and writing $a \sim b$ instead of $(a,b) \in E$.

Definition 4 (Biregularity). *Let A/B be a bipartite graph and fix $\eta > 0$. We say that A/B is $\eta-$*biregular *if the distribution which draws $a \sim A$, $b \sim B(a)$, and outputs (a,b) is within statistical distance η of the distribution which gives the same output by drawing $b \sim B$, $a \sim A(b)$.*[3]

Biregularity ensures that for any $B' \subset B$ of size $|B'| = \lambda \cdot |B|$, the expectation (over $a \sim A$) of $\Pr_{b \sim B(a)}[b \in B']$ is close to λ. We say that A/B is *sampling* if a concentration bound holds.

Definition 5 (Sampler Graph [Zuc97]**).** *Fix $\varepsilon, \delta > 0$. We say that the bipartite graph A/B is $(\varepsilon,\delta)-$*sampling *if for all subsets $B' \subset B$ of size $|B'| = \lambda \cdot |B|$,*

$$\Pr_{a \sim A}\left[\left|\Pr_{b \sim B(a)}\left[b \in B'\right] - \lambda\right| > \varepsilon\right] \leq \delta.$$

Double Samplers. A triple (A,B,C) is called a *double sampler* if B/C is sampling and for all $c \in C$, $A(c)/B(c)$ is sampling. Double samplers have been used implicitly in several works prior to their formalization in [DK17]. We use them implicitly in this work as well. The construction in [DK17] is of a double sampler of linear size (*i.e.*, $|A| \approx |B| \approx |C|$) based on high-dimensional expanders. The double samplers used in this work are built from elementary means and are not linear size (our double samplers have $|A| \gg |B| \gg |C|$). Importantly, a random object in our parameter regime is a double sampler with good probability, while this is not true in the linear size regime.

Fact 1 (Properties of Samplers). *Suppose A/B is $\eta-$biregular and $(\varepsilon,\delta)-$sampling. We have the following.*

1. *For any $\rho > 0$ and $f : B \to [0,1]$,*

$$\Pr_{a \sim A}\left[\left|\mathbb{E}_{b \sim B(a)}\left[f(b)\right] - \mathbb{E}_{b \sim B}\left[f(b)\right]\right| > \varepsilon + 2\rho\right] \leq \delta/\rho.$$

[3] This is related to the usual notion of biregularity; specifically, if A/B is biregular in the usual sense, then it is $0-$biregular in the sense of Definition 4.

2. *For any $\rho > 0$, B/A is $\big(\rho, 2(\varepsilon+\delta+\eta)/\rho\big)-$sampling.*
3. *For any $B' \subset B$ of size $|B'| = \lambda \cdot |B|$ with $\lambda > \varepsilon$,*

$$\Delta\Big(\big\{(a,b) : \substack{b\sim B' \\ a\sim A(b)}\big\}, \big\{(a,b) : \substack{a\sim A \\ b\sim B'(a)}\big\}\Big) \leq \delta + \eta/\varepsilon,$$

where $B'(a)$ denotes the distribution which draws $b \sim B(a)$ and outputs if $b \in B'$, else resamples (or if $B(a) \cap B' = \emptyset$, $B'(a)$ outputs an arbitrary $b \in B$).

The facts above are all well-known. See, for example, [Zuc97, IKW12, BDN17] for proofs of points 1, 2, and 3, respectively.

Fact 2 (Extending Sampling via Biregularity). *Fix $\varepsilon, \varepsilon', \delta, \delta', \eta > 0$. Suppose $A/B/C$ are such that $B(a)/C(a)$ is $\eta-$biregular and $C(a,b) = C(b)$ for all $a \in A$ and $b \in B(a)$. The following hold.*

1. *If B/C is $(\varepsilon', \delta')-$sampling and A/B is $\eta-$biregular, then A/C is $(\varepsilon, \delta)-$sampling, where $\delta \geq \varepsilon^{-1} \cdot (2\eta + \varepsilon' + \delta')$.*
2. *If A/B is $(\varepsilon', \delta')-$sampling and B/C is $\eta-$biregular, then A/C is $(\varepsilon, \delta)-$sampling, where $\varepsilon \geq 3\varepsilon' + 2\eta$ and $\delta \geq \delta'/\varepsilon'$.*

Proof. Assume $A/B/C$ are such that for all $a \in A$, $B(a)/C(a)$ is $\eta-$biregular, and also that $C(a,b) = C(b)$. Let $C' \subset C$ be a subset of size $|C'| = \lambda \cdot |C|$. By $\eta-$biregularity,

$$\Big|\Pr_{c\sim C(a)}(c \in C') - \lambda\Big| \leq \Big|\mathbb{E}_{b\sim B(a)}\big[\Pr_{c\sim C(b)}(c \in C')\big] - \lambda\Big| + \eta$$

holds for all $a \in A$. Now, let $\mathsf{val} := \Pr_{a\sim A}\big[|\Pr_{c\sim C(a)}(c \in C') - \lambda| > \varepsilon\big]$ be the quantity we have to bound. For the first point we have

$$\mathsf{val} \leq \varepsilon^{-1} \cdot \left(\mathbb{E}_{\substack{a\sim A \\ b\sim B(a)}}\Big[\big|\Pr_{c\sim C(b)}(c \in C') - \lambda\big|\Big] + \eta\right) \leq \varepsilon^{-1} \cdot (2\eta + \varepsilon' + \delta'),$$

by Markov's inequality, the $\eta-$biregularity of A/B and the $(\varepsilon', \delta')-$sampling of B/C. For the second point we have

$$\mathsf{val} \leq \Pr_{a\sim A}\left[\Big|\mathbb{E}_{b\sim B(a)}\big[\lambda(b)\big] - \mathbb{E}_{b\sim B}\big[\lambda(b)\big]\Big| > \varepsilon - 2\eta \geq 3\varepsilon'\right] \leq \delta'/\varepsilon',$$

where $\lambda(b) := \Pr_{c\sim C(b)}(c \in C')$. We have used the $\eta-$biregularity of B/C to say that $\mathbb{E}_{b\sim B}\big[\lambda(b)\big]$ is in $\lambda \pm \eta$, and the $(\varepsilon', \delta')-$sampling of A/B combined with the first point of Fact 1. □

Fact 3 (Replacement Product). *Let $\varepsilon, \varepsilon', \delta, \delta' > 0$ be such that $\delta \cdot (\varepsilon - 5\varepsilon') \geq 2\delta'/\varepsilon'$. Suppose $A/B/C$ is such that:*

- *A/C, B/C and $B(a)/C(a)$ are $0-$biregular for all $a \in A$; and*
- *A/C and $A(c)/B(c)$ are $(\varepsilon', \delta')-$sampling for all $c \in C$.*

Then A/B is $(\varepsilon, \delta)-$sampling.

The replacement product was originally proved in [WZ93] in the context of seeded randomness extractors (which are equivalent to sampler graphs). We give the proof ported over to the language of samplers in Appendix A for completeness.

2.3 Polynomials Over Finite Fields and Incidence Geometry

Let $\mathbb{F}$ be a finite field, and let $k \geq 4$ and $d \geq 2$ be dimension and degree parameters, respectively. We denote by A the set of affine $3-$planes in $\mathbb{F}^k$, $\mathsf{C} = \mathbb{F}^k$. The edge relation in the bipartite graph A/C is incidence: $\mathsf{a} \sim \mathsf{c}$ iff $\mathsf{c} \in \mathsf{a}$. Let Γ and Γ_{A} be the sets of $k-$variate and $3-$variate polynomials of degree at most d over $\mathbb{F}$, respectively; let $\Gamma_{\mathsf{C}} = \mathbb{F}$. This defines an incidence $\times$ agreement bipartite graph $\overline{\mathsf{A}}/\overline{\mathsf{C}}$ where $\overline{\mathsf{A}} = \mathsf{A} \times \Gamma_{\mathsf{A}}$, $\overline{\mathsf{C}} = \mathsf{C} \times \Gamma_{\mathsf{C}}$ and the edge relation is "incidence $\times$ agreement": $\overline{\mathsf{a}} = (\mathsf{a}, \alpha) \sim (\mathsf{c}, \gamma) = \overline{\mathsf{c}}$ iff $\mathsf{c} \in \mathsf{a}$ and $\alpha|_{\mathsf{c}} = \gamma$. We show in Section 5 that $\overline{\mathsf{A}}/\overline{\mathsf{C}}$ has similar sampling properties to A/C, whose sampling properties are well known.

3 Locally Testable, Non-Malleable Codes

3.1 Definition

Just as how we defined non-malleable codes via non-malleable reductions, likewise locally testable, non-malleable codes are a special case of locally testable, non-malleable reductions which we now define. As mentioned previously, this

Definition 6 (Locally Testable, Non-Malleable Reductions). *Fix parameters $\ell \in \mathbb{N}$, $\varepsilon > 0$, and function families $\mathcal{F} \subset \{f : \Gamma^n \to \Gamma^n\}$ and $\mathcal{G} \subset \{g : \Gamma^n \to \Gamma^n\}$. We say that a LTC $(\mathsf{Enc}, \mathsf{Dec}, \mathsf{Test})$ is an $(\ell, \varepsilon)-$*locally-testable non-malleable reduction *from $\mathcal{F}$ to $\mathcal{G}$ if for all $f \in \mathcal{F}$, there exists $\mathsf{L}_f = \{g^{(1)}, \ldots, g^{(\ell)}\} \subset \mathcal{G}$ of size $|\mathsf{L}_f| = \ell$ and a function $g : \Gamma^n \to \Gamma^n$ such that:*

1. *for all $i \in [n]$ and $\mathbf{y} \in \Gamma^n$, $g(\mathbf{y})_i \in \{g^{(j)}(\mathbf{y})_i : g^{(j)} \in \mathsf{L}_f\}$; and*
2. *for all $m \in \Gamma^k$, $\Pr_{\mathbf{x} \sim \mathsf{Enc}(m), I}\Big[\mathsf{Test}\big(f(\mathbf{x}); I\big) = 1 \;\&\; f(\mathbf{x})_I \neq g(\mathbf{x})_I\Big] \leq \varepsilon$.*

*As before, if $\mathcal{G}$ is the family of trivial tampering functions consisting just of the identity and constants, then $(\mathsf{Enc}, \mathsf{Dec}, \mathsf{Test})$ is called an $(\ell, \varepsilon)-$*locally-testable non-malleable code.

Remark. Some remarks are in order.

1. The list-decoding intuition is captured by the shortness of L_f: nearly all of the test passing probability of an $f-$tampered codeword is explained by f's agreement with g, which in turn, always agrees with one of the functions in $\mathsf{L}_f \subset \mathcal{G}$. Note that each coordinate of g is a (possibly different) convex combination of the corresponding coordinates of the $g^{(j)}$. The non-malleability intuition is captured because the function g is defined given f, and the agreement guarantee of point 2 holds for all $m \in \Gamma^k$.

2. Unlike Definition 3, the functions in $\mathcal{G}$ in Definition 6 map codewords to codewords, rather than messages to messages. This modification is so that we can meaningfully compare $f(\mathbf{x})_I$ with $g(\mathbf{x})_I$, an important feature of local-testing definitions. The family $\mathsf{Dec} \circ \mathcal{G} \circ \mathsf{Enc}$ would be the corresponding distribution on message-to-message functions. In this work, $\mathcal{G}$ will always be either the family of trivial tampering functions, or the family of affine tampering functions. In either case, $\mathsf{Dec} \circ \mathcal{G} \circ \mathsf{Enc}$ is also trivial or affine. The distribution G_f of

Definition 3 outputs $g^{(j)}$ with probability proportional to the probability that $\mathcal{D}_f^{\mathsf{SIM}}$ agrees with $g^{(j)}$.

3. Composing two standard non-malleable reductions $-$ one from $\mathcal{F}$ to $\mathcal{G}$, one from $\mathcal{G}$ to $\mathcal{H}$ $-$ yields a non-malleable reduction from $\mathcal{F}$ to $\mathcal{H}$. The same composition theorem does not hold generically for locally testable, non-malleable reductions. We use a non-generic composition theorem to combine a locally testable, non-malleable reduction from $\mathcal{F}$ to $\mathcal{G}$ with a non-malleable code against $\mathcal{G}$ (for specific $\mathcal{F}$ and $\mathcal{G}$) to obtain a locally testable, non-malleable code against $\mathcal{F}$. The test of our composed code involves locally decoding a symbol of the outer code so it can be checked for validity by the inner code. This idea is often used to compose locally testable codes and PCPs.

3.2 Discussion

In this section, we briefly discuss how our new coding gadgets relate to other nearby members in the coding theory tree. Additionally we discuss a few naive attempts at building LTNMCs by combining other coding objects.

First, it is clear that any LTNMC is also a LTC. On the other hand, LTNMCs do not seem to immediately give NMCs. Essentially, this is because the tester for LTNMC does not distinguish between the case when the tampered codeword is valid and when it is very close but not equal to a valid codeword. Thus, the definition 6 does not prevent "selective bot attacks" where the probability of decoding failure varies very slightly with the message. For example, a tampering function might be able to tamper an encoding of $m = 0$ to a valid codeword, and an encoding of $m = 1$ to a valid codeword except with a single incorrect symbol. In this case the tester will not notice the difference, but a decoding algorithm will have to output 0 in one case and $\perp$ in the other.

In the other direction, NMCs also do not readily give LTNMC because they might not be locally testable. One might try composed a NMC with an outer LTC to obtain a code with a local tester and (hopefully) some non-malleability properties. However, in order to show that the concatenated code is non-malleable, one basically has to show that if the outer LTC is tampered, the resulting tampering on the inner NMC is precisely what it is secure against. Thus, this requires the outer LTC to already have some non-malleability.

One notable exception to this is the case of linear (or affine) tampering. If an LTC has an encoding algorithm which is linear and the inner NMC is non-malleable against affine tampering, then the composed code will be a LTNMC against affine tampering as well, since an affine attack on the outer code translates (by linearity) to an affine attack on the inner one.

3.3 Sufficient Conditions for LTNMCs

The following claim gives a useful set of sufficient conditions for a LTC being non-malleable. For the codes used in this work, the first condition will be more or less trivial to establish. Thus, Claim 1 essentially reduces proving non-malleability to the task of establishing condition 2. This will simplify our proofs considerably.

Claim 1 (Sufficient Conditions for Non-Malleability in LTCs.). *Let* $(\mathsf{Enc}, \mathsf{Dec}, \mathsf{Test})$ *be a LTC with* $\mathsf{Enc} : \Gamma^k \to \Gamma^n$*, and let*

$$\mathcal{F} \subset \{f : \Gamma^n \to \Gamma^n\} \text{ and } \mathcal{G} \subset \left\{g : \Gamma^n \to (\Gamma \cup \{\perp\})^n\right\}$$

be function families. Suppose the following two conditions hold.

1. **Tampering function distance:** *For all distinct function pairs* $g, g' \in \mathcal{G}$ *and* $m \in \Gamma^k$*, we have*
$$\Pr_{\mathbf{x} \sim \mathsf{Enc}(m), I}\left[g(\mathbf{x})_I = g'(\mathbf{x})_I\right] \leq \varepsilon/\ell^2$$
2. **List decoding:** *For all* $f \in \mathcal{F}$ *there exists a list* $\mathsf{L}_f = \{g^{(1)}, \ldots, g^{(\ell)}\} \subset \mathcal{G}$ *of size* $|\mathsf{L}_f| = \ell$ *such that for all* $m \in \Gamma^k$*,*
$$\Pr_{\mathbf{x} \sim \mathsf{Enc}(m), I}\left[\mathsf{Test}\big(f(\mathbf{x}); I\big) = 1 \;\&\; f(\mathbf{x})_I \notin \left\{g^{(j)}(\mathbf{x})_I : g^{(j)} \in \mathsf{L}_f\right\}\right] \leq \varepsilon.$$

Then $(\mathsf{Enc}, \mathsf{Dec}, \mathsf{Test})$ *is an* $(\ell, 2\varepsilon)-$*locally testable, non-malleable reduction from* $\mathcal{F}$ *to* $\mathcal{G}$*.*

Proof. Define $g : \Gamma^n \to (\Gamma \cup \{\perp\})^n$ by setting $g(\mathbf{x})_i = f(\mathbf{x})_i$ if there is a unique $g^{(j)} \in \mathsf{L}_f$ such that $f(\mathbf{x})_i = g^{(j)}(\mathbf{x})_i$, and $g(\mathbf{x})_i = \perp$ otherwise. Note every coordinate of g is a convex combination of the corresponding coordinate functions in $\mathsf{L}_f \cup \{\perp\}$, a subset of $\mathcal{G}$ of size $\ell + 1$. Thus, for all $m \in \Gamma^k$,

$$\begin{aligned}
&\Pr_{\mathbf{x} \sim \mathsf{Enc}(m), I}\left[\mathsf{Test}\big(f(\mathbf{x}); I\big) = 1 \;\&\; f(\mathbf{x})_I \neq g(\mathbf{x})_I\right] \\
&\qquad \leq \Pr_{\mathbf{x} \sim \mathsf{Enc}(m), I}\left[\exists\, g^{(j)}, g^{(j')} \in \mathsf{L}_f \text{ st } g^{(j)}(\mathbf{x})_I = g^{(j')}(\mathbf{x})_I\right] \\
&\qquad + \Pr_{\mathbf{x} \sim \mathsf{Enc}(m), I}\left[\mathsf{Test}\big(f(\mathbf{x}); I\big) = 1 \;\&\; f(\mathbf{x})_I \notin \left\{g^{(j)}(\mathbf{x})_I : g^{(j)} \in \mathsf{L}_f\right\}\right] \\
&\qquad \leq \binom{\ell}{2} \cdot \frac{\varepsilon}{\ell^2} + \varepsilon \leq 2\varepsilon.
\end{aligned}$$

We used the two given conditions to bound the two terms in second line. The claim follows. □

3.4 Our Code and Main Theorem

Main Construction. We choose a dimension parameter $k \geq 4$ and the degree parameter $d \geq 2$. With notations as defined in section 2.3:

- $\underline{\mathsf{Enc}(m)}$**:** For $m \in \mathbb{F}$, draw $\Phi \sim \Gamma$ such that $\Phi(\mathbf{0}) = m$ and output $\{\Phi|_{\mathsf{a}}\}_{\mathsf{a} \in \mathsf{A}} \in \Gamma_{\mathsf{A}}^{|\mathsf{A}|}$. We will often write codewords as $\big\{(\mathsf{a}, \alpha)\big\}_{\mathsf{a} \in \mathsf{A}}$ with the understanding that $\alpha = \Phi|_{\mathsf{a}}$.
- $\underline{\mathsf{Dec}\big(\{(\mathsf{a}, \alpha)\}_{\mathsf{a} \in \mathsf{A}}\big)}$**:** Given $\big\{(\mathsf{a}, \alpha)\big\}_{\mathsf{a} \in \mathsf{A}}$, find $\Phi \in \Gamma$ such that $(\mathsf{a}, \alpha) = (\mathsf{a}, \Phi|_{\mathsf{a}})$ for all $\mathsf{a} \in \mathsf{A}$.[4] If such Φ exists, output $m = \Phi(\mathbf{0})$, otherwise output $\perp$.[5]

[4] Such Φ, if it exists, can be found in time $\mathsf{poly}\big(|\mathbb{F}|\big)$ by interpolation.

[5] As written, decoding runs in time $\mathsf{poly}(|\mathbb{F}|)$, which is exponential in the message length. However, local decoding algorithms exist which run in time $\mathsf{poly}\big(\lambda, \log|\mathbb{F}|, 1/\delta\big)$ and output m (or a list containing m) with probability $1 - 2^{-\lambda}$ whenever the input is within distance δ of a valid encoding of m. See for example [Sud97].

• $\underline{\mathsf{Test}\big(\{(\mathsf{a},\alpha)\}_{\mathsf{a}\in\mathsf{A}}\big)}$**:** Draw $\mathsf{c}\sim\mathsf{C}$, $\mathsf{a},\mathsf{a}'\sim\mathsf{A}(\mathsf{c})$; read (a,α) and (a',α'), and output 1 if $\alpha|_{\mathsf{c}}=\alpha'|_{\mathsf{c}}$ ($\alpha|_{\mathsf{c}}$ denotes the evaluation of α at c), 0 otherwise.

The above code is known to be a $\big(2,|\mathbb{F}|^{-\Omega(1)}\big)-$locally testable code. This was proven originally the influential works [AS97, RS97]. Our main theorem is that this code also possesses non-malleability guarantees. Before we state this formally, we introduce the tampering function families.

Tampering Function Families. We identify three types of tampering.

• **Coordinate-Wise:** $\mathcal{F}:=\Big\{\{\mathsf{f}_{\mathsf{a}}\}_{\mathsf{a}\in\mathsf{A}}\Big|\mathsf{f}_{\mathsf{a}}:\Gamma_{\mathsf{A}}\to\Gamma_{\mathsf{A}}\Big\}$ tampers codewords via

$$\{\mathsf{f}_{\mathsf{a}}\}_{\mathsf{a}}:\{(\mathsf{a},\alpha)\}_{\mathsf{a}}\mapsto\Big\{\big(\mathsf{a},\mathsf{f}_{\mathsf{a}}(\alpha)\big)\Big\}_{\mathsf{a}}.$$

• **Affine:** We say that $\mathsf{T}:\Gamma\to\Gamma$ is *affine* if $\exists\ (s,\Phi_0)\in\mathbb{F}\times\Gamma$ such that $\mathsf{T}(\Phi)=s\cdot\Phi+\Phi_0$. We define $\mathcal{G}$ to be the family of coordinate-wise restrictions of global affine maps:

$$\mathcal{G}:=\Big\{\{\mathsf{g}_{\mathsf{a}}\}_{\mathsf{a}\in\mathsf{A}}\Big|\ \exists\ (s,\Phi_0)\in\mathbb{F}\times\Gamma\text{ st }\mathsf{g}_{\mathsf{a}}(\alpha)=s\cdot\alpha+\Phi_0|_{\mathsf{a}}\ \forall\ \mathsf{a}\in\mathsf{A}\Big\}\subset\mathcal{F}.$$

• **Trivial:** We say that $\mathsf{T}:\Gamma\to\Gamma$ is *trivial* if either $\mathsf{T}(\Phi)=\Phi$ or if $\exists\ \ \Phi_0\in\Gamma$ such that $\mathsf{T}(\Phi)=\Phi_0$. We define $\mathcal{H}$ to be the family of coordinate-wise restrictions of trivial maps:

$$\mathcal{H}:=\Big\{\{\mathsf{h}_{\mathsf{a}}\}_{\mathsf{a}\in\mathsf{A}}\Big|\text{ either }\mathsf{h}_{\mathsf{a}}(\alpha)=\alpha\ \forall\ (\mathsf{a},\alpha)\text{ or }\exists\ \Phi_0\in\Gamma\text{ st }\mathsf{h}_{\mathsf{a}}(\alpha)=\Phi_0|_{\mathsf{a}}\ \forall\ (\mathsf{a},\alpha)\Big\}\subset\mathcal{G}.$$

We also include the "all $\perp$ function" (maps every coordinate to $\perp$) in $\mathcal{G}$ and $\mathcal{H}$.

Theorem 1. *The code above is an $(\ell,\varepsilon)-$locally-testable, non-malleable reduction from $\mathcal{F}$ to $\mathcal{G}$ where $\varepsilon=|\mathbb{F}|^{-\Omega(1)}$ and $\ell=4/\varepsilon$.*

We use this locally testable, non-malleable reduction to build a locally testable, non-malleable code against $\mathcal{F}$. The explicit construction is given in section 8.

Theorem 2 (Main). *There exists an explicit $(\ell,\varepsilon)-$locally testable, non-malleable code against $\mathcal{F}$, the family of coordinate-wise tampering functions where $\varepsilon=|\mathbb{F}|^{-\Omega(1)}$ and $\ell=4/\varepsilon$.*

4 The Affine Agreement Theorem

In this section we state the affine agreement theorem, which is at the core of the proof of Theorem 1. Theorem 1 follows from our affine agreement theorem in much the same way as list-decoding theorems often follow from agreement theorems.

Theorem 3 (Affine Agreement). *There exists $\varepsilon=|\mathbb{F}|^{-\Omega(1)}$ such that for all $\{\mathsf{f}_{\mathsf{a}}\}\in\mathcal{F}$, the following holds. If*

$$\Pr_{\Phi,(\mathsf{c},\mathsf{a},\mathsf{a}')}\big[\tilde{\alpha}|_{\mathsf{c}}=\tilde{\alpha}'|_{\mathsf{c}}\big]\geq 6\varepsilon,$$

where the probability is over $\Phi\sim\Gamma$, $\mathsf{c}\sim\mathsf{C}$, $\mathsf{a},\mathsf{a}'\sim\mathsf{A}(\mathsf{c})$ and where $(\tilde{\alpha},\tilde{\alpha}')=\big(\mathsf{f}_{\mathsf{a}}(\Phi|_{\mathsf{a}}),\mathsf{f}_{\mathsf{a}'}(\Phi|_{\mathsf{a}'})\big)$, then there exists an affine $\mathsf{T}:\Gamma\to\Gamma$ such that $\Pr_{(\Phi,\mathsf{a})\sim\Gamma\times\mathsf{A}}\big[\tilde{\alpha}=\mathsf{T}(\Phi)|_{\mathsf{a}}\big]\geq\varepsilon$.

Proof of Theorem 1 Assuming Theorem 3. Let ε be as in Theorem 3 above and fix $\mathsf{f} = \{\mathsf{f}_\mathsf{a}\}_\mathsf{a} \in \mathcal{F}$. We prove that the code is non-malleable by establishing the sufficient conditions of Claim 1. The first condition is immediate. Because: For all distinct $\{\mathsf{g}_\mathsf{a}\}_\mathsf{a}, \{\mathsf{g}'_\mathsf{a}\}_\mathsf{a} \in \mathcal{G}$, $\mathsf{g}_\mathsf{a}(\Phi|_\mathsf{a}) = \mathsf{g}'_\mathsf{a}(\Phi|_\mathsf{a})$ holds only if either $\mathsf{g}_\mathsf{a} = \mathsf{g}'_\mathsf{a}$ (occurs with probability $\mathcal{O}\big(|\mathbb{F}|^{-1}\big)$ when $\{\mathsf{g}_\mathsf{a}\}_\mathsf{a} \neq \{\mathsf{g}'_\mathsf{a}\}_\mathsf{a}$), or if $\mathsf{g}_\mathsf{a} \neq \mathsf{g}'_\mathsf{a}$ but $\mathsf{g}_\mathsf{a}(\Phi|_\mathsf{a}) = \mathsf{g}'_\mathsf{a}(\Phi|_\mathsf{a})$ (also probability $\mathcal{O}\big(|\mathbb{F}|^{-1}\big)$). Thus, $\Pr_{\Phi,(\mathsf{c},\mathsf{a},\mathsf{a}')}\big[\mathsf{g}_\mathsf{a}(\Phi|_\mathsf{a}) = \mathsf{g}'_\mathsf{a}(\Phi|_\mathsf{a})\big] = \mathcal{O}\big(|\mathbb{F}|^{-1}\big) \ll \varepsilon/\ell^2$. For the second condition, we show that there exists $\mathsf{L}_\mathsf{f} \subset \mathcal{G}$ of size at most ℓ such that

$$\Pr_{\Phi,(\mathsf{c},\mathsf{a},\mathsf{a}')}\Big[\tilde{\alpha}|_\mathsf{c} = \tilde{\alpha}'|_\mathsf{c} \;\&\; (\tilde{\alpha},\tilde{\alpha}') \notin \big\{\big(\mathsf{g}_\mathsf{a}(\alpha), \mathsf{g}_{\mathsf{a}'}(\alpha')\big) : \{\mathsf{g}_\mathsf{a}\}_\mathsf{a} \in \mathsf{L}_\mathsf{f}\big\}\Big] < 6\varepsilon, \tag{2}$$

where $(\tilde{\alpha},\tilde{\alpha}') = \big(\mathsf{f}_\mathsf{a}(\alpha), \mathsf{f}_{\mathsf{a}'}(\alpha')\big)$ for $(\alpha,\alpha') = \big(\Phi|_\mathsf{a}, \Phi|_{\mathsf{a}'}\big)$, and where $\Phi \sim \Gamma$.[6] Towards this end, let $\mathsf{L}_\mathsf{f} := \big\{\{\mathsf{g}_\mathsf{a}\}_\mathsf{a} \in \mathcal{G} :\ \Pr_{(\Phi,\mathsf{a})\sim\Gamma\times\mathsf{A}}\big[\tilde{\alpha} = \mathsf{g}_\mathsf{a}(\alpha)\big] \geq \varepsilon/2\big\}$.

Small List Size. Assume for contradiction that $|\mathsf{L}_f| \geq \ell = 4/\varepsilon + 1$, and so contains a set $\big\{\{\mathsf{g}^1_\mathsf{a}\}_\mathsf{a}, \ldots, \{\mathsf{g}^\ell_\mathsf{a}\}_\mathsf{a}\big\}$. By inclusion-exclusion,

$$\begin{aligned} 1 \;&\geq\; \Pr_{(\Phi,\mathsf{a})\sim\Gamma\times\mathsf{A}}\Big[\tilde{\alpha} \in \big\{\mathsf{g}^i_\mathsf{a}(\alpha) : i = 1,\ldots,\ell\big\}\Big] \\ &\geq\; \frac{\ell\cdot\varepsilon}{2} - \sum_{1\leq i<j\leq\ell} \Pr_{\Phi,\mathsf{a}}\Big[\mathsf{g}^i_\mathsf{a}(\alpha) = \mathsf{g}^j_\mathsf{a}(\alpha)\Big] > 2 - \binom{\ell}{2}\cdot\left(\frac{1}{|\Gamma|} + \frac{d}{|\mathbb{F}|}\right). \end{aligned}$$

The last inequality used $\ell\varepsilon > 4$, and the bound on $\Pr_{\Phi,\mathsf{a}}\big[\mathsf{g}^i_\mathsf{a}(\Phi|_\mathsf{a}) = \mathsf{g}^j_\mathsf{a}(\Phi|_\mathsf{a})\big]$ from point 2 above. The right hand side simplifies to $2 - o(1) > 1$, a contradiction.

Proximity Implies List Decoding. Suppose $\{\mathsf{f}_\mathsf{a}\}$ is such that (2) does not hold. Define $\{\mathsf{f}'_\mathsf{a}\}_\mathsf{a} \in \mathcal{F}$ as follows: $\mathsf{f}'_\mathsf{a}(\alpha) = \mathsf{f}_\mathsf{a}(\alpha)$, unless $\mathsf{f}_\mathsf{a}(\alpha) = \mathsf{g}_\mathsf{a}(\alpha)$ for some $\{\mathsf{g}_\mathsf{a}\}_\mathsf{a} \in \mathsf{L}_\mathsf{f}$ in which case $\mathsf{f}'_\mathsf{a}(\alpha)$ outputs a random $\tilde{\alpha} \notin \big\{\mathsf{g}_\mathsf{a}(\alpha) : \{\mathsf{g}_\mathsf{a}\}_\mathsf{a} \in \mathsf{L}_\mathsf{f}\big\}$. Note

$$\Pr_{\Phi,(\mathsf{c},\mathsf{a},\mathsf{a}')}\big[\mathsf{f}'_\mathsf{a}(\alpha)|_\mathsf{c} = \mathsf{f}'_{\mathsf{a}'}(\alpha')|_\mathsf{c}\big] \geq 6\varepsilon$$

since (2) does not hold. Therefore, by Theorem 3, there exists an affine $\mathsf{T} : \Gamma \to \Gamma$ such that $\Pr_{\Phi,\mathsf{a}}\big[\mathsf{f}'_\mathsf{a}(\Phi|_\mathsf{a}) = \mathsf{T}(\Phi)|_\mathsf{a}\big] \geq \varepsilon$. Thus $\Pr_{\Phi,\mathsf{a}}\big[\mathsf{f}_\mathsf{a}(\Phi|_\mathsf{a}) = \mathsf{T}(\Phi)|_\mathsf{a}\big] \geq \varepsilon - \ell/|\Gamma_\mathsf{A}| \geq \varepsilon/2$, and so the coordinate-wise version of T is in L_f. This is a contradiction since by construction, for every $\{\mathsf{g}_\mathsf{a}\}_\mathsf{a} \in \mathsf{L}_\mathsf{f}$, $\mathsf{f}'_\mathsf{a}(\alpha) \neq \mathsf{g}_\mathsf{a}(\alpha)$ holds for all $\mathsf{a} \in \mathsf{A}$ and $\alpha \in \Gamma_\mathsf{A}$. □

4.1 Overview of the Proof of Theorem 3

Theorem 3 roughly says if any random tampered codeword passes the the plane-point-plane agreement test with good probability then the tampering function, f, must be close to some affine mapping. To see intuitively why this theorem holds, lets first consider a more favorable situation where the test passes with close to one probability (over random codewords and test indices). Since our code is already known to be locally testable, any tampered codeword that passes the test with high probability, must agree with some polynomial $\tilde{\Phi}$ at many co-oridnates. Thus it must be the case

[6] The difference in probability caused by drawing $\Phi \sim \Gamma$ such that $\Phi(\mathbf{0}) = m$ instead is $\mathcal{O}(|\mathbb{F}|^{-1})$, thus negligible.

that our tampering f is mapping random codewords to close to valid codewords or equivalently, we can think the tampering as defining a function that sends a random Φ to some $\tilde{\Phi}$. The key of our theorem is showing that if f is a coordinate-wise function that maps random polynomials to polynomials, then it must be affine. This completes the very high level proof sketch of Theorem 3 modulo the low error assumption. Now, to prove the theorem in the high error regime, we first show that whenever low agreement holds, there is a small fraction of planes where very high agreement holds. Thus, we can essentially perform the proof mentioned above on this small set of planes to get our result. The lemmas described in the next section capture these two pieces of intuition.

4.2 Reducing the NM Agreement Theorem to Two Lemmas

The proof of Theorem 3 will occupy much of the rest of this paper. In this section, we separate the proof into two parts by stating two lemmas which combine to immediately prove the theorem.

Proof of Theorem 3. Suppose $\varepsilon = |\mathbb{F}|^{-\Omega(1)}$ is chosen so it satisfies Lemmas 1 and 2, below. Let $\{\mathsf{f}_\mathsf{a}\}_\mathsf{a} \in \mathcal{F}$ be such that

$$\Pr_{\Phi,(\mathsf{c},\mathsf{a},\mathsf{a}')}\left[\tilde{\alpha}|_\mathsf{c} = \tilde{\alpha}'|_\mathsf{c}\right] \geq 6\varepsilon. \tag{3}$$

By Lemma 1 below, there exists a function $\mathsf{h} : \overline{\mathsf{C}} \to \Gamma_\mathsf{C}$ such that

$$\Pr_{(\mathsf{a},\Phi)\sim \mathsf{A}\times\Gamma}\left[\Pr_{\bar{\mathsf{c}}\sim\overline{\mathsf{C}}(\bar{\mathsf{a}})}\left[\tilde{\alpha}|_\mathsf{c} = \tilde{\gamma}\right] \geq 1-\zeta\right] \geq 2\varepsilon, \tag{4}$$

where $\tilde{\gamma} = \mathsf{h}(\bar{\mathsf{c}})$, $\bar{\mathsf{a}} = (\mathsf{a}, \Phi|_\mathsf{a})$, and where $\zeta = |\mathbb{F}|^{-\Omega(1)}$ is specified precisely in Section 6. By Lemma 2, there exists an affine map $\mathsf{T} : \Gamma \to \Gamma$ such that

$$\Pr_{(\mathsf{a},\Phi)\sim \mathsf{A}\times\Gamma}\left[\tilde{\alpha} = \mathsf{T}(\Phi)|_\mathsf{a}\right] \geq \varepsilon. \tag{5}$$

□

Lemma 1 (Global Agreement). *There exists $\varepsilon = |\mathbb{F}|^{-\Omega(1)}$ such that whenever $\{\mathsf{f}_\mathsf{a}\}_\mathsf{a} \in \mathcal{F}$ is such that* (3) *holds, there exists $\mathsf{h} : \overline{\mathsf{C}} \to \Gamma_\mathsf{C}$ such that* (4) *holds.*

Lemma 2 (Affine Agreement). *There exists $\varepsilon = |\mathbb{F}|^{-\Omega(1)}$ such that whenever $\{\mathsf{f}_\mathsf{a}\}_\mathsf{a} \in \mathcal{F}$ and $\mathsf{h} : \overline{\mathsf{C}} \to \Gamma_\mathsf{C}$ are such that* (4) *holds, there exists an affine $\mathsf{T} : \Gamma \to \Gamma$ such that* (5) *holds.*

Lemma 1 is proved in Section 6 using a sampler-based decoding argument similar to ones which have appeared in several recent works, for example [BDN17]. The linearity test analyzed in the proof of Lemma 2 in Section 7 is new to this work.

5 Sampler Graph Preliminaries

5.1 Why Samplers Play a Role

Here we briefly discuss how sampler graphs serve as an important component in our analysis. We begin by recalling the 'plane vs plane' low degree testing model from PCP literature [RS97,

BDN17]. In this model, a test algorithm gets oracle access to a 'planes' table where to each plane, $\mathsf{a} \in \mathsf{A}$, the table contains a polynomial, α, defined on that plane. Then the test algorithm's task is to decide if the table is close to any global low degree polynomial Φ. The final step is then to prove an agreement theorem that says if the test passes with good probability, then there exists a polynomial that agrees with the table on many planes. In literature, these agreement theorems are proven using essentially two ingredients: sampling properties of planes and facts about low degree polynomials. Now, its easy to see that our tampering and testing model is very similar to the 'plane vs plane' model. The only difference is that in our model as we are looking at coordinate-wise tampering $f_a(\alpha) = \tilde{\alpha}$, we have a 'plane$\times$polynomial table' where to (a, α) the table contains a polynomial $\tilde{\alpha}$. Thus, to prove an agreement theorem in our setting, we wind up using sampling of 'planes$\times$polynomials' [see section 5.2 below] and the same facts about polynomials.

5.2 Incidence $\times$ Agreement Samplers

Sampler graphs play a big role in the proofs in the following sections. In this section we introduce the graphs whose sampling will be used, and various properties of sampler graphs. All of the graphs are what we call "incidence $\times$ agreement" graphs, such as $\overline{\mathsf{A}}/\overline{\mathsf{C}}$ from last section. We begin with some notation.

Notation. Recall $\mathbb{F}$ is a finite field, $k \geq 4$, $d \geq 2$, A is the set of $3-$planes in $\mathbb{F}^k$, $\mathsf{C} = \mathbb{F}^k$, Γ and Γ_{A} are the sets of $k-$variate and $3-$variate polynomials of degree at most d over $\mathbb{F}$, respectively, $\Gamma_{\mathsf{C}} = \mathbb{F}$. This defines an incidence $\times$ agreement bipartite graph $\overline{\mathsf{A}}/\overline{\mathsf{C}}$ where $\overline{\mathsf{A}} = \mathsf{A}\times\Gamma_{\mathsf{A}}$, $\overline{\mathsf{C}} = \mathsf{C}\times\Gamma_{\mathsf{C}}$ and the edge relation is "incidence $\times$ agreement": $\overline{\mathsf{a}} = (\mathsf{a}, \alpha) \sim (\mathsf{c}, \gamma) = \overline{\mathsf{c}}$ iff $\mathsf{c} \in \mathsf{a}$ and $\alpha|_{\mathsf{c}} = \gamma$. For $r = 1, 2$, let B_r denote the set of affine $r-$dimensional planes in $\mathbb{F}^k$, let Γ_{B_r} be the set of $r-$variate polynomials of degree at most d over $\mathbb{F}$, and let $\overline{\mathsf{B}}_r = \mathsf{B}_r \times \Gamma_{\mathsf{B}_r}$. At various points during the proof, we will use that $\overline{\mathsf{A}}/\overline{\mathsf{B}}_r/\overline{\mathsf{C}}$ is a double sampler. The incidence $\times$ agreement edge relation extends naturally to $\overline{\mathsf{A}}/\overline{\mathsf{B}}_r$, $\overline{\mathsf{B}}_r/\overline{\mathsf{C}}$, and $\overline{\mathsf{B}}_2/\overline{\mathsf{B}}_1$. For example, if $\overline{\mathsf{a}} = (\mathsf{a}, \alpha) \in \overline{\mathsf{A}}$ and $\overline{\mathsf{b}} = (\mathsf{b}, \beta) \in \overline{\mathsf{B}}_2$, then $\overline{\mathsf{a}} \sim \overline{\mathsf{b}}$ iff $\mathsf{b} \subset \mathsf{a}$ and $\alpha|_{\mathsf{b}} = \beta$.

We begin by listing the incidence $\times$ agreement samplers we will need in the remainder of the paper and proving they are sampling. In the claim statement below, $\overline{\mathsf{A}}(\overline{\mathsf{c}})$, for $\overline{\mathsf{c}} \in \overline{\mathsf{C}}$, denotes the set of $\overline{\mathsf{a}} \in \overline{\mathsf{A}}$ such that $\overline{\mathsf{a}} \sim \overline{\mathsf{c}}$. In the proof which follows, we use $\overline{\mathsf{A}}(\overline{\mathsf{c}})$ to mean either this set, or the uniform distribution on this set; in all cases, our intention should be clear from the context.

Claim 2. *The following graphs are all $\mathcal{O}\big(|\mathbb{F}|^{-1}\big)-$biregular and $\big(12\cdot|\mathbb{F}|^{-1/15}, |\mathbb{F}|^{-1/15}\big)-$sampling:*

(1) $\overline{\mathsf{B}}_1/\overline{\mathsf{C}}$ (2) $\overline{\mathsf{A}}(\overline{\mathsf{c}})/\overline{\mathsf{B}}_2(\overline{\mathsf{c}})\ \forall\ \overline{\mathsf{c}} \in \overline{\mathsf{C}}$ (3) $\overline{\mathsf{A}}/\overline{\mathsf{C}}$ (4) $\overline{\mathsf{A}}(\overline{\mathsf{c}}, \overline{\mathsf{c}}')/\overline{\mathsf{C}}\ \forall\ \overline{\mathsf{c}}, \overline{\mathsf{c}}' \in \overline{\mathsf{C}}$
(5) $\overline{\mathsf{A}}/\overline{\mathsf{C}}^2$ (6) $\overline{\mathsf{A}}(\overline{\mathsf{c}})/\overline{\mathsf{C}}^2\ \forall\ \overline{\mathsf{c}} \in \overline{\mathsf{C}}$ (7) $\mathsf{A}\times\Gamma/\overline{\mathsf{C}}$ (8) $\overline{\mathsf{B}}_2(\overline{\mathsf{c}})/\overline{\mathsf{C}}\ \forall\ \overline{\mathsf{c}} \in \overline{\mathsf{C}}$
(9) $\overline{\mathsf{A}}(\overline{\mathsf{b}})/\overline{\mathsf{C}}\ \forall\ \overline{\mathsf{b}} \in \overline{\mathsf{B}}_1$ (10) $\mathsf{A}\times\Gamma/\overline{\mathsf{B}}_1$

Proof. It is easy to see that all of the graphs in the Claim statement are $\mathcal{O}\big(|\mathbb{F}|^{-1}\big)-$biregular, as per Definition 4. By symmetry, graphs (1), (2), (3), (5), (7), (10) are actually $0-$biregular. The others have a slight error introduced by the fact, for example, that the distribution which draws $\overline{\mathsf{a}} \sim \overline{\mathsf{A}}(\overline{\mathsf{c}})$ and outputs a random element of $\overline{\mathsf{C}}(\overline{\mathsf{a}})$ is more likely to output $\overline{\mathsf{c}}$ than $\overline{\mathsf{c}}' \neq \overline{\mathsf{c}}$. However, an easy calculation shows that the statistical distance between the required distributions is $\mathcal{O}\big(|\mathbb{F}|^{-1}\big)$; the

same is true for all examples in the list. The rest of the proof is divided into two stages. First, we use a pairwise independence argument to show that $\overline{\mathsf{B}}_1/\overline{\mathsf{C}}$, $\overline{\mathsf{B}}_2/\overline{\mathsf{C}}$, $\overline{\mathsf{A}}(\overline{\mathsf{b}}_1)/\overline{\mathsf{B}}_2(\overline{\mathsf{b}}_1)$ for all $\overline{\mathsf{b}}_1 \in \overline{\mathsf{B}}_1$ and $\overline{\mathsf{B}}_2(\overline{\mathsf{c}})/\overline{\mathsf{B}}_1(\overline{\mathsf{c}})$, $\overline{\mathsf{A}}(\overline{\mathsf{c}})/\overline{\mathsf{B}}_1(\overline{\mathsf{c}})$ for all $\overline{\mathsf{c}} \in \overline{\mathsf{C}}$ are $(|\mathbb{F}|^{-1/5}, |\mathbb{F}|^{-3/5})-$sampling. Then we reduce the sampling of every graph above to the sampling of these five.

We phrase the pairwise independence argument for a generic bipartite graph A/B. The key feature we need involves a set X which parametrizes the neighborhoods $B(a)$ for all $a \in A$. Given $x \in X$ and $a \in A$, we write the $x-$th neighbor of a as $a(x) \in B$, so X parametrizes neighborhoods as $B(a) = \{a(x) : x \in X\}$ for all $a \in A$. The property we require is that for all $x_1 \neq x_2 \in X$, the random variable $\big(a(x_1), a(x_2)\big)$ (randomness over $a \sim A$) is uniform on B^2. For $\overline{\mathsf{B}}_1/\overline{\mathsf{C}}$, $X = \mathbb{F}$ since $\overline{\mathsf{C}}(\overline{\mathsf{b}})$ is parametrized by the points on the line b. Likewise, for $\overline{\mathsf{B}}_2/\overline{\mathsf{C}}$, $X = \mathbb{F}^2$. For $\overline{\mathsf{A}}(\overline{\mathsf{b}}_1)/\overline{\mathsf{B}}_2(\overline{\mathsf{b}}_1)$,the neighborhood $\overline{\mathsf{B}}_2(\overline{\mathsf{b}}_1, \overline{\mathsf{a}})$ is paramterized by all possible planes in a through b_1, so we have $|X| = |\mathbb{F}| + 1$. For $\overline{\mathsf{B}}_2(\overline{\mathsf{c}})/\overline{\mathsf{B}}_1(\overline{\mathsf{c}})$, $X = \mathbb{F} \cup \{\infty\}$, since $\overline{\mathsf{B}}_1(\overline{\mathsf{c}}, \overline{\mathsf{b}}_2)$ is parametrized by all possible slopes of a line in b_2 through c. Finally, for $\overline{\mathsf{A}}(\overline{\mathsf{c}})/\overline{\mathsf{B}}_1(\overline{\mathsf{c}})$ the neighborhood $\overline{\mathsf{B}}_1(\overline{\mathsf{c}}, \overline{\mathsf{a}})$ is paramterized by all possible lines in a through c, so we have $|X| = |\mathbb{F}|^2 + |\mathbb{F}| + 1$. In all cases, independence follows from the fact that for every $b_1 \in B$, the distribution which draws $a \sim A(b_1)$ and outputs $b_2 \sim B(a) \setminus \{b_1\}$ is the uniform distribution on B.

So now, let A/B be a bipartite graph which satisfies the pairwise independent parametrized neighborhood property described above. Let $B' \subset B$ be a subset of size $|B'| = \lambda \cdot |B|$. For $b \in B$, let $\mathbb{1}_{B'}(b)$ indicate whether $b \in B'$ or not, and let $\hat{\mathbb{1}}_{B'}(b) := \mathbb{1}_{B'}(b) - \lambda$. Note $\mathbb{E}_{b\sim B}\big[\hat{\mathbb{1}}_{B'}(b)\big] = 0$. Finally, define $f : A \to [0,1]$ by $f(a) := \mathbb{E}_{b\sim B(a)}\big[\hat{\mathbb{1}}_{B'}(b)\big]$. We will show $\mathbb{E}_{a\sim A}\big[f(a)^2\big] \leq |\mathbb{F}|^{-1}$. This suffices by Markov's inequality:

$$\Pr_{a\sim A}\left[\left|\Pr_{b\sim B(a)}\big(b \in B'\big) - \lambda\right| > |\mathbb{F}|^{-1/5}\right] \leq \Pr_{a\sim A}\Big[f(a)^2 > |\mathbb{F}|^{-2/5}\Big] \leq |\mathbb{F}|^{2/5} \cdot \mathbb{E}_{a\sim A}\big[f(a)^2\big].$$

We use the pairwise independence property to conclude:

$$\begin{aligned}\mathbb{E}_{a\sim A}\big[f(a)^2\big] &= \mathbb{E}_{a\sim A}\Big[\mathbb{E}_{x_1,x_2\sim X}\big[\hat{\mathbb{1}}_{B'}\big(a(x_1)\big) \cdot \hat{\mathbb{1}}_{B'}\big(a(x_2)\big)\big]\Big] \\ &\leq \frac{1}{|X|} + \mathbb{E}_{b_1,b_2\sim B}\Big[\hat{\mathbb{1}}_{B'}(b_1) \cdot \hat{\mathbb{1}}_{B'}(b_2)\Big] = \frac{1}{|X|}.\end{aligned}$$

For the reductions in the second phase, we use the generic facts about samplers stated in Section 2.2. Since $\overline{\mathsf{B}}_2/\overline{\mathsf{C}}$ and $\overline{\mathsf{B}}_2(\overline{\mathsf{c}})/\overline{\mathsf{B}}_1(\overline{\mathsf{c}})$ for all $\overline{\mathsf{c}} \in \overline{\mathsf{C}}$ are each $\big(|\mathbb{F}|^{-1/5}, |\mathbb{F}|^{-3/5}\big)-$sampling, $\overline{\mathsf{B}}_1/\overline{\mathsf{C}}$ and $\overline{\mathsf{B}}_2/\overline{\mathsf{B}}_1$ are both $\big(7 \cdot |\mathbb{F}|^{-1/5}, |\mathbb{F}|^{-1/5}\big)-$sampling (we have already shown sampling of $\overline{\mathsf{B}}_1/\overline{\mathsf{C}}$ with better parameters, sampling of $\overline{\mathsf{B}}_2/\overline{\mathsf{B}}_1$ follows from Fact 3. We have also shown that $\overline{\mathsf{A}}(\overline{\mathsf{b}}_1)/\overline{\mathsf{B}}_2(\overline{\mathsf{b}}_1)$ for all $\overline{\mathsf{b}}_1 \in \overline{\mathsf{B}}_1$ and $\overline{\mathsf{A}}(\overline{\mathsf{c}})/\overline{\mathsf{B}}_1(\overline{\mathsf{c}})$ for all $\overline{\mathsf{c}} \in \overline{\mathsf{C}}$ are both $\big(|\mathbb{F}|^{-1/5}, |\mathbb{F}|^{-3/5}\big)-$sampling. This fact combined with Fact 3 proves sampling of $\overline{\mathsf{A}}(\overline{\mathsf{c}})/\overline{\mathsf{B}}_2(\overline{\mathsf{c}})$ for all $\overline{\mathsf{c}} \in \overline{\mathsf{C}}$. The first point of Fact 2 says that any time we have Z such that $Z/\overline{\mathsf{B}}_1$ or $Z/\overline{\mathsf{B}}_2$ is $\mathcal{O}\big(|\mathbb{F}|^{-1}\big)-$biregular, then $Z/\overline{\mathsf{C}}$ or $Z/\overline{\mathsf{B}}_1$ is $\big(3 \cdot |\mathbb{F}|^{-1/15}, 3 \cdot |\mathbb{F}|^{-2/15}\big)-$sampling. This proves the sampling of all graphs except for (5) and (6): $\overline{\mathsf{A}}/\overline{\mathsf{C}}^2$ and $\overline{\mathsf{A}}(\overline{\mathsf{c}})/\overline{\mathsf{C}}^2$ for all $\overline{\mathsf{c}} \in \overline{\mathsf{C}}$, so it remains to prove sampling of these. Note $\overline{\mathsf{A}}(\overline{\mathsf{c}})/\overline{\mathsf{B}}_1$ for all $\overline{\mathsf{c}} \in \overline{\mathsf{C}}$ and $\overline{\mathsf{A}}/\overline{\mathsf{B}}_1$ are $\big(3 \cdot |\mathbb{F}|^{-1/15}, 3 \cdot |\mathbb{F}|^{-2/15}\big)-$samplers, since $\overline{\mathsf{A}}(\overline{\mathsf{c}})/\overline{\mathsf{B}}_2$ and $\overline{\mathsf{A}}/\overline{\mathsf{B}}_2$ are $\mathcal{O}\big(|\mathbb{F}|^{-1}\big)-$biregular. Thus we can use the second point Fact 2 to get $\big(12 \cdot |\mathbb{F}|^{-1/15}, |\mathbb{F}|^{-1/15}\big)-$sampling of graphs (5) and (6) because $\overline{\mathsf{B}}_1/\overline{\mathsf{C}}^2$ is $\mathcal{O}\big(|\mathbb{F}|^{-1}\big)-$biregular. $\square$

Notational Conventions and Example Use. Our proofs in the next sections rely heavily, and often implicitly, on the fact that the graphs of Claim 2 are samplers, and on the properties of sampler graphs stated in Fact 1. To facilitate readability, from here on, we reserve the quantity $\delta > 0$ for the loss introduced any time a sampling argument is used. As an example of how this looks in the body of the paper, let $\overline{\mathsf{C}}' \subset \overline{\mathsf{C}}$ be a set with $|\overline{\mathsf{C}}'| \geq \lambda \cdot |\overline{\mathsf{C}}|$, and let $\mathbf{E}$ be some event. Then we might deduce: $\mathbb{E}_{\overline{\mathsf{c}},\overline{\mathsf{c}}' \sim \overline{\mathsf{C}}'}\big[\Pr_{\overline{\mathsf{a}} \sim \overline{\mathsf{A}}(\overline{\mathsf{c}},\overline{\mathsf{c}}')}(\mathbf{E})\big] \geq \mathbb{E}_{\overline{\mathsf{a}} \sim \overline{\mathsf{A}}}\big[\Pr_{\overline{\mathsf{c}},\overline{\mathsf{c}}' \sim \overline{\mathsf{C}}'(\overline{\mathsf{a}})}(\mathbf{E})\big] - \delta$, "because of the sampling of $\overline{\mathsf{A}}/\overline{\mathsf{C}}^2$." Formally, we are using the third point of Fact 1, the fact that $\overline{\mathsf{A}}/\overline{\mathsf{C}}^2$ is $\eta'-$biregular, $(\varepsilon', \delta')-$sampling with $\lambda > \varepsilon'$ and that $\delta \geq \delta' + \eta'/\varepsilon'$.

Setting the Sampling Parameter. In the example use mentioned above, $\eta' = \mathcal{O}\big(|\mathbb{F}|^{-1}\big)$ and $\varepsilon', \delta' = \mathcal{O}\big(|\mathbb{F}|^{-1/15}\big)$. Thus, $\delta = \mathcal{O}\big(|\mathbb{F}|^{-1/15}\big)$ is sufficient for $\delta \geq \delta' + \eta'/\varepsilon'$ to hold. In general, each sampler property use will put a lower bound on δ, and so we simply set δ large enough so that they all hold. Explicitly, $\delta = 3 \cdot |\mathbb{F}|^{-1/60}$ is sufficient for our purposes.

We conclude this section with a claim listing two sampler-based facts which will be useful in the calculations in the next section.

Claim 3. *Let the notations be as above, and let $\delta = 3 \cdot |\mathbb{F}|^{-1/60}$ and $\eta = \mathcal{O}\big(|\mathbb{F}|^{-1}\big)$. Let $\overline{\mathsf{C}}' \subset \overline{\mathsf{C}}$ be a subset of size $|\overline{\mathsf{C}}'|/|\overline{\mathsf{C}}| \geq 12 \cdot |\mathbb{F}|^{-1/15}$. We have the following.*

1. $$\left\{ (\overline{\mathsf{c}}, \overline{\mathsf{b}}, \overline{\mathsf{c}}') \,\middle|\, \begin{array}{c} \overline{\mathsf{a}} \sim \overline{\mathsf{A}} \\ \overline{\mathsf{c}} \sim \overline{\mathsf{C}}(\overline{\mathsf{a}}) \\ \overline{\mathsf{c}}' \sim \overline{\mathsf{C}}'(\overline{\mathsf{a}}) \\ \mathsf{b} \sim \mathsf{B}_2(\mathsf{a}, \mathsf{c}, \mathsf{c}') \end{array} \right\} \approx_\delta \left\{ (\overline{\mathsf{c}}, \overline{\mathsf{b}}, \overline{\mathsf{c}}') \,\middle|\, \begin{array}{c} \overline{\mathsf{c}} \sim \overline{\mathsf{C}} \\ \overline{\mathsf{b}} \sim \overline{\mathsf{B}}_2(\overline{\mathsf{c}}) \\ \overline{\mathsf{c}}' \sim \overline{\mathsf{C}}'(\overline{\mathsf{b}}) \end{array} \right\},$$
   where in the first distribution $\overline{\mathsf{b}} = (\mathsf{b}, \alpha|_{\mathsf{b}})$, where $\overline{\mathsf{a}} = (\mathsf{a}, \alpha)$.

2. $$\left\{ (\overline{\mathsf{a}}, \overline{\mathsf{b}}, \overline{\mathsf{c}}') \,\middle|\, \begin{array}{c} \overline{\mathsf{a}} \sim \overline{\mathsf{A}} \\ \overline{\mathsf{c}} \sim \overline{\mathsf{C}}(\overline{\mathsf{a}}) \\ \overline{\mathsf{c}}' \sim \overline{\mathsf{C}}'(\overline{\mathsf{a}}) \\ \mathsf{b} \sim \mathsf{B}_2(\mathsf{c}, \mathsf{c}') \end{array} \right\} \approx_\delta \left\{ (\overline{\mathsf{a}}, \overline{\mathsf{b}}, \overline{\mathsf{c}}') \,\middle|\, \begin{array}{c} \overline{\mathsf{c}}' \sim \overline{\mathsf{C}}' \\ \overline{\mathsf{b}} \sim \overline{\mathsf{B}}_2(\overline{\mathsf{c}}') \\ \overline{\mathsf{a}} \sim \overline{\mathsf{A}}(\overline{\mathsf{b}}) \end{array} \right\},$$
   where in the first distribution $\overline{\mathsf{b}} = \overline{\mathsf{a}}|_{\mathsf{b}}$.

In both (1) and (2) above, $\approx_\delta$ denotes that the two distributions are within statistical distance δ of one another.

Proof. For the first part, we have

$$\left\{ \begin{array}{c} \overline{\mathsf{a}} \sim \overline{\mathsf{A}} \\ \overline{\mathsf{c}} \sim \overline{\mathsf{C}}(\overline{\mathsf{a}}) \\ \overline{\mathsf{c}}' \sim \overline{\mathsf{C}}'(\overline{\mathsf{a}}) \end{array} \right\} \approx_{\delta/3} \left\{ \begin{array}{c} \overline{\mathsf{c}}' \sim \overline{\mathsf{C}}' \\ \overline{\mathsf{a}} \sim \overline{\mathsf{A}}(\overline{\mathsf{c}}') \\ \overline{\mathsf{c}} \sim \overline{\mathsf{C}}(\overline{\mathsf{a}}) \end{array} \right\} \approx_{\eta} \left\{ \begin{array}{c} \overline{\mathsf{c}}' \sim \overline{\mathsf{C}}' \\ \overline{\mathsf{b}} \sim \overline{\mathsf{B}}_2(\overline{\mathsf{c}}') \\ \overline{\mathsf{c}} \sim \overline{\mathsf{C}}(\overline{\mathsf{b}}) \end{array} \right\} \approx_{\delta/3} \left\{ \begin{array}{c} \overline{\mathsf{b}} \sim \overline{\mathsf{B}}_2 \\ \overline{\mathsf{c}}' \sim \overline{\mathsf{C}}'(\overline{\mathsf{b}}) \\ \overline{\mathsf{c}} \sim \overline{\mathsf{C}}(\overline{\mathsf{b}}) \end{array} \right\},$$

where each distribution outputs $(\overline{\mathsf{c}}, \overline{\mathsf{b}}, \overline{\mathsf{c}}')$ and where $\overline{\mathsf{b}} = \overline{\mathsf{a}}|_{\mathsf{b}}$ for $\mathsf{b} \sim \mathsf{B}_2(\mathsf{a}, \mathsf{c}, \mathsf{c}')$ is implied in the first two distributions. The first relation follows from sampling of $\overline{\mathsf{A}}/\overline{\mathsf{C}}$; the second follows from the

$\eta-$biregularity of $\overline{\mathsf{B}}_2(\bar{\mathsf{a}},\bar{\mathsf{c}}')/\overline{\mathsf{C}}(\bar{\mathsf{a}})$ for all $\bar{\mathsf{a}}\in\overline{\mathsf{A}}$ and $\bar{\mathsf{c}}'\in\overline{\mathsf{C}}(\bar{\mathsf{a}})$, and the $0-$biregularity of $\overline{\mathsf{A}}(\bar{\mathsf{c}}')/\overline{\mathsf{B}}(\bar{\mathsf{c}}')$ for all $\bar{\mathsf{c}}'\in\overline{\mathsf{C}}$; the third follows from the sampling of $\overline{\mathsf{B}}_2/\overline{\mathsf{C}}$. Finally, the last distribution is identical to the desired distribution on the right of point 1 because of the $0-$biregularity of $\overline{\mathsf{B}}_2/\overline{\mathsf{C}}$. We work similarly for the second point:

$$\left\{\begin{array}{c}\bar{\mathsf{a}}\sim\overline{\mathsf{A}}\\ \bar{\mathsf{c}}\sim\overline{\mathsf{C}}(\bar{\mathsf{a}})\\ \bar{\mathsf{c}}'\sim\overline{\mathsf{C}}'(\bar{\mathsf{a}})\end{array}\right\}\approx_{\delta/2}\left\{\begin{array}{c}\bar{\mathsf{c}}'\sim\overline{\mathsf{C}}'\\ \bar{\mathsf{a}}\sim\overline{\mathsf{A}}(\bar{\mathsf{c}}')\\ \bar{\mathsf{c}}\sim\overline{\mathsf{C}}(\bar{\mathsf{a}})\end{array}\right\}\approx_{\eta}\left\{\begin{array}{c}\bar{\mathsf{c}}'\sim\overline{\mathsf{C}}'\\ \bar{\mathsf{a}}\sim\overline{\mathsf{A}}(\bar{\mathsf{c}}')\\ \bar{\mathsf{b}}\sim\overline{\mathsf{B}}_2(\bar{\mathsf{a}},\bar{\mathsf{c}}')\end{array}\right\}\equiv\left\{\begin{array}{c}\bar{\mathsf{c}}'\sim\overline{\mathsf{C}}'\\ \bar{\mathsf{b}}\sim\overline{\mathsf{B}}_2(\bar{\mathsf{c}}')\\ \bar{\mathsf{a}}\sim\overline{\mathsf{A}}(\bar{\mathsf{b}})\end{array}\right\},$$

where each distribution outputs $(\bar{\mathsf{a}},\bar{\mathsf{b}},\bar{\mathsf{c}}')$ and where $\bar{\mathsf{b}}=\bar{\mathsf{a}}|_{\mathsf{b}}$ (as above, $\mathsf{b}\sim\mathsf{B}_2(\mathsf{a},\mathsf{c},\mathsf{c}')$ is implicit in the first two distributions). We have used the sampling of $\overline{\mathsf{A}}/\overline{\mathsf{C}}$, $\eta-$biregularity of $\overline{\mathsf{B}}_2(\bar{\mathsf{a}},\bar{\mathsf{c}}')/\overline{\mathsf{C}}(\bar{\mathsf{a}})$ for all $\bar{\mathsf{a}}\in\overline{\mathsf{A}}$ and $\bar{\mathsf{c}}'\in\overline{\mathsf{C}}(\bar{\mathsf{a}})$, and $0-$biregularity of $\overline{\mathsf{A}}(\bar{\mathsf{c}}')/\overline{\mathsf{B}}_2(\bar{\mathsf{c}}')$ for all $\bar{\mathsf{c}}'\in\overline{\mathsf{C}}$. □

6 Global Agreement

In this section we prove Lemma 1, restated below in a quantitative form.

Lemma 1 (Restated). *Suppose $\varepsilon\geq\mathbb{F}^{-1/1000}$, and fix parameters $\eta=|\mathbb{F}|^{-9/10}$, $\delta=3\cdot|\mathbb{F}|^{-1/60}$, and $\tau=\mathcal{O}\big(\delta/\varepsilon^6+\eta/\varepsilon^{11}\big)$. Suppose $\{\mathsf{f}_{\mathsf{a}}\}_{\mathsf{a}}\subset\{\mathsf{f}:\Gamma_{\mathsf{A}}\to\Gamma_{\mathsf{A}}\}$ is such that*

$$\Pr_{\Phi,(\mathsf{c},\mathsf{a},\mathsf{a}')}\Big[\tilde{\alpha}|_{\mathsf{c}}=\tilde{\alpha}'|_{\mathsf{c}}\Big]=6\varepsilon \tag{6}$$

where the probability is over $\Phi\sim\Gamma$, $\mathsf{c}\sim\mathsf{C}$, $\mathsf{a},\mathsf{a}'\sim\mathsf{A}(\mathsf{c})$, and where $(\tilde{\alpha},\tilde{\alpha}')=\big(\mathsf{f}_{\mathsf{a}}(\Phi|_{\mathsf{a}}),\mathsf{f}_{\mathsf{a}'}(\Phi|_{\mathsf{a}'})\big)$. Then there exists a set $\mathsf{G}\subset\mathsf{A}\times\Gamma$ of size at least $|\mathsf{G}|\geq2\varepsilon\cdot|\mathsf{A}\times\Gamma|$ and a function $\mathsf{h}:\overline{\mathsf{C}}\to\Gamma_{\mathsf{C}}$ such that: $\Pr_{\substack{(\mathsf{a},\Phi)\sim\mathsf{G}\\ \mathsf{c}\sim\mathsf{C}(\mathsf{a})}}\big[\tilde{\gamma}\sim\tilde{\alpha}\big]\geq1-\zeta$, where $\tilde{\gamma}=\mathsf{h}(\mathsf{c},\Phi|_{\mathsf{c}})$ and $\zeta:=\varepsilon^{-2}\cdot(\tau+\delta)+\varepsilon^{-1}\cdot(\eta+\delta)$.

Remark. *Many different parameters are introduced during the course of our analysis which are all $\mathcal{O}\big(|\mathbb{F}|^{-1}\big)$. We encourage the reader to think of two levels of parameters: level one consists of ε only; all other parameters are in level 2 and are much smaller. The level two parameters are each defined to be smaller than ε^c for some constant $c=\mathcal{O}(1)$ which arises during our analysis. So in the above theorem, for example, in order for τ to be level 2, it must be that $\delta\ll\varepsilon^6$ and $\eta\ll\varepsilon^{11}$; additionally, for ζ to be level 2, $\tau\ll\varepsilon^2$ is required. We remark that the analysis prioritizes modularity and succinctness, rather than optimizing constants. As a result, the small constant $1/1000$ is suboptimal.*

We begin by introducing the notation and ideas needed to prove Lemma 1 in Section 6.1. The actual proof appears in Section 6.2, conditioned on two claims which we state in Section 6.1 and prove in Section 6.3.

6.1 Proof Setup.

Notations. In this section B denotes the set of $2-$dimensional planes in $\mathbb{F}^k$, and Γ_{B} is the set of $2-$variate polynomials over $\mathbb{F}$ of degree at most d, and $\overline{\mathsf{B}}=\mathsf{B}\times\Gamma_{\mathsf{B}}$. The sets $\overline{\mathsf{A}},\overline{\mathsf{C}},\Gamma$ are as usual.

We will take advantage of the sampling properties of the triple $\overline{\mathsf{A}}/\overline{\mathsf{B}}/\overline{\mathsf{C}}$. When considering two polynomials whose domains intersect, we write $\sim$ to indicate that they agree on the intersection. For example, given $\tilde{\alpha}, \tilde{\alpha}' \in \Gamma_{\mathsf{A}}$ defined on $\mathsf{a}, \mathsf{a}' \in \mathsf{A}(\overline{\mathsf{c}})$ we write $\tilde{\alpha} \sim \tilde{\alpha}'$ if $\tilde{\alpha}|_{\mathsf{c}} = \tilde{\alpha}'|_{\mathsf{c}}$.

We say that $(\mathsf{c}, \gamma, \tilde{\gamma})$ is *good* if $\Pr_{(\mathsf{a},\Phi)}\big[\tilde{\alpha} \sim \tilde{\gamma}\big] \geq 4\varepsilon$, where the probability is over $\mathsf{a} \sim \mathsf{A}(\mathsf{c})$ and $\Phi \sim \Gamma(\overline{\mathsf{c}})$. We say $\overline{\mathsf{c}} = (\mathsf{c}, \gamma)$ is *good* if there exists $\tilde{\gamma}$ such that $(\mathsf{c}, \gamma, \tilde{\gamma})$ is. Note that $\Pr_{\overline{\mathsf{c}} \sim \overline{\mathsf{C}}}\big[\overline{\mathsf{c}} \text{ good}\big] \geq 2\varepsilon$. To see this, let $\mathsf{p}_{\mathsf{c},\gamma,\tilde{\gamma}} := \Pr_{(\mathsf{a},\Phi)}\big[\tilde{\alpha} \sim \tilde{\gamma}\big]$. Then (6) gives

$$6\varepsilon = \mathbb{E}_{\overline{\mathsf{c}} \sim \overline{\mathsf{C}}}\left[\sum_{\tilde{\gamma}} \mathsf{p}_{\mathsf{c},\gamma,\tilde{\gamma}} \cdot \Pr_{\mathsf{a}' \sim \mathsf{A}(\mathsf{c})}\big[\tilde{\alpha}' \sim \tilde{\gamma}\big]\right] \leq \mathbb{E}_{\overline{\mathsf{c}} \sim \overline{\mathsf{C}}}\Big[\max_{\tilde{\gamma}}\big\{\mathsf{p}_{\mathsf{c},\gamma,\tilde{\gamma}}\big\}\Big].$$

We have used that $\sum_{\tilde{\gamma}} \Pr_{\mathsf{a}' \sim \mathsf{A}(\mathsf{c})}\big[\tilde{\alpha}' \sim \tilde{\gamma}\big] = 1$ for all $\overline{\mathsf{c}}$.

Local Functions. Let $\mathsf{h}_0 : \overline{\mathsf{C}} \to \Gamma_{\mathsf{C}}$ be the randomized function which sends $\overline{\mathsf{c}} = (\mathsf{c}, \gamma)$ to a random $\tilde{\gamma}$ such that $(\mathsf{c}, \gamma, \tilde{\gamma})$ is good if such $\tilde{\gamma}$ exists, and to an arbitrary $\tilde{\gamma} \in \Gamma_{\mathsf{C}}$ if not. For $\overline{\mathsf{c}} \in \overline{\mathsf{C}}$, let $\mathsf{g}_{\overline{\mathsf{c}}} : \overline{\mathsf{B}}(\overline{\mathsf{c}}) \to \Gamma_{\mathsf{B}}$ be the randomized function where $\mathsf{g}_{\overline{\mathsf{c}}}(\overline{\mathsf{b}})$ is the distribution on Γ_{B} which draws $\overline{\mathsf{a}} \sim \overline{\mathsf{A}}(\overline{\mathsf{b}})$ such that $\tilde{\alpha} \sim \mathsf{h}_0(\overline{\mathsf{c}})$, and outputs $\tilde{\beta} = \tilde{\alpha}|_{\mathsf{b}}$.

Definition 7 (Well-Defined). *Let $\eta = |\mathbb{F}|^{-9/10}$. We say that $\mathsf{g}_{\overline{\mathsf{c}}}$ is* well-defined *if*

$$\Pr_{\substack{\overline{\mathsf{b}} \sim \overline{\mathsf{B}}(\overline{\mathsf{c}}) \\ \overline{\mathsf{a}}, \overline{\mathsf{a}}' \sim \overline{\mathsf{A}}(\overline{\mathsf{b}})}}\Big[\tilde{\alpha} \approx \tilde{\alpha}' \Big| \tilde{\alpha} \sim \mathsf{h}_0(\overline{\mathsf{c}}) \sim \tilde{\alpha}'\Big] \geq 1 - \eta,$$

where $\tilde{\alpha} \approx \tilde{\alpha}'$ indicates that $\tilde{\alpha}|_{\mathsf{b}} = \tilde{\alpha}'|_{\mathsf{b}}$.

Previous work [IKW12, BDN17] refers to the good $\overline{\mathsf{c}} \in \overline{\mathsf{C}}$ for which $\mathsf{g}_{\overline{\mathsf{c}}}$ is well-defined as *excellent*; the fact that the excellent points comprise a non-negligible fraction of $\overline{\mathsf{C}}$ is a crucial component of the proofs in these papers. We require one extra property from our specialized subset of $\overline{\mathsf{C}}$ which simplifies the remainder of our proof greatly. The following is proved in Section 6.3.

Claim 4. *There exists a set $\overline{\mathsf{C}}' \subset \overline{\mathsf{C}}$ such that the following hold: 1) $|\overline{\mathsf{C}}'| \geq \varepsilon^3 |\overline{\mathsf{C}}|$; 2) every $\overline{\mathsf{c}} \in \overline{\mathsf{C}}'$ is good and such that $\mathsf{g}_{\overline{\mathsf{c}}}$ is well-defined; 3)*

$$\Pr_{\overline{\mathsf{c}}, \overline{\mathsf{c}}' \sim \overline{\mathsf{C}}'}\Big[\Pr_{\overline{\mathsf{a}} \sim \overline{\mathsf{A}}(\overline{\mathsf{c}}, \overline{\mathsf{c}}')}\big[\mathsf{h}_0(\overline{\mathsf{c}}) \sim \tilde{\alpha} \sim \mathsf{h}_0(\overline{\mathsf{c}}')\big] \geq \varepsilon^5\Big] \geq 1 - \sigma,$$

where $\sigma := \delta/\varepsilon^3 + \delta/\varepsilon^6 + \eta/\varepsilon^{11}$.

Intuitively, the extra property captured by (3) demands that the set of excellent points can be partitioned into large sets of *mutually compatible* points; the set $\overline{\mathsf{C}}'$ is any member of this partition.

The Global Function. Let $\mathsf{h} : \overline{\mathsf{C}} \to \Gamma_{\mathsf{C}}$ be the randomized function where $\mathsf{h}(\overline{\mathsf{c}})$ draws $\overline{\mathsf{b}} \sim \overline{\mathsf{B}}(\overline{\mathsf{c}})$, $\overline{\mathsf{c}}' \sim \overline{\mathsf{C}}'(\overline{\mathsf{b}})$ and outputs $\tilde{\beta}|_{\mathsf{c}}$ where $\tilde{\beta} = \mathsf{g}_{\overline{\mathsf{c}}'}(\overline{\mathsf{b}})$. The following is also proved in Section 6.3.

Claim 5. *We have $\Pr_{(\overline{\mathsf{c}}, \overline{\mathsf{b}}, \overline{\mathsf{c}}')}\big[\mathsf{h}(\overline{\mathsf{c}}) \sim \tilde{\beta}\big] \geq 1 - \tau$, where $\tau := \big(\sigma + 2\varepsilon^{-5}(\eta + \delta) + 2\delta\big)$, $\tilde{\beta} = \mathsf{g}_{\overline{\mathsf{c}}'}(\overline{\mathsf{b}})$ and the probability is over $\overline{\mathsf{c}} \sim \overline{\mathsf{C}}$, $\overline{\mathsf{b}} \sim \overline{\mathsf{B}}(\overline{\mathsf{c}})$, $\overline{\mathsf{c}}' \sim \overline{\mathsf{C}}'(\overline{\mathsf{b}})$.*

6.2 Proof of Lemma 1

Notational Convention. Let $\mathsf{h}_0, \mathsf{h} : \overline{\mathsf{C}} \to \Gamma_{\mathsf{C}}$ be the functions defined in Section 6.1. In this section if we write $\tilde{\gamma}$ when working with $\overline{\mathsf{c}} \in \overline{\mathsf{C}}$, it should be understood that $\tilde{\gamma} = \mathsf{h}(\overline{\mathsf{c}})$. We will always refer to $\mathsf{h}_0(\overline{\mathsf{c}})$ explicitly.

Proof. Suppose $(\varepsilon, \{\mathsf{f}_\mathsf{a}\})$ are such that (6) holds; let $\overline{\mathsf{C}}' \subset \overline{\mathsf{C}}$ be the set guaranteed by Claim 4. We define G to be the set of $(\mathsf{a}, \Phi) \in \mathsf{A} \times \Gamma$ such that $\Pr_{\overline{\mathsf{c}} \sim \overline{\mathsf{C}}'(\overline{\mathsf{a}})}\big[\tilde{\alpha} \sim \mathsf{h}_0(\overline{\mathsf{c}})\big] \geq \varepsilon$. We have,

$$\mathbb{E}_{(\mathsf{a},\Phi)\sim\mathsf{A}\times\Gamma}\Big[\Pr_{\overline{\mathsf{c}}\sim\overline{\mathsf{C}}'(\overline{\mathsf{a}})}\big[\tilde{\alpha} \sim \mathsf{h}_0(\overline{\mathsf{c}})\big]\Big] \geq \mathbb{E}_{\overline{\mathsf{c}}\sim\overline{\mathsf{C}}'}\Big[\Pr_{\substack{\mathsf{a}\sim\mathsf{A}(\mathsf{c}) \\ \Phi\sim\Gamma(\overline{\mathsf{c}})}}\big[\tilde{\alpha} \sim \mathsf{h}_0(\overline{\mathsf{c}})\big]\Big] - \delta \geq 3\varepsilon$$

We have used the sampling of $\mathsf{A} \times \Gamma / \overline{\mathsf{C}}$ for the first inequality, and that all $\overline{\mathsf{c}} \in \overline{\mathsf{C}}'$ are good for the second (and $4\varepsilon - \delta \geq 3\varepsilon$). It follows that $|\mathsf{G}| \geq 2\varepsilon|\mathsf{A} \times \Gamma|$. Thus, it remains to prove that $\Pr_{(\mathsf{a},\Phi),\mathsf{c}}\big[\tilde{\gamma} \sim \tilde{\alpha}\big] \geq 1 - \zeta$, where the probability is over $(\mathsf{a}, \Phi) \sim \mathsf{G}$, $\mathsf{c} \sim \mathsf{C}(\mathsf{a})$ and where $\tilde{\gamma} = \mathsf{h}(\mathsf{c}, \Phi|_{\mathsf{c}})$, where h is the global function defined in Section 6.1.

So let $\mathsf{p} := \Pr_{(\mathsf{a},\Phi),\mathsf{c}}\big[\tilde{\gamma} \sim \tilde{\alpha}\big]$ be the probability we are trying to bound. We have

$$\mathsf{p} \geq \Pr_{\substack{(\mathsf{a},\Phi) \\ \mathsf{b},\mathsf{c},\overline{\mathsf{c}}'}}\big[\tilde{\gamma} \sim \tilde{\beta} \sim \tilde{\alpha} \big| \tilde{\alpha} \sim \mathsf{h}_0(\overline{\mathsf{c}}')\big] \geq \Pr_{\substack{(\mathsf{a},\Phi) \\ \mathsf{b},\mathsf{c},\overline{\mathsf{c}}'}}\big[\tilde{\gamma} \sim \tilde{\beta} \big| \tilde{\alpha} \sim \mathsf{h}_0(\overline{\mathsf{c}}')\big] - \Pr_{\substack{(\mathsf{a},\Phi) \\ \mathsf{b},\mathsf{c},\overline{\mathsf{c}}'}}\big[\tilde{\beta} \not\sim \tilde{\alpha} \big| \tilde{\alpha} \sim \mathsf{h}_0(\overline{\mathsf{c}}')\big],$$

where the probabilities are over $(\mathsf{a}, \Phi) \sim \mathsf{G}$, $\mathsf{c} \sim \mathsf{C}(\mathsf{a})$, $\overline{\mathsf{c}}' \sim \overline{\mathsf{C}}'(\overline{\mathsf{a}})$, $\mathsf{b} \sim \mathsf{B}(\mathsf{a}, \mathsf{c}, \mathsf{c}')$, and where $\tilde{\beta} = \mathsf{g}_{\overline{\mathsf{c}}'}(\overline{\mathsf{b}})$, for $\overline{\mathsf{b}} = (\mathsf{b}, \Phi|_{\mathsf{b}})$. We conclude by bounding both probabilities on the right; denoted RHS_1 and RHS_2, respectively. We have

$$\begin{aligned} 1 - \mathsf{RHS}_1 &= \Pr_{\substack{(\mathsf{a},\Phi)\sim\mathsf{G} \\ \mathsf{b},\mathsf{c},\overline{\mathsf{c}}'}}\big[\tilde{\gamma} \not\sim \tilde{\beta} \big| \tilde{\alpha} \sim \mathsf{h}_0(\overline{\mathsf{c}}')\big] \leq \frac{\Pr_{(\mathsf{a},\Phi),\mathsf{b},\mathsf{c},\overline{\mathsf{c}}'}\big[\tilde{\gamma} \not\sim \tilde{\beta}\big]}{\min_{(\mathsf{a},\Phi)\in\mathsf{G}}\big\{\Pr_{\overline{\mathsf{c}}'\sim\overline{\mathsf{C}}'(\overline{\mathsf{a}})}\big[\tilde{\alpha} \sim \mathsf{h}_0(\overline{\mathsf{c}}')\big]\big\}} \\ &\leq \frac{\varepsilon^{-2}}{2} \cdot \Pr_{\substack{\overline{\mathsf{a}}\sim\overline{\mathsf{A}} \\ \mathsf{b},\mathsf{c},\overline{\mathsf{c}}'}}\big[\tilde{\gamma} \not\sim \tilde{\beta}\big] < \varepsilon^{-2} \cdot \Big(\Pr_{\substack{\overline{\mathsf{c}}\sim\overline{\mathsf{C}} \\ \overline{\mathsf{b}}\sim\overline{\mathsf{B}}(\overline{\mathsf{c}}) \\ \overline{\mathsf{c}}'\sim\overline{\mathsf{C}}'(\overline{\mathsf{b}})}}\big[\tilde{\gamma} \not\sim \tilde{\beta}\big] + \delta\Big) \leq \varepsilon^{-2} \cdot (\tau + \delta). \end{aligned}$$

The first inequality on the second line used the definition of G and that $|\mathsf{G}| \geq 2\varepsilon \cdot |\mathsf{A} \times \Gamma|$; the second used Claim 3, point 1; and the last used Claim 5. Finally,

$$\begin{aligned} \mathsf{RHS}_2 &\leq \frac{\varepsilon^{-1}}{2} \cdot \Pr_{\substack{\overline{\mathsf{a}}\sim\overline{\mathsf{A}} \\ \overline{\mathsf{c}}'\sim\overline{\mathsf{C}}'(\overline{\mathsf{a}}) \\ \overline{\mathsf{b}}\sim\overline{\mathsf{B}}(\overline{\mathsf{c}}',\overline{\mathsf{a}})}}\big[\tilde{\beta} \not\sim \tilde{\alpha} \big| \tilde{\alpha} \sim \mathsf{h}_0(\overline{\mathsf{c}}')\big] \\ &\leq \varepsilon^{-1} \cdot \left(\max_{\overline{\mathsf{c}}'\in\overline{\mathsf{C}}'}\Big\{\Pr_{\substack{\overline{\mathsf{b}}\sim\overline{\mathsf{B}}(\overline{\mathsf{c}}') \\ \overline{\mathsf{a}}\sim\overline{\mathsf{A}}(\overline{\mathsf{b}})}}\big[\tilde{\beta} \not\sim \tilde{\alpha} \big| \tilde{\alpha} \sim \mathsf{h}_0(\overline{\mathsf{c}}')\big]\Big\} + \delta\right) \leq \varepsilon^{-1}\big(\eta + \delta\big). \end{aligned}$$

We have used Claim 3 point 2 and the fact that $\mathsf{g}_{\overline{\mathsf{c}}'}$ is well-defined for all $\overline{\mathsf{c}}' \in \overline{\mathsf{C}}'$. The result follows. □

6.3 Proving the Claims

Starting Assumption and Notational Conventions. Throughout this section, we assume the hypotheses of Lemma 1, namely $(\varepsilon, \{\mathsf{f_a}\})$ are such that $\Pr_{\Phi,(\mathsf{c},\mathsf{a},\mathsf{a}')}\big[\tilde{\alpha} \sim \tilde{\alpha}'\big] = 6\varepsilon$ (*i.e.*, such that (6) holds). Let $\mathsf{h}_0, \mathsf{h} : \overline{\mathsf{C}} \to \Gamma_{\mathsf{C}}$ be the functions defined in Section 6.1. In this section if we write $\tilde{\gamma}$ when working with $\overline{\mathsf{c}} \in \overline{\mathsf{C}}$, it should be understood that $\tilde{\gamma} = \mathsf{h}_0(\overline{\mathsf{c}})$. We will refer to $\mathsf{h}(\overline{\mathsf{c}})$ explicitly (note, this is opposite to the convention of Section 6.2). Given $\overline{\mathsf{c}}, \overline{\mathsf{c}}' \in \overline{\mathsf{C}}$ set $\mu_{\overline{\mathsf{c}}}$, $\mathsf{p}(\overline{\mathsf{c}})$ and $\mathsf{q}(\overline{\mathsf{c}}, \overline{\mathsf{c}}')$ to:

$$\Pr_{\substack{\mathsf{a}\sim\mathsf{A}(\mathsf{c})\\ \Phi\sim\Gamma(\overline{\mathsf{c}})}}\big[\tilde{\gamma} \sim \tilde{\alpha}\big];\ \Pr_{\substack{\overline{\mathsf{b}}\sim\overline{\mathsf{B}}(\overline{\mathsf{c}})\\ \overline{\mathsf{a}}\sim\overline{\mathsf{A}}(\overline{\mathsf{b}})}}\big[\tilde{\beta} \sim \tilde{\alpha} \big| \tilde{\gamma} \sim \tilde{\alpha}\big];\ \Pr_{\overline{\mathsf{a}}\sim\overline{\mathsf{A}}(\overline{\mathsf{c}},\overline{\mathsf{c}}')}\big[\tilde{\gamma} \sim \tilde{\alpha} \sim \tilde{\gamma}'\big].$$

In Section 6.1 we called $\overline{\mathsf{c}} \in \overline{\mathsf{C}}$ such that $\mu_{\overline{\mathsf{c}}} \geq 4\varepsilon$ *good*. Also for $\overline{\mathsf{c}} \in \overline{\mathsf{C}}$ we defined local functions $\mathsf{g}_{\overline{\mathsf{c}}} : \overline{\mathsf{B}}(\overline{\mathsf{c}}) \to \Gamma_{\mathsf{B}}$ and said that $\mathsf{g}_{\overline{\mathsf{c}}}$ was *well-defined* if $\mathsf{p}(\overline{\mathsf{c}}) \geq 1 - \eta$, where $\eta = |\mathbb{F}|^{-9/10}$. In the remainder of this section we prove three claims; the first two combine to prove Claim 4, the last is Claim 5.

Claim 6. *There exists a set $\overline{\mathsf{C}}_0' \subset \overline{\mathsf{C}}$ such that the following hold: 1) $|\overline{\mathsf{C}}_0'| \geq \varepsilon|\overline{\mathsf{C}}|$; 2) $\mu_{\overline{\mathsf{c}}} \geq 4\varepsilon$ for every $\overline{\mathsf{c}} \in \overline{\mathsf{C}}_0'$; 3) $\mathsf{p}(\overline{\mathsf{c}}) \geq 1 - \eta$ for every $\overline{\mathsf{c}} \in \overline{\mathsf{C}}_0'$.*

Proof. Let $\overline{\mathsf{C}}_0' \subset \overline{\mathsf{C}}$ be the set of $\overline{\mathsf{c}} \in \overline{\mathsf{C}}$ for which $\mu_{\overline{\mathsf{c}}} \geq 4\varepsilon$ and $\mathsf{p}(\overline{\mathsf{c}}) \geq 1 - \eta$ (*i.e.*, $\overline{\mathsf{c}} \in \overline{\mathsf{C}}_0'$ if $\overline{\mathsf{c}}$ is good and such that $\mathsf{g}_{\overline{\mathsf{c}}}$ is well-defined). We bound $|\overline{\mathsf{C}}_0'|$ using three observations. First, as noted in Section 6.1, $\Pr_{\overline{\mathsf{c}}\sim\overline{\mathsf{C}}}\big[\mu_{\overline{\mathsf{c}}} \geq 4\varepsilon\big] \geq 2\varepsilon$. Second, for all $\overline{\mathsf{c}} \in \overline{\mathsf{C}}$ such that $\mu_{\overline{\mathsf{c}}} \geq 4\varepsilon$:

$$\Pr_{\substack{\overline{\mathsf{b}}\sim\overline{\mathsf{B}}(\overline{\mathsf{c}})\\ \overline{\mathsf{a}},\overline{\mathsf{a}}'\sim\overline{\mathsf{A}}(\overline{\mathsf{b}})}}\big[\tilde{\alpha} \sim \tilde{\gamma} \sim \tilde{\alpha}'\big] = \mathbb{E}_{\overline{\mathsf{b}}\sim\overline{\mathsf{B}}(\overline{\mathsf{c}})}\Big[\mu_{\overline{\mathsf{c}}}(\overline{\mathsf{b}})^2\Big] \geq \Pr_{\overline{\mathsf{b}}\sim\overline{\mathsf{B}}(\overline{\mathsf{c}})}\Big[\big|\mu_{\overline{\mathsf{c}}}(\overline{\mathsf{b}}) - \mu_{\overline{\mathsf{c}}}\big| \leq \varepsilon\Big] \cdot 9\varepsilon^2 \geq \varepsilon^2,$$

where $\mu_{\overline{\mathsf{c}}}(\overline{\mathsf{b}}) := \Pr_{\overline{\mathsf{a}}\sim\overline{\mathsf{A}}(\overline{\mathsf{b}})}\big[\tilde{\alpha} \sim \tilde{\gamma}\big]$ is shorthand. We have used the sampling of $\overline{\mathsf{A}}(\overline{\mathsf{c}})/\overline{\mathsf{B}}(\overline{\mathsf{c}})$ to (crudely) lower bound $\Pr_{\overline{\mathsf{b}}\sim\overline{\mathsf{B}}(\overline{\mathsf{c}})}\big[|\mu_{\overline{\mathsf{c}}}(\overline{\mathsf{b}}) - \mu_{\overline{\mathsf{c}}}| \leq \varepsilon\big]$. Finally, by Markov's inequality and Schwartz-Zippel:

$$\Pr_{\overline{\mathsf{c}}\sim\overline{\mathsf{C}}}\left[\Pr_{\substack{\overline{\mathsf{b}}\sim\overline{\mathsf{B}}(\overline{\mathsf{c}})\\ \overline{\mathsf{a}},\overline{\mathsf{a}}'\sim\overline{\mathsf{A}}(\overline{\mathsf{b}})}}\Big[\tilde{\alpha} \not\approx \tilde{\alpha}' \ \&\ \tilde{\alpha} \sim \tilde{\gamma} \sim \tilde{\alpha}'\Big] > \eta\varepsilon^2\right] \leq \frac{d}{\eta\varepsilon^2|\mathbb{F}|}.$$

Putting these together gives

$$\begin{aligned}
\frac{|\overline{\mathsf{C}}_0'|}{|\overline{\mathsf{C}}|} &= \Pr_{\overline{\mathsf{c}}\sim\overline{\mathsf{C}}}\left[\mu_{\overline{\mathsf{c}}} \geq 4\varepsilon \ \&\ \Pr_{\substack{\overline{\mathsf{b}}\sim\overline{\mathsf{B}}(\overline{\mathsf{c}})\\ \overline{\mathsf{a}},\overline{\mathsf{a}}'\sim\overline{\mathsf{A}}(\overline{\mathsf{b}})}}\Big[\tilde{\alpha} \not\approx \tilde{\alpha}' \Big| \tilde{\alpha} \sim \tilde{\gamma} \sim \tilde{\alpha}'\Big] \leq \eta\right] \\
&\geq \Pr_{\overline{\mathsf{c}}\sim\overline{\mathsf{C}}}\big[\mu_{\overline{\mathsf{c}}} \geq 4\varepsilon\big] - \Pr_{\overline{\mathsf{c}}\sim\overline{\mathsf{C}}}\left[\Pr_{(\overline{\mathsf{b}},\overline{\mathsf{a}},\overline{\mathsf{a}}')}\Big[\tilde{\alpha} \not\approx \tilde{\alpha}' \ \&\ \tilde{\alpha} \sim \tilde{\gamma} \sim \tilde{\alpha}'\Big] > \eta\varepsilon^2\right] \\
&\geq 2\varepsilon - \frac{d}{\eta\varepsilon^2|\mathbb{F}|} \geq \varepsilon.
\end{aligned}$$

□

Claim 4 (Restated). There exists a set $\overline{\mathsf{C}}' \subset \overline{\mathsf{C}}$ such that the following hold: 1) $|\overline{\mathsf{C}}'| \geq \varepsilon^3|\overline{\mathsf{C}}|$; 2) $\mu_{\overline{\mathsf{c}}} \geq 4\varepsilon$ for every $\overline{\mathsf{c}} \in \overline{\mathsf{C}}'$; 3) $\mathsf{p}(\overline{\mathsf{c}}) \geq 1 - \eta$ for every $\overline{\mathsf{c}} \in \overline{\mathsf{C}}'$; 4) $\Pr_{\overline{\mathsf{c}},\overline{\mathsf{c}}'\sim\overline{\mathsf{C}}'}\big[\mathsf{q}(\overline{\mathsf{c}}, \overline{\mathsf{c}}') \geq \varepsilon^5\big] \geq 1 - \sigma$, where $\sigma := \delta/\varepsilon^3 + \delta/\varepsilon^6 + \eta/\varepsilon^{11}$.

Proof. By Claim 6 it suffices to construct a large subset of $\overline{\mathsf{C}}_0'$ such that the fourth property holds. For this purpose, we equip $\overline{\mathsf{C}}_0'$ with a graph structure: $\overline{\mathsf{c}}, \overline{\mathsf{c}}' \in \overline{\mathsf{C}}_0'$ are adjacent if $\mathsf{q}(\overline{\mathsf{c}}, \overline{\mathsf{c}}') \geq \varepsilon^2$. Our final set $\overline{\mathsf{C}}'$ will be the neighborhood, $\mathsf{N}(\overline{\mathsf{c}}') := \{\overline{\mathsf{c}} \in \overline{\mathsf{C}}_0' : \mathsf{q}(\overline{\mathsf{c}}, \overline{\mathsf{c}}') \geq \varepsilon^2\}$ of some $\overline{\mathsf{c}}' \in \overline{\mathsf{C}}_0'$. In order for this to work, $\overline{\mathsf{c}}'$ should satisfy: 1) $\big|\mathsf{N}(\overline{\mathsf{c}}')\big|$ must be large; 2) $\Pr_{\overline{\mathsf{c}}, \overline{\mathsf{c}}'' \sim \mathsf{N}(\overline{\mathsf{c}}')}[\mathsf{q}(\overline{\mathsf{c}}, \overline{\mathsf{c}}'') < \varepsilon^5]$ must be small. We show there exists such a $\overline{\mathsf{c}}' \in \overline{\mathsf{C}}_0'$. Specifically we prove

1. $\mathbb{E}_{\overline{\mathsf{c}}, \overline{\mathsf{c}}' \sim \overline{\mathsf{C}}_0'}\big[\mathsf{q}(\overline{\mathsf{c}}, \overline{\mathsf{c}}')\big] \geq 3\varepsilon^2$; and
2. $\Pr_{\substack{\overline{\mathsf{c}}' \sim \overline{\mathsf{C}}_0' \\ \overline{\mathsf{c}}, \overline{\mathsf{c}}'' \sim \mathsf{N}(\overline{\mathsf{c}}')}}\Big[\mathsf{q}(\overline{\mathsf{c}}, \overline{\mathsf{c}}'') \geq \varepsilon^5 \Big| |\mathsf{N}(\overline{\mathsf{c}}')| > \varepsilon^3 |\overline{\mathsf{C}}|\Big] \geq 1 - \sigma.$

It follows from the first point that $\Pr_{\overline{\mathsf{c}}' \sim \overline{\mathsf{C}}_0'}\big[|\mathsf{N}(\overline{\mathsf{c}}')| \geq \varepsilon^3 |\overline{\mathsf{C}}|\big] > \varepsilon^2$ (using $|\overline{\mathsf{C}}_0'| \geq \varepsilon |\overline{\mathsf{C}}|$). Thus, the two points together guarantee the existence of some $\overline{\mathsf{c}}' \in \overline{\mathsf{C}}_0'$ such that $|\mathsf{N}(\overline{\mathsf{c}}')| \geq \varepsilon^3 |\overline{\mathsf{C}}|$ and $\Pr_{\overline{\mathsf{c}}, \overline{\mathsf{c}}'' \sim \mathsf{N}(\overline{\mathsf{c}}')}\big[\mathsf{q}(\overline{\mathsf{c}}, \overline{\mathsf{c}}'') \geq \varepsilon^5\big] \geq 1 - \sigma$. Setting $\overline{\mathsf{C}}' = \mathsf{N}(\overline{\mathsf{c}}')$ for such a $\overline{\mathsf{c}}' \in \overline{\mathsf{C}}_0'$ completes the proof. So it remains to establish the above two bounds.

For the first, we have

$$\begin{aligned}
\mathbb{E}_{\overline{\mathsf{c}}, \overline{\mathsf{c}}' \sim \overline{\mathsf{C}}_0'}\big[\mathsf{q}(\overline{\mathsf{c}}, \overline{\mathsf{c}}')\big] &\geq \mathbb{E}_{\overline{\mathsf{a}} \sim \overline{\mathsf{A}}}\Big[\Pr_{\overline{\mathsf{c}} \sim \overline{\mathsf{C}}_0'(\overline{\mathsf{a}})}\big[\tilde{\gamma} \sim \tilde{\alpha}\big]^2\Big] - \delta \geq \mathbb{E}_{\overline{\mathsf{a}} \sim \overline{\mathsf{A}}}\Big[\Pr_{\overline{\mathsf{c}} \sim \overline{\mathsf{C}}_0'(\overline{\mathsf{a}})}\big[\tilde{\gamma} \sim \tilde{\alpha}\big]\Big]^2 - \delta \\
&\geq \mathbb{E}_{\overline{\mathsf{c}} \sim \overline{\mathsf{C}}_0'}\big[\mu_{\overline{\mathsf{c}}}\big]^2 - 3\delta \geq 16\varepsilon^2 - 3\delta \geq 3\varepsilon^2.
\end{aligned}$$

We have used the sampling of $\overline{\mathsf{A}}/\overline{\mathsf{C}}^2$, Jensen's inequality, the sampling of $\overline{\mathsf{A}}/\overline{\mathsf{C}}$, and the fact that $\mu_{\overline{\mathsf{c}}} \geq 4\varepsilon$ for all $\overline{\mathsf{c}} \in \overline{\mathsf{C}}_0'$. Establishing the second bound is more involved. Towards this end, we define three quantities, shorthanded as $\mathsf{val}_1, \mathsf{val}_2, \mathsf{val}_3$; each is a function of $(\overline{\mathsf{c}}, \overline{\mathsf{c}}', \overline{\mathsf{c}}'')$:

- $\mathsf{val}_1 := \Big|\Pr_{\overline{\mathsf{a}}' \sim \overline{\mathsf{A}}(\overline{\mathsf{c}}, \overline{\mathsf{c}}', \overline{\mathsf{c}}'')}\big[\tilde{\gamma} \sim \tilde{\alpha}' \sim \tilde{\gamma}''\big] - \mathsf{q}(\overline{\mathsf{c}}, \overline{\mathsf{c}}'')\Big|;$
- $\mathsf{val}_2 := \Big|\Pr_{\overline{\mathsf{a}}' \sim \overline{\mathsf{A}}(\overline{\mathsf{c}}, \overline{\mathsf{c}}', \overline{\mathsf{c}}'')}[\tilde{\gamma}' \sim \tilde{\alpha}'] - \mu_{\overline{\mathsf{c}}'}\Big|;$
- $\mathsf{val}_3 := \Pr_{\substack{\overline{\mathsf{b}} \sim \overline{\mathsf{B}}(\overline{\mathsf{c}}, \overline{\mathsf{c}}') \\ \overline{\mathsf{a}} \sim \overline{\mathsf{A}}(\overline{\mathsf{b}}) \\ \overline{\mathsf{a}}' \sim \overline{\mathsf{A}}(\overline{\mathsf{b}}, \overline{\mathsf{c}}'')}}\big[\tilde{\alpha} \not\approx \tilde{\alpha}' \big| \tilde{\alpha} \sim \tilde{\gamma}' \sim \tilde{\alpha}'\big] + \Pr_{\substack{\overline{\mathsf{b}}'' \sim \overline{\mathsf{B}}(\overline{\mathsf{c}}', \overline{\mathsf{c}}'') \\ \overline{\mathsf{a}}'' \sim \overline{\mathsf{A}}(\overline{\mathsf{b}}'') \\ \overline{\mathsf{a}}' \sim \overline{\mathsf{A}}(\overline{\mathsf{b}}'', \overline{\mathsf{c}})}}\big[\tilde{\alpha}' \not\approx \tilde{\alpha}'' \big| \tilde{\alpha}' \sim \tilde{\gamma}' \sim \tilde{\alpha}''\big].$

We show that each val_i is small with very high probability over $(\overline{\mathsf{c}}, \overline{\mathsf{c}}', \overline{\mathsf{c}}'')$ drawn as follows: $\overline{\mathsf{c}}' \sim \overline{\mathsf{C}}_0'$ such that $|\mathsf{N}(\overline{\mathsf{c}}')| \geq \varepsilon^3 |\overline{\mathsf{C}}|$, $\overline{\mathsf{c}}, \overline{\mathsf{c}}'' \sim \mathsf{N}(\overline{\mathsf{c}}')$. These bounds will be used in the computation which follows. We have

$$\Pr_{(\overline{\mathsf{c}}, \overline{\mathsf{c}}', \overline{\mathsf{c}}'')}\big[\mathsf{val}_1 > \delta\big] \leq \varepsilon^{-3} \cdot \max_{\overline{\mathsf{c}}, \overline{\mathsf{c}}'' \in \overline{\mathsf{C}}}\left\{\Pr_{\overline{\mathsf{c}}' \sim \overline{\mathsf{C}}}\left[\left|\mathbb{E}_{\overline{\mathsf{a}}' \sim \overline{\mathsf{A}}(\overline{\mathsf{c}}, \overline{\mathsf{c}}', \overline{\mathsf{c}}'')}\big[f_1(\overline{\mathsf{a}}')\big] - \mathbb{E}_{\overline{\mathsf{a}}' \sim \overline{\mathsf{A}}(\overline{\mathsf{c}}, \overline{\mathsf{c}}'')}\big[f_1(\overline{\mathsf{a}}')\big]\right| > \delta\right]\right\},$$

where $f_1(\overline{\mathsf{a}}') = 1$ if $\tilde{\gamma} \sim \tilde{\alpha}' \sim \tilde{\gamma}''$, 0 otherwise. Thus $\Pr_{(\overline{\mathsf{c}}, \overline{\mathsf{c}}', \overline{\mathsf{c}}'')}\big[\mathsf{val}_1 > \delta\big] \leq \delta/\varepsilon^3$, by the sampling of $\overline{\mathsf{A}}(\overline{\mathsf{c}}, \overline{\mathsf{c}}'')/\overline{\mathsf{C}}$ for all $\overline{\mathsf{c}}, \overline{\mathsf{c}}'' \in \overline{\mathsf{C}}$. Likewise, $\Pr_{(\overline{\mathsf{c}}, \overline{\mathsf{c}}', \overline{\mathsf{c}}'')}\big[\mathsf{val}_2 > \delta\big] \leq \delta/\varepsilon^6$ follows from the same reasoning using the sampling of $\overline{\mathsf{A}}(\overline{\mathsf{c}}')/\overline{\mathsf{C}}^2$ and the function $f_2(\overline{\mathsf{a}}') = 1$ iff $\tilde{\gamma}' \sim \tilde{\alpha}'$. Finally,

$$\Pr_{(\overline{\mathsf{c}}, \overline{\mathsf{c}}', \overline{\mathsf{c}}'')}\big[\mathsf{val}_3 > 2\varepsilon^5\big] \leq \varepsilon^{-6} \cdot \max_{\overline{\mathsf{c}}' \in \overline{\mathsf{C}}_0'}\Big\{\Pr_{\overline{\mathsf{c}}, \overline{\mathsf{c}}'' \sim \overline{\mathsf{C}}}\big[\mathsf{val}_3 > 2\varepsilon^5\big]\Big\} \leq \frac{\varepsilon^{-11}}{2} \cdot \max_{\overline{\mathsf{c}}' \in \overline{\mathsf{C}}_0'}\Big\{\mathbb{E}_{\overline{\mathsf{c}}, \overline{\mathsf{c}}'' \sim \overline{\mathsf{C}}}\big[\mathsf{val}_3\big]\Big\}$$

$$= \frac{\varepsilon^{-11}}{2} \cdot \max_{\overline{c}' \in \overline{C}'_0} \left\{ 2 \cdot \big(1 - \mathsf{p}(\overline{c}')\big) \right\} \leq \eta/\varepsilon^{11}.$$

Now we show how these values figure into deriving the bound we need. The key point is that they let us bound $\mathsf{q}(\overline{c}, \overline{c}'')$ in terms of $\mathsf{q}(\overline{c}, \overline{c}') \cdot \mathsf{q}(\overline{c}', \overline{c}'') \cdot \mu_{\overline{c}'}$, which is large when $\overline{c}, \overline{c}'' \in \mathsf{N}(\overline{c}')$ and $\overline{c}' \in \overline{C}'_0$. We have:

$$\begin{aligned}
\mathsf{q}(\overline{c}, \overline{c}'') &= \Pr_{\overline{a}' \sim \overline{A}(\overline{c}, \overline{c}'')} \big[\tilde{\gamma} \sim \tilde{\alpha}' \sim \tilde{\gamma}''\big] \geq \Pr_{\overline{a}' \sim \overline{A}(\overline{c}, \overline{c}', \overline{c}'')} \big[\tilde{\gamma} \sim \tilde{\alpha}' \sim \tilde{\gamma}''\big] - \mathsf{val}_1 \\
&\geq \Pr_{\substack{\overline{a}' \sim \overline{A}(\overline{c}, \overline{c}', \overline{c}'') \\ \overline{a} \sim \overline{A}(\overline{b}) \\ \overline{a}'' \sim \overline{A}(\overline{b}'')}} \big[\tilde{\gamma} \sim \tilde{\alpha} \sim \tilde{\gamma}' \sim \tilde{\alpha}'' \sim \tilde{\gamma}'' \;\&\; \tilde{\alpha} \approx \tilde{\alpha}' \approx \tilde{\alpha}'' \;\&\; \tilde{\gamma}' \sim \tilde{\alpha}'\big] - \mathsf{val}_1 \\
&\geq \Pr_{\substack{\overline{a} \sim \overline{A}(\overline{c}, \overline{c}') \\ \overline{a}' \sim \overline{A}(\overline{c}, \overline{c}', \overline{c}'') \\ \overline{a}'' \sim \overline{A}(\overline{c}', \overline{c}'')}} \big[\tilde{\gamma} \sim \tilde{\alpha} \sim \tilde{\gamma}' \;\&\; \tilde{\gamma}' \sim \tilde{\alpha}' \;\&\; \tilde{\gamma}' \sim \tilde{\alpha}'' \sim \tilde{\gamma}''\big] - \mathsf{val}_1 - \mathsf{val}_3 \\
&\geq \mathsf{q}(\overline{c}, \overline{c}') \cdot \mathsf{q}(\overline{c}', \overline{c}'') \cdot \mu_{\overline{c}'} - \mathsf{val}_1 - \mathsf{val}_2 - \mathsf{val}_3 \geq 4\varepsilon^5 - \mathsf{val}_1 - \mathsf{val}_2 - \mathsf{val}_3.
\end{aligned}$$

In the probability subscript in the second line, $\overline{b}$ and $\overline{b}''$ are the restrictions of $\overline{a}'$ to the lines spanned by (c, c') and (c', c''), respectively. The result follows:

$$\Pr_{\substack{\overline{c}' \sim \overline{C}'_0 \\ \overline{c}, \overline{c}'' \sim \mathsf{N}(\overline{c}')}} \left[\mathsf{q}(\overline{c}, \overline{c}'') \geq \varepsilon^5 \middle| |\mathsf{N}(\overline{c}')| > \varepsilon^3 |\overline{C}| \right] \geq \Pr_{(\overline{c}, \overline{c}', \overline{c}'')} \left[\mathsf{val}_1 + \mathsf{val}_2 + \mathsf{val}_3 \leq 3\varepsilon^5 \right] \geq 1 - \sigma.$$

□

Claim 5 (Restated). *We have*

$$\Pr_{\substack{\overline{c} \sim \overline{C} \\ \overline{b}_1 \sim \overline{B}(\overline{c}) \\ \overline{c}'_1 \sim \overline{C}'(\overline{b}_1)}} \big[\mathsf{h}(\overline{c}) \sim \tilde{\beta}_1\big] \geq 1 - \tau,$$

where $\tilde{\beta} = \mathsf{g}_{\overline{c}'}(\overline{b})$*, and where* $\tau := \big(\sigma + 2\varepsilon^{-5}(\eta + \delta) + 2\delta\big)$*. Recall* $\mathsf{h}(\overline{c})$ *is the distribution on* Γ_{C} *which draws* $\overline{b}'_2 \sim \overline{B}(\overline{c})$*,* $\overline{c}'_2 \sim \overline{C}'(\overline{b}_2)$ *and outputs* $\mathsf{g}_{\overline{c}'_2}(\overline{b}_2)|_{\mathsf{c}}$*.*

Proof. We show $\Pr_{(\overline{c}, \overline{c}'_1, \overline{c}'_2, \overline{b}_1, \overline{b}_2)} \big[\tilde{\beta}_1 \sim \tilde{\beta}_2\big] \geq 1 - \big(\sigma + 2\varepsilon^{-5}(\eta + \delta)\big)$, where the probability is over $\overline{c} \sim \overline{C}$, $\overline{c}'_1, \overline{c}'_2 \sim \overline{C}'$, $\overline{b}_1 \sim \overline{B}(\overline{c}, \overline{c}'_1)$, $\overline{b}_2 \sim \overline{B}(\overline{c}, \overline{c}'_2)$ and where $\tilde{\beta}_1 \sim \tilde{\beta}_2$ means that $\mathsf{g}_{\overline{c}'_1}(\overline{b}_1)$ and $\mathsf{g}_{\overline{c}'_2}(\overline{b}_2)$ agree at $\overline{c}$. The result then follows by the sampling of $\overline{B}(\overline{c})/\overline{C}$ for all $\overline{c} \in \overline{C}$. We have

$$\begin{aligned}
\Pr_{(\overline{c}, \overline{c}'_1, \overline{c}'_2, \overline{b}_1, \overline{b}_2)} \left[\tilde{\beta}_1 \sim \tilde{\beta}_2\right] &\geq \mathbb{E}_{\overline{c}'_1, \overline{c}'_2 \sim \overline{C}'} \left[\Pr_{(\overline{c}, \overline{b}_1, \overline{b}_2)} \big[\exists\, \overline{a} \in \overline{A}(\overline{b}_1, \overline{b}_2) \text{ st } \tilde{\gamma}'_1 \sim \tilde{\alpha} \sim \tilde{\gamma}'_2 \;\&\; \tilde{\beta}_1 \sim \tilde{\alpha} \sim \tilde{\beta}_2\big] \right] \\
&\geq \mathbb{E}_{\overline{c}'_1, \overline{c}'_2 \sim \overline{C}'} \left[\Pr_{\substack{(\overline{c}, \overline{b}_1, \overline{b}_2) \\ \overline{a} \sim \overline{A}(\overline{b}_1, \overline{b}_2)}} \left[\tilde{\beta}_1 \sim \tilde{\alpha} \sim \tilde{\beta}_2 \middle| \tilde{\gamma}'_1 \sim \tilde{\alpha} \sim \tilde{\gamma}'_2 \right] \right].
\end{aligned}$$

Let $\mathsf{val} := \Pr_{(\overline{c}, \overline{b}_1, \overline{b}_2, \overline{a})} \big[\tilde{\beta}_1 \sim \tilde{\alpha} \sim \tilde{\beta}_2 \big| \tilde{\gamma}'_1 \sim \alpha \sim \tilde{\gamma}'_2\big]$ be shorthand for the quantity inside the expectation. We have

$$\mathsf{val} \;\geq\; 1 - \left[\Pr_{(\overline{c}, \overline{b}_1, \overline{b}_2, \overline{a})} \left[\tilde{\beta}_1 \not\sim \tilde{\alpha} \middle| \tilde{\gamma}'_1 \sim \tilde{\alpha} \sim \tilde{\gamma}'_2 \right] + \Pr_{(\overline{c}, \overline{b}_1, \overline{b}_2, \overline{a})} \left[\tilde{\beta}_2 \not\sim \tilde{\alpha} \middle| \tilde{\gamma}'_1 \sim \tilde{\alpha} \sim \tilde{\gamma}'_2 \right] \right]$$

$$\geq \; 1 - \frac{1}{\mathsf{q}(\overline{\mathsf{c}}'_1, \overline{\mathsf{c}}'_2)} \cdot \left[\Pr_{\substack{\overline{\mathsf{b}}_1 \sim \overline{\mathsf{B}}(\overline{\mathsf{c}}'_1) \\ \overline{\mathsf{a}} \sim \overline{\mathsf{A}}(\overline{\mathsf{b}}_1, \overline{\mathsf{c}}'_2)}} \left[\tilde{\beta}_1 \not\sim \tilde{\alpha} \middle| \tilde{\gamma}'_1 \sim \tilde{\alpha} \right] + \Pr_{\substack{\overline{\mathsf{b}}_2 \sim \overline{\mathsf{B}}(\overline{\mathsf{c}}'_2) \\ \overline{\mathsf{a}} \sim \overline{\mathsf{A}}(\overline{\mathsf{b}}_2, \overline{\mathsf{c}}'_1)}} \left[\tilde{\beta}_2 \not\sim \tilde{\alpha} \middle| \tilde{\gamma}'_2 \sim \tilde{\alpha} \right] \right]$$

By definition of $\overline{\mathsf{C}}'$, we have $\Pr_{\overline{\mathsf{c}}'_1, \overline{\mathsf{c}}'_2 \sim \overline{\mathsf{C}}'} \left[\mathsf{q}(\overline{\mathsf{c}}'_1, \overline{\mathsf{c}}'_2) < \varepsilon^5 \right] \leq \sigma$ and also

$$\begin{aligned} \mathbb{E}_{\overline{\mathsf{c}}'_1, \overline{\mathsf{c}}'_2 \sim \overline{\mathsf{C}}'} \left[\Pr_{\substack{\overline{\mathsf{b}}_1 \sim \overline{\mathsf{B}}(\overline{\mathsf{c}}'_1) \\ \overline{\mathsf{a}} \sim \overline{\mathsf{A}}(\overline{\mathsf{b}}_1, \overline{\mathsf{c}}'_2)}} \left[\tilde{\beta}_1 \not\sim \tilde{\alpha} \middle| \tilde{\gamma}'_1 \sim \tilde{\alpha} \right] \right] \; &\leq \; \max_{\overline{\mathsf{c}}'_1 \in \overline{\mathsf{C}}'} \left\{ \Pr_{\substack{\overline{\mathsf{b}}_1 \sim \overline{\mathsf{B}}(\overline{\mathsf{c}}'_1) \\ \overline{\mathsf{a}} \sim \overline{\mathsf{A}}(\overline{\mathsf{b}}_1)}} \left[\tilde{\beta}_1 \not\sim \tilde{\alpha} \middle| \tilde{\gamma}'_1 \sim \tilde{\alpha} \right] + \delta \right\} \\ &= \; \max_{\overline{\mathsf{c}}'_1 \in \overline{\mathsf{C}}'} \left\{ 1 - \mathsf{p}(\overline{\mathsf{c}}'_1) + \delta \right\} \leq \eta + \delta. \end{aligned}$$

We have used the sampling of $\overline{\mathsf{A}}(\overline{\mathsf{b}})/\overline{\mathsf{C}}$ for all $\overline{\mathsf{b}} \in \overline{\mathsf{B}}$, and that $\mathsf{p}(\overline{\mathsf{c}}'_1) \geq 1 - \eta$ since $\overline{\mathsf{c}}'_1 \in \overline{\mathsf{C}}'$. The result follows:

$$\mathbb{E}_{\overline{\mathsf{c}}'_1, \overline{\mathsf{c}}'_2 \sim \overline{\mathsf{C}}'} \left[\mathsf{val} \right] \geq \big(1 - \sigma\big) \cdot \big(1 - 2\varepsilon^{-5}(\eta + \delta)\big) \geq 1 - \big(\sigma + 2\varepsilon^{-5}(\eta + \delta)\big).$$

□

7 Affine Agreement

In this section we prove Lemma 2, restated in an expanded form below. We begin here by reducing Lemma 2 to Claims 7, 8 and 9, which we will prove in Section 7.2 after gathering some background on linearity/low-degree tests in Section 7.1. Recall that a function $\mathsf{T} : \Gamma \to \Gamma$ is *affine* if there exists $u \in \mathbb{F}$ and $\Phi_0 \in \Gamma$ such that $\mathsf{T}(\Phi) = u \cdot \Phi + \Phi_0$.

Lemma 2 (Restated). *Suppose* $\{\mathsf{f}_\mathsf{a}\}_\mathsf{a} \subset \{\mathsf{f} : \Gamma_\mathsf{A} \to \Gamma_\mathsf{A}\}$, $\mathsf{h} : \Gamma_\mathsf{C} \to \Gamma_\mathsf{C}$ *and* $\mathsf{G} \subset \mathsf{A} \times \Gamma$ *are such that* $|\mathsf{G}| \geq 2\varepsilon \cdot |\mathsf{A} \times \Gamma|$*, and*

$$\Pr_{\substack{(\mathsf{a}, \Phi) \sim \mathsf{G} \\ \overline{\mathsf{c}} \sim \mathsf{C}(\overline{\mathsf{a}})}} \left[\tilde{\gamma} \sim \tilde{\alpha} \right] \geq 1 - \zeta, \tag{7}$$

where (ε, ζ) *are as in Lemma 1. Then there exists an affine map* $\mathsf{T} : \Gamma \to \Gamma$ *such that*

$$\Pr_{(\mathsf{a}, \Phi) \sim \mathsf{G}} \left[\tilde{\alpha} = \mathsf{T}(\Phi)\big|_\mathsf{a} \right] \geq 1/2.$$

Claim 7. *Let* $\big(\varepsilon, \zeta, \{\mathsf{f}_\mathsf{a}\}, \mathsf{h}, \mathsf{G}\big)$ *be as in the hypothesis of Lemma 2, so that* (7) *holds. Then there exist affine maps* $\{\mathsf{T}_\mathsf{c}\}_{\mathsf{c} \in \mathsf{C}}$ *with* $\mathsf{T}_\mathsf{c} : \Gamma_\mathsf{C} \to \Gamma_\mathsf{C}$ *such that* $\Pr_{\overline{\mathsf{c}} \sim \overline{\mathsf{C}}} \left[\tilde{\gamma} = \mathsf{T}_\mathsf{c}(\gamma) \right] \geq 1 - \xi_7$ *holds, where* $\xi_7^2 := 32(d+1)(\zeta + \delta)$.

Claim 8. *Let* $\big(\varepsilon, \zeta, \{\mathsf{f}_\mathsf{a}\}, \mathsf{h}, \mathsf{G}\big)$ *be as in the hypothesis of Lemma 2, so that* (7) *holds, and let* $\{\mathsf{T}_\mathsf{c}\}$ *be the family of affine maps promised by Claim 7. For each* $\mathsf{c} \in \mathsf{C}$*, let* $u_\mathsf{c}, v_\mathsf{c} \in \mathbb{F}$ *be the scalars defining* T_c*, so* $\mathsf{T}_\mathsf{c}(\gamma) := u_\mathsf{c} \cdot \gamma + v_\mathsf{c}$*. Then there exists* $u \in \mathbb{F}$ *such that* $\Pr_{\mathsf{c} \sim \mathsf{C}} \left[u_\mathsf{c} = u \right] \geq 1 - \xi_8$*, where* $\xi_8 := (d+2)(\zeta + \delta) + 4\xi_7 + 2/|\mathbb{F}|$.

Claim 9. *Let* $\big(\varepsilon, \zeta, \{\mathsf{f}_\mathsf{a}\}, \mathsf{h}, \mathsf{G}\big)$ *be as in the hypothesis of Lemma 2, so that* (7) *holds, and let* $\{\mathsf{T}_\mathsf{c}\}$ *be the family of affine maps promised by Claim 7, with* $\mathsf{T}_\mathsf{c}(\gamma) := u_\mathsf{c} \cdot \gamma + v_\mathsf{c}$*, as in Claim 8. Then there exists* $\Phi_0 \in \Gamma$ *such that* $\Pr_{\mathsf{c} \sim \mathsf{C}} \left[v_\mathsf{c} = \Phi_0(\mathsf{c}) \right] \geq 1 - \xi_9$*, where* $\xi_9^2 := 8(d+3)^2(\zeta + \xi_7 + \xi_8)$.

Proof of Lemma 2 Assuming Claims 7, 8 and 9. Let $\left(\varepsilon, \zeta, \{\mathsf{f_a}\}, \mathsf{h}, \mathsf{G}\right)$ be as in the hypothesis of Lemma 2, so that (7) holds, and let $\{\mathsf{T_c}\}$ be the family of affine maps promised by Claim 7. Define the affine map $\mathsf{T} : \Gamma \to \Gamma$ by $\mathsf{T}(\Phi) := u \cdot \Phi + \Phi_0$, where $u \in \mathbb{F}$ and $\Phi_0 \in \Gamma$ are the quantities guaranteed by Claims 8 and 9, respectively. We have

$$\frac{3}{4} \leq \Pr_{\substack{(\mathsf{a},\Phi)\sim\mathsf{G} \\ \mathsf{c}\sim\mathsf{C}(\mathsf{a})}}\Big[\tilde{\gamma} \sim \tilde{\alpha} \;\&\; \tilde{\gamma} = \mathsf{T_c}(\gamma) \;\&\; u_{\mathsf{c}} = u \;\&\; v_{\mathsf{c}} = \Phi_0(\mathsf{c})\Big] \leq \Pr_{\substack{(\mathsf{a},\Phi)\sim\mathsf{G} \\ \mathsf{c}\sim\mathsf{C}(\mathsf{a})}}\Big[\tilde{\alpha}|_{\mathsf{c}} = \mathsf{T}(\Phi)|_{\mathsf{c}}\Big].$$

This follows from (7), Claims 7, 8, 9 and the sampling of $\mathsf{A} \times \Gamma/\overline{\mathsf{C}}$. We have used the loose bound $1/4 \leq (\zeta + \xi_7 + \xi_8 + \xi_9 + \delta)$ where $\zeta > 0$ (resp. ξ_7, ξ_8, ξ_9) are the quantities from the statement of Lemma 2 (resp. Claims 7, 8, and 9), and $\delta > 0$ is the sampling parameter. It follows that $\Pr_{(\mathsf{a},\Phi)\sim\mathsf{G}}\big[\tilde{\alpha} = \mathsf{T}(\Phi)|_{\mathsf{a}}\big] \geq 1/2$, since whenever $\tilde{\alpha}$ and $\mathsf{T}(\Phi)|_{\mathsf{a}}$ agree on half of the $\mathsf{c} \in \mathsf{C}(\mathsf{a})$, they must be equal as they are both low degree. The lemma follows. □

7.1 Linearity Testing Background

In this section we state three facts which we use in the next section to prove the claims. Throughout this section we use notations consistent with the rest of the paper. Additionally, in this section we use B as the set of lines in $\mathbb{F}^k$ and Γ_{B} is the set of univariate polynomials over $\mathbb{F}$ of degree at most d. Recall $\mathsf{T} : \Gamma_{\mathsf{C}} \to \Gamma_{\mathsf{C}}$ is *affine* if there exist coefficients $u, v \in \mathbb{F}$ such that $\mathsf{T}(x) = u \cdot x + v$ for all $x \in \Gamma_{\mathsf{C}}$. The first fact is standard and can be proved using linear algebraic methods.

Fact 4 (Linear Dependence of Polynomial Evaluations). *Suppose* $|\mathbb{F}| \geq d + 2$. *For any* $\mathsf{b} \in \mathsf{B}$ *and distinct* $\mathsf{c}_0, \ldots \mathsf{c}_{d+1} \in \mathsf{C}(\mathsf{b})$*, there exist non-zero coefficients* $r_0, r_1, \ldots, r_{d+1} \in \mathbb{F}$ *such that for all* $\beta \in \Gamma_{\mathsf{B}}$*,*

$$\sum_{i=0}^{d+1} r_i \cdot \beta|_{\mathsf{c}_i} = 0.$$

The second and third facts are proved in [RS96]. The second fact gives a sufficient condition for a function $f : \mathbb{F}^k \to \mathbb{F}$ being close to a multivariate low-degree polynomial.

Fact 5 (Robust Characterization of Low-Degree Functions). *Fix* $\kappa > 0$ *such that* $\kappa \leq \frac{1}{2(d+2)^2}$. *If* $f : \mathsf{C} \to \mathbb{F}$ *is such that*

$$\Pr_{\mathsf{b}\sim\mathsf{B}}\Big[\exists\, \beta \in \Gamma_{\mathsf{B}} \text{ st } \Pr_{\mathsf{c}\sim\mathsf{C}(\mathsf{b})}\big[f(\mathsf{c}) = \beta|_{\mathsf{c}}\big] \geq 1 - \kappa\Big] \geq 1 - \kappa,$$

then there exists $\Phi \in \Gamma$ *such that* $\Pr_{\mathsf{c}\sim\mathsf{C}}\big[f(\mathsf{c}) = \Phi(\mathsf{c})\big] \geq 1 - 2(d+3)\kappa$.

Fact 6 (Testing Affine Maps over Large Fields in High Soundness Regime). *Fix* $\kappa > 0$ *such that* $\kappa \leq \frac{1}{18}$. *If* $f : \Gamma_{\mathsf{C}} \to \Gamma_{\mathsf{C}}$ *is such that*

$$\Pr_{x,y,z\sim\Gamma_{\mathsf{C}}}\Big[f(x) + f(y+z) = f(x+y) + f(z)\Big] \geq 1 - \kappa,$$

then there exists an affine $\mathsf{T} : \Gamma_{\mathsf{C}} \to \Gamma_{\mathsf{C}}$ *such that* $\Pr_{x\sim\Gamma_{\mathsf{C}}}\big[f(x) = \mathsf{T}(x)\big] \geq 1 - 2\kappa$.

7.2 Proving the Claims

In this section we restate and prove the claims used to prove Lemma 2.

Notation. Throughout this section, we assume $\{\mathsf{f}_\mathsf{a}\}_\mathsf{a} \subset \{\mathsf{f} : \Gamma_\mathsf{A} \to \Gamma_\mathsf{A}\}$, $\mathsf{h} : \Gamma_\mathsf{C} \to \Gamma_\mathsf{C}$ and $\mathsf{G} \subset \mathsf{A} \times \Gamma$ with $|\mathsf{G}| \geq 2\varepsilon \cdot |\mathsf{A} \times \Gamma|$ are such that (7) holds. Namely, we assume that the hypotheses of Lemma 2. We also use $\tilde{\gamma} = \mathsf{h}(\overline{\mathsf{c}})$ throughout.

Claim 7 (Restated). *There exist affine maps* $\{\mathsf{T}_\mathsf{c}\}_{\mathsf{c}\in\mathsf{C}}$ *such that* $\Pr_{\overline{\mathsf{c}}\sim\overline{\mathsf{C}}}\big[\tilde{\gamma} = \mathsf{T}_\mathsf{c}(\gamma)\big] \geq 1 - \xi_7$.

Proof. Consider the following distribution, $\mathcal{D}$ on $\mathsf{C} \times \Gamma_\mathsf{C}^3$. Ultimately, the output of $\mathcal{D}$ is just uniform, however the internal choices of $\mathcal{D}$ help in our analysis. $\mathcal{D}$ works as follows:

1. draw $\mathsf{b} \sim \mathsf{B}$ and distinct $\mathsf{c}_0, \mathsf{c}_1, \ldots, \mathsf{c}_{d+1} \sim \mathsf{C}(\mathsf{b})$; let $r_0, \ldots, r_{d+1} \in \mathbb{F}$ be the coefficients guaranteed by Fact 4;
2. draw $\gamma_0^0, \gamma_0^1, \ldots, \gamma_d^0, \gamma_d^1 \sim \Gamma_\mathsf{C}$; let $\overline{\mathsf{c}}_{i,k} = (\mathsf{c}_k, \gamma_k^i)$, and $\tilde{\gamma}_k^i = \mathsf{h}(\overline{\mathsf{c}}_{i,k})$ for $i = 0, 1$ and $k = 0, \ldots, d$;
3. for $i, j \in \{0, 1\}$, let $\beta^{i,j} \in \Gamma_\mathsf{B}$ be the unique polynomial that agrees with γ_0^i at c_0 and γ_k^j at c_k for all $k = 1, \ldots, d$; let $\overline{\mathsf{b}}_{i,j} = (\mathsf{b}, \beta^{i,j})$;
4. for $i, j \in \{0, 1\}$, draw $(\mathsf{a}_{i,j}, \Phi^{i,j}) \sim \mathsf{G}(\overline{\mathsf{b}}_{i,j})$ and set $\tilde{\alpha}^{i,j} = \mathsf{f}_{\mathsf{a}_{i,j}}(\Phi^{i,j}|_{\mathsf{a}_{i,j}})$ and $\tilde{\beta}^{i,j} = \tilde{\alpha}^{i,j}|_\mathsf{b}$;
5. let $(\tilde{\gamma}, \tilde{\gamma}', \tilde{\gamma}'', \tilde{\gamma}''') = \big(\mathsf{h}(\mathsf{c}_{d+1}, \gamma), \mathsf{h}(\mathsf{c}_{d+1}, \gamma'), \mathsf{h}(\mathsf{c}_{d+1}, \gamma''), \mathsf{h}(\mathsf{c}_{d+1}, \gamma''')\big)$, where
$$(\gamma, \gamma', \gamma'', \gamma''') = \Big(\beta^{0,0}|_{\mathsf{c}_{d+1}}, \beta^{1,0}|_{\mathsf{c}_{d+1}}, \beta^{0,1}|_{\mathsf{c}_{d+1}}, \beta^{1,1}|_{\mathsf{c}_{d+1}}\Big);$$
here $\beta|_\mathsf{c}$ denotes the evaluation of the polynomial β at the point c;
6. output $(\mathsf{c}, x, y, z) = (\mathsf{c}_{d+1}, \gamma, \gamma' - \gamma, \gamma'')$.

Note that the output of $\mathcal{D}$ is uniform on $\mathsf{C} \times \Gamma_\mathsf{C}^3$. Indeed, c_{d+1} drawn in Step 1 is uniform since B/C is biregular. Moreover, given any fixed $\gamma_1^1, \ldots, \gamma_k^1$, γ'' varies uniformly as γ_0^0 does. Then, given any fixing of $(\gamma_0^0, \gamma_1^1, \ldots, \gamma_k^1)$, γ varies uniformly as $(\gamma_1^0, \ldots, \gamma_k^0)$ does. Finally, given any fixing of γ_0^0 and $(\gamma_0^0, \gamma_1^1, \ldots, \gamma_k^0, \gamma_k^1)$, γ' varies uniformly as γ_0^1 does.

Now, let $\mathbf{E}$ be the event: $\tilde{\gamma}_0^i \sim \tilde{\beta}^{i,j} \sim \tilde{\gamma}_k^j \; \forall\, (i, j, k) \in \{0, 1\}^2 \times \{1, \ldots, d\}$, where the $\tilde{\gamma}_0^i$, $\tilde{\beta}^{i,j}$, and $\tilde{\gamma}_k^j$ are the internal values drawn during steps 2 and 4. By the assumptions of Lemma 2 and the sampling of $\mathsf{A} \times \Gamma/\overline{\mathsf{B}}$, we have $\Pr_{\overline{\mathsf{b}},\overline{\mathsf{c}},(\mathsf{a},\Phi)}\big[\tilde{\gamma} \sim \tilde{\alpha}\big] \geq 1 - \zeta - \delta$, where the probability is over $\overline{\mathsf{b}} \sim \overline{\mathsf{B}}$, $\overline{\mathsf{c}} \sim \overline{\mathsf{C}}(\overline{\mathsf{b}})$, $(\mathsf{a}, \Phi) \sim \mathsf{G}(\overline{\mathsf{b}})$. It follows from the union bound that $\Pr_\mathcal{D}\big[\mathbf{E}\big] \geq 1 - \xi_7^2/8$ (substituting $\xi_7^2 = 32(d+1)(\zeta + \delta)$), since each $(\overline{\mathsf{b}}_{i,j}, \overline{\mathsf{c}}_{i,0}, \mathsf{a}_{i,j}, \Phi^{i,j})$ and $(\overline{\mathsf{b}}_{i,j}, \overline{\mathsf{c}}_{j,k}, \mathsf{a}_{i,j}, \Phi^{i,j})$ are, individually, drawn in this way for each $(i, j, k) \in \{0, 1\}^2 \times \{0, \ldots, d\}$.

We complete the proof by showing that whenever the sampling of $(\mathsf{c}, x, y, z) \sim \mathcal{D}$ is such that $\mathbf{E}$ occurs, it holds that $\mathsf{h}(\mathsf{c}, x) + \mathsf{h}(\mathsf{c}, y + z) = \mathsf{h}(\mathsf{c}, x + y) + \mathsf{h}(\mathsf{c}, z)$. Together with Fact 6, this implies that there is a family of affine maps $\{\mathsf{T}_\mathsf{c}\}_{\mathsf{c}\in\mathsf{C}}$ such that

$$\Pr_{\mathsf{c}\sim\mathsf{C}}\left[\Pr_{\gamma\sim\Gamma_\mathsf{c}}\big[\tilde{\gamma} = \mathsf{T}_\mathsf{c}(\gamma)\big] \geq 1 - \frac{\xi_7}{2}\right] \geq 1 - \frac{\xi_7}{2},$$

which implies the claim.

So it suffices to show that

$$\gamma - \gamma' = \gamma'' - \gamma''' \text{ and } \tilde{\gamma} - \tilde{\gamma}' = \tilde{\gamma}'' - \tilde{\gamma}'''$$

both hold whenever $\mathbf{E}$ occurs (the first equality always holds, the second holds whenever $\mathbf{E}$ occurs). This follows from Fact 4. The first equality holds since the $\beta^{i,j}$ are low-degree and for all (i,j,k), γ_0^i and γ_k^j are the evaluations of $\beta^{i,j}$ at c_0 and c_k, respectively. Thus Fact 4 gives

$$r_0 \cdot \gamma_0^0 + \left(\sum_{k=1}^d r_k \cdot \gamma_k^0\right) + r_{d+1} \cdot \gamma = 0; \quad r_0 \cdot \gamma_0^1 + \left(\sum_{k=1}^d r_k \cdot \gamma_k^0\right) + r_{d+1} \cdot \gamma' = 0;$$
$$r_0 \cdot \gamma_0^0 + \left(\sum_{k=1}^d r_k \cdot \gamma_k^1\right) + r_{d+1} \cdot \gamma'' = 0; \quad r_0 \cdot \gamma_0^1 + \left(\sum_{k=1}^d r_k \cdot \gamma_k^1\right) + r_{d+1} \cdot \gamma''' = 0,$$

which simplifies to $\gamma - \gamma' = \gamma'' - \gamma'''$ since $r_{d+1} \neq 0$. Likewise, for the second equality, the $\tilde{\beta}^{i,j}$ are low degree and when $\mathbf{E}$ occurs, the $\tilde{\gamma}_0^i$ and $\tilde{\gamma}_k^j$ are the evaluations of $\tilde{\beta}^{i,j}$ at c_0 and c_k. As above, this implies $\tilde{\gamma} - \tilde{\gamma}' = \tilde{\gamma}'' - \tilde{\gamma}'''$. □

Claim 8 (Restated). *Let $\{\mathsf{T}_\mathsf{c}\}$ be the family of affine maps promised by Claim 7; for each $\mathsf{c} \in \mathsf{C}$, let $\mathsf{T}_\mathsf{c}(\gamma) := u_\mathsf{c} \cdot \gamma + v_\mathsf{c}$ for $u_\mathsf{c}, v_\mathsf{c} \in \mathbb{F}$. Then there exists $u \in \mathbb{F}$ such that $\Pr_{\mathsf{c}\sim\mathsf{C}}\big[u_\mathsf{c} = u\big] \geq 1 - \xi_8$, where $\xi_8 = (d+2)(\zeta+\delta) + 4\xi_7 + 2/|\mathbb{F}|$.*

Proof. We prove that $\Pr_{\mathsf{c},\mathsf{c}'\sim\mathsf{C}}\big[u_\mathsf{c} = u_{\mathsf{c}'}\big] \geq 1 - \xi_8$ which suffices since

$$\Pr_{\mathsf{c},\mathsf{c}'\sim\mathsf{C}}\big[u_\mathsf{c} = u_{\mathsf{c}'}\big] = \sum_{u\in\mathbb{F}} \mathsf{p}_u^2 \leq \max\big\{\mathsf{p}_u : u \in \mathbb{F}\big\},$$

where $\mathsf{p}_u := \Pr_{\mathsf{c}\sim\mathsf{C}}\big[u_\mathsf{c} = u\big]$ is shorthand. As in the previous proof, we describe a distribution $\mathcal{D}'$ on C^2:

1. draw $\mathsf{b} \sim \mathsf{B}$ and distinct $\mathsf{c}_0, \mathsf{c}_1, \ldots, \mathsf{c}_{d+1} \sim \mathsf{C}(\mathsf{b})$; let $r_0, \ldots, r_{d+1} \in \mathbb{F}$ be the coefficients guaranteed by Fact 4; let $u_0, u_{d+1} \in \mathbb{F}$ denote the linear terms of $\mathsf{T}_{\mathsf{c}_0}$ and $\mathsf{T}_{\mathsf{c}_{d+1}}$, respectively;
2. draw $\gamma_0^0, \gamma_0^1, \gamma_k \sim \Gamma_\mathsf{C}$ for $k = 1, \ldots, d$; let $\overline{\mathsf{c}}_{i,0} = (\mathsf{c}_0, \gamma_0^i)$ for $i = 0, 1$ and $\overline{\mathsf{c}}_k = (\mathsf{c}_k, \gamma_k)$ for $k = 1, \ldots, d$; let $\tilde{\gamma}_0^i = \mathsf{h}(\overline{\mathsf{c}}_{i,0})$ and $\tilde{\gamma}_k = \mathsf{h}(\overline{\mathsf{c}}_k)$;
3. for $i \in \{0,1\}$, let $\beta^i \in \Gamma_\mathsf{B}$ be the unique polynomial that agrees with γ_0^i at c_0 and γ_k at c_k for all $k = 1, \ldots, d$; let $\overline{\mathsf{b}}_i = (\mathsf{b}, \beta^i)$;
4. for $i \in \{0,1\}$, draw $(\mathsf{a}_i, \Phi^i) \sim \mathsf{G}(\overline{\mathsf{b}}_i)$ and set $\tilde{\alpha}^i = \mathsf{f}_{\mathsf{a}_i}(\Phi^i|_{\mathsf{a}_i})$ and $\tilde{\beta}^i = \tilde{\alpha}^i|_\mathsf{b}$;
5. let $(\tilde{\gamma}, \tilde{\gamma}') = \big(\mathsf{h}(\mathsf{c}_{d+1}, \gamma), \mathsf{h}(\mathsf{c}_{d+1}, \gamma')\big)$, where $(\gamma, \gamma') = \big(\beta^0|_{\mathsf{c}_{d+1}}, \beta^1|_{\mathsf{c}_{d+1}}\big)$;
6. output $(\mathsf{c}, \mathsf{c}') = (\mathsf{c}_0, \mathsf{c}_{d+1})$.

Note that $\mathcal{D}'$ outputs two random points on a random line, which is within statistical distance $2/|\mathbb{F}|$ of uniform on C^2. Let $\mathbf{E}'$ be the event:

1. $\tilde{\gamma}_0^i \sim \tilde{\beta}^i \sim \tilde{\gamma}_k \; \forall \; (i,k) \in \{0,1\} \times \{1,\ldots,d\}$; and
2. $(\tilde{\gamma}_0^0, \tilde{\gamma}_0^1, \tilde{\gamma}, \tilde{\gamma}') = \big(\mathsf{T}_{\mathsf{c}_0}(\gamma_0^0), \mathsf{T}_{\mathsf{c}_0}(\gamma_0^1), \mathsf{T}_{\mathsf{c}_{d+1}}(\gamma), \mathsf{T}_{\mathsf{c}_{d+1}}(\gamma')\big)$

The first condition occurs with probability at least $1-(d+2)(\zeta+\delta)$; as in the proof of Claim 7, this follows from (7), the sampling of $\mathsf{A} \times \Gamma/\overline{\mathsf{B}}$, and a union bound. The second condition occurs with probability at least $1-4\xi_7$, by Claim 7. Upon substituting $\xi_8 = (d+2)(\zeta+\delta)+4\xi_7+2/|\mathbb{F}|$, we get $\Pr_{(\mathsf{c},\mathsf{c}')\sim\mathsf{C}^2}\big[\mathbf{E}'\big] \geq \Pr_{\mathcal{D}'}\big[\mathbf{E}'\big] - 2/|\mathbb{F}| \geq 1-\xi_8$. As in the proof of Claim 7, Fact 4 gives

$$r_0 \cdot (\gamma_0^0 - \gamma_0^1) + r_{d+1} \cdot (\gamma - \gamma') = 0 = r_0 \cdot (\tilde{\gamma}_0^0 - \tilde{\gamma}_0^1) + r_{d+1} \cdot (\tilde{\gamma} - \tilde{\gamma}').$$

Substituting $(\tilde{\gamma}_0^0 - \tilde{\gamma}_0^1) = u_0 \cdot (\gamma_0^0 - \gamma_0^1)$ and $(\tilde{\gamma} - \tilde{\gamma}') = u_{d+1} \cdot (\gamma - \gamma')$ gives $r_{d+1}(u_{d+1} - u_0)(\gamma - \gamma') = 0$ which means $u_{d+1} = u_0$ since $r_{d+1} \neq 0$ and $\gamma \neq \gamma'$. Thus, $\Pr_{\mathsf{c},\mathsf{c}'\sim\mathsf{C}}\big[u_{\mathsf{c}} = u_{\mathsf{c}'}\big] \geq 1-\xi_8$. □

Claim 9 (Restated). *Let $\{\mathsf{T}_{\mathsf{c}}\}$ be the family of affine maps promised by Claim 7. Then there exists $\Phi_0 \in \Gamma$ with $\Pr_{\mathsf{c}\sim\mathsf{C}}\big[\mathsf{T}_{\mathsf{c}}(\mathbf{0}) = \Phi_0(\mathsf{c})\big] \geq 1-\xi_9$, where $\xi_9^2 = 8(d+3)^2(\zeta+\xi_7+\xi_8)$.*

Proof. Let $v : \mathsf{C} \to \mathbb{F}$ as a function mapping $\mathsf{c} \mapsto v_{\mathsf{c}} = \mathsf{T}_{\mathsf{c}}(\mathbf{0})$. Let $\xi := \frac{\xi_9}{2(d+3)}$. We will show that

$$\Pr_{\mathsf{b}\sim\mathsf{B}}\Big[\exists\, \tilde{\beta}' \in \Gamma_{\mathsf{B}} \text{ st } \Pr_{\mathsf{c}\sim\mathsf{C}(\mathsf{b})}\big[v_{\mathsf{c}} = \tilde{\beta}'|_{\mathsf{c}}\big] \geq 1-\xi\Big] \geq 1-\xi. \tag{8}$$

The claim then follows from Fact 5. Towards establishing (8), note that

$$\Pr_{\substack{(\mathsf{a},\Phi)\sim\mathsf{G}\\ \mathsf{b}\sim\mathsf{B}(\mathsf{a})\\ \mathsf{c}\sim\mathsf{C}(\mathsf{b})}}\big[v_{\mathsf{c}} = \tilde{\beta}|_{\mathsf{c}} - u\cdot\beta|_{\mathsf{c}}\big] \geq 1-(\zeta+\xi_7+\xi_8) \geq 1-\xi(\xi-\delta),$$

where $\beta = \Phi|_{\mathsf{b}}$ and $\tilde{\beta} = \tilde{\alpha}|_{\mathsf{b}}$; we have used $\xi(\xi-\delta) \geq \xi^2/2 = \zeta+\xi_7+\xi_8$. This follows immediately from (7) and Claims 7 and 8. By an averaging argument,

$$\Pr_{\substack{(\mathsf{a},\Phi)\sim\mathsf{G}\\ \mathsf{b}\sim\mathsf{B}(\mathsf{a})}}\Big[\Pr_{\mathsf{c}\sim\mathsf{C}(\mathsf{b})}\big[v_{\mathsf{c}} = \tilde{\beta}'|_{\mathsf{c}}\big] \geq 1-\xi\Big] \geq 1-\xi+\delta,$$

where $\tilde{\beta}' = \tilde{\beta} - u\cdot\beta$. The bound (8) now follows from the sampling of $\mathsf{A} \times \Gamma/\mathsf{B}$. □

8 A Locally Testable, Non-Malleable Code

In this section, we give a construction of a locally testable non-malleable code against coordinate wise tampering. To build our code, we take the LTNM reduction, $(\mathsf{E}_{\mathsf{LTNM}}, \mathsf{D}_{\mathsf{LTNM}}, \mathsf{T}_{\mathsf{LTNM}})$, from coordinate-wise tampering to affine tampering, from section 3.4 and compose it with a new non-malleable code, $(\mathsf{E}_{\mathsf{aff}}, \mathsf{D}_{\mathsf{aff}})$, against affine tampering.

8.1 A Simple Non-malleable Code against Affine Tampering

We begin with a new constant rate, non-malleable code against affine tampering. This result is not new, several prior works [ADL14, CZ14, Li16, CL17] give such codes, however, our construction is considerably simpler than those prior.

Notations. Let $\mathbb{F}$ be a finite field and $\mathbb{K}/\mathbb{F}$ a degree 3 extension, so $\mathbb{K} = \mathbb{F}[x]/\big(p(x)\big)$ for an irreducible cubic polynomial $p(x) = x^3 - e_2x^2 - e_1x - e_0$. Thus $\mathbb{K}$ is a 3-dimensional $\mathbb{F}-$vector space with basis $\{1, \sigma, \sigma^2\}$, where $\sigma \in \mathbb{K}$ is a root of $p(x)$. The 'multiplication by σ' map $\mathbb{F}^3 \to \mathbb{F}^3$ is linear, specified over this basis by the matrix

$$\Sigma = \begin{bmatrix} 0 & 0 & e_0 \\ 1 & 0 & e_1 \\ 0 & 1 & e_2 \end{bmatrix} \in \mathbb{F}^{3\times 3}.$$

Our code makes use of an $\varepsilon-$*high entropy encoding*, (E, D), with codeword space $\mathbb{F}$, such that for all m, c^*, $\Pr_{c\sim\mathsf{E}(m)}\big[c = c^*\big] \leq \varepsilon$. Such codes can be trivilally constructed by appending a message with a random string of length $\log\big(1/\varepsilon\big)$.

Construction. Let (E, D) be an $\varepsilon-$high entropy code with message space $\mathcal{M}$ and codeword space $\mathbb{F}$, and let $m \in \mathcal{M}$.

• $\underline{\mathsf{E}_{\mathsf{aff}}(m)}$**:** Draw $r \sim \mathbb{F}; w \sim \mathsf{E}(m)$ and output $w + r \cdot \sigma + wr \cdot \sigma^2 \in \mathbb{K}$.

• $\underline{\mathsf{D}_{\mathsf{aff}}(c)}$**:** Parse $c = c_0 + c_1 \cdot \sigma + c_2 \cdot \sigma^2$; if $c_0 \cdot c_1 = c_2$, output $m = \mathsf{D}(c_0)$; if not, output $\perp$.

Theorem 4. *Fix $\varepsilon > 0$, and let (E, D) be an $\varepsilon-$high entropy code with message space $\mathcal{M}$ and codeword space $\mathbb{F}$. Then $(\mathsf{E}_{\mathsf{aff}}, \mathsf{D}_{\mathsf{aff}})$ is a $(2\varepsilon+2/|\mathbb{F}|)-$non-malleable code against affine tampering functions.*

Proof. Fix an affine map f given by $f(x) = sx+t$ where $s, t, x \in \mathbb{K}$ and fix any message $m \in \mathcal{M}$. Parse $s = s_0+s_1\cdot\sigma+s_2\cdot\sigma^2$ and $t = t_0+t_1\cdot\sigma+t_2\cdot\sigma^2$. To prove the theorem, we exhibbit a trivial tampering function g_f (*i.e.*, either constant or the identity) such that the tampering distribution $\big(\mathsf{D}_{\mathsf{aff}} \circ f \circ \mathsf{E}_{\mathsf{aff}}\big)(m)$ outputs $g_f(m)$ with probability at least $1 - 2\varepsilon - 2/|\mathbb{F}|$. The trivial function g_f is f if f is either the identity or a constant function mapping to a valid codeword, and is the constant $\perp$ function otherwise. Specifically, if $(s, t) = (1, 0)$, g_f is the identity; if $s = 0$ and $t_0 \cdot t_1 = t_2$, g_f is the constant function mapping everything to t; otherwise g_f is the constant $\perp$ function. The key point, is that for all $m \in \mathcal{M}$, the distribution $f\big(\mathsf{E}_{\mathsf{aff}}(m)\big)$ draws $w \sim \mathsf{E}(m)$, $r \sim \mathbb{F}$ and outputs

$$S\begin{bmatrix} w \\ r \\ wr \end{bmatrix} + \begin{bmatrix} t_0 \\ t_1 \\ t_2 \end{bmatrix} = \begin{bmatrix} t_0 + s_0w + e_0s_2r + (e_0s_1 + e_0e_2s_2)wr \\ t_1 + s_1w + (s_0 + e_1s_2)r + (e_1s_1 + s_2e_0 + s_2e_1e_2)wr \\ t_2 + s_2w + (s_1 + e_2s_2)r + (s_0 + e_2s_1 + s_2e_2^2 + s_2e_1)wr \end{bmatrix} =: \begin{bmatrix} \mathsf{C}_0(w, r) \\ \mathsf{C}_1(w, r) \\ \mathsf{C}_2(w, r) \end{bmatrix},$$

where $S \in \mathbb{F}^{3\times 3}$ is the 'multiplication by s' matrix: $S = s_0 \cdot \mathbb{1} + s_1 \cdot \Sigma + s_2 \cdot \Sigma^2$. In the above, we have defined bilinear (*i.e.*, of the form $a + bx + cy + dxy$) polynomials $\mathsf{C}_0, \mathsf{C}_1, \mathsf{C}_2 \in \mathbb{F}[x, y]$. Note that if $\mathsf{C}_0(x, y)\cdot\mathsf{C}_1(x, y) \not\equiv \mathsf{C}_2(x, y)$ as polynomials, then $\mathsf{C}_0(w, r)\cdot\mathsf{C}_1(w, r) = \mathsf{C}_2(w, r)$ holds with probability at most $2\varepsilon + 2/|\mathbb{F}|$, in which case $\big(\mathsf{D}_{\mathsf{aff}} \circ f \circ \mathsf{E}_{\mathsf{aff}}\big)(m) = \perp$ with high probability. This follows immediately from Schwartz-Zippel and the low entropy property of (E, D). Therefore, in order to prove the theorem, it suffices to show that if $\mathsf{C}_0(x, y) \cdot \mathsf{C}_1(x, y) \equiv \mathsf{C}_2(x, y)$ holds, then either $s = 0$ or $(s, t) = (1, 0)$. We assume $\mathsf{C}_0(x, y) \cdot \mathsf{C}_1(x, y) \equiv \mathsf{C}_2(x, y)$ holds, and we prove the following three items:

1. either $s_1 = 0$ or $s_2 = 0$;
2. $s_1 = 0 \Leftrightarrow s_2 = 0$;
3. if $s_1 = s_2 = 0$ then either $s_0 = 0$ or $s_0 = 1$ and $t_0 = t_1 = t_2 = 0$.

The third point is easiest: if $\mathsf{C}_0(x,y) \cdot \mathsf{C}_1(x,y) \equiv \mathsf{C}_2(x,y)$ and $s_1 = s_2 = 0$ then plugging gives

$$(t_0 + s_0 x) \cdot (t_1 + s_0 y) = t_2 + s_0 xy,$$

from which it follows that either $s_0 = 0$ or $s_0 = 1$ and $t_i = 0$ for all $i = 0, 1, 2$. To prove the first point, note that if $\mathsf{C}_0(x,y) \cdot \mathsf{C}_1(x,y) \equiv \mathsf{C}_2(x,y)$, then $s_0 \cdot s_1 = 0$ (since the x^2 coefficient in C_2 is zero). If $s_1 = 0$ we are done; if $s_0 = 0$ then $e_0 e_1 s_2^2 = 0$ (since y^2 coefficient in C_2 is zero), which implies $e_1 s_2 = 0$ since $e_0 \neq 0$ (else $p(x)$ is reducible). If $s_2 = 0$ we are done; if $e_1 = 0$ then $e_0^2 s_2^2 = 0$ (since xy^2 coefficient in C_2 is zero). Again, $e_0 \neq 0$ so $s_2 = 0$ so the first point follows.

Finally, for the second point, assume $s_1 = 0$. Then $s_0 s_2 \cdot (e_0 + e_1 e_2) = 0$ since the coefficient of $x^2 y = 0$ in C_2. Note $e_0 \neq -e_1 e_2$ since otherwise $p(x)$ is reducible: $p(x) = (x - e_2)(x^2 - e_1)$. However, if $s_0 = 0$ then, as shown in the proof of the first point, $s_2 = 0$; therefore $s_1 = 0$ implies $s_2 = 0$. Conversely, if $s_2 = 0$ then $e_0 s_0 s_1 = 0$ (coefficient of xy^2 in C_2 is zero), so $s_0 s_1 = 0$. If $s_0 = 0$ then $e_0 s_1^2 = 0$ (coefficient of $x^2 y$ in C_2 is zero). Thus $s_2 = 0$ implies $s_1 = 0$, and we are done. □

Remark. In our LTNM code in the next section, we will use $(\mathsf{E}_{\mathsf{aff}}, \mathsf{D}_{\mathsf{aff}})$ to encode a random $w \in \mathbb{F}$ and so the high entropy encoding is not necessary. The precise claim we use is stated below. The proof is the same as above since if $\mathsf{C}_0(x,y) \cdot \mathsf{C}_1(x,y) \not\equiv \mathsf{C}_2(x,y)$ as polynomials, then $\mathsf{C}_0(w,r) \cdot \mathsf{C}_1(w,r) = \mathsf{C}_2(w,r)$ holds with probability at most $4/|\mathbb{F}|$ over $w, r \sim \mathbb{F}$.

Claim 10. *Let $f : \mathbb{K} \to \mathbb{K}$ be affine of the form $f(x) = sx + t$ for $s, t \in \mathbb{K}$ such that $s \neq 0$ and $(s,t) \neq (1,0)$. Then* $\Pr_{w,r \sim \mathbb{F}}\Big[\mathsf{D}_{\mathsf{aff}}\big(f(w + r \cdot \sigma + wr \cdot \sigma^2)\big) \neq \perp\Big] \leq 4/|\mathbb{F}|$.

8.2 A LTNM Code via Composition

Composition Overview. The local test of our main construction from Section 3.4 passes whenever codewords are tampered by a coordinate-wise affine function. Thus, in order to use our main construction to build a fully LTNM code, we must modify the test in such a way so that it fails whenever a non-trivial affine tampering function is used. We do this in two steps. First, we modify the local tester so that it locally decodes a specified polynomial evaluation. Second, the tester checks that the evaluation recovered is a valid codeword of $(\mathsf{E}_{\mathsf{aff}}, \mathsf{D}_{\mathsf{aff}})$, if not it outputs $\perp$. Essentially, the reason this works is that the local decoder will output $\perp$ unless the codeword is tampered with an affine function, in which case the evaluation recovered is an affine function of the original evaluation. If the original evaluation is a random valid codeword of $(\mathsf{E}_{\mathsf{aff}}, \mathsf{D}_{\mathsf{aff}})$ then by Claim 10, the recovered evaluation is a valid codeword only if the affine tampering function is trivial.

Notations. As in the previous section, let $\mathbb{K}/\mathbb{F}$ be a degree 3 extension with $\mathbb{F}-$basis $\{1, \sigma, \sigma^2\}$. Let $k \geq 5$ and $d \geq 2$. As in the rest of the paper, let A be the set of $3-$planes in $\mathbb{F}^k$ and $\mathsf{C} = \mathbb{F}^k$. In this section, we use B and $\bar{\mathsf{A}}$ to denote the set of lines and $4-$planes respectively(note, the second usage is different from rest of the paper where we used $\bar{\mathsf{A}}$ to denote $\mathsf{A} \times \Gamma_{\mathsf{A}}$). Let $\mathsf{p} = (1, 0, \ldots, 0) \in \mathbb{F}^k$.

Construction. Let $\mathsf{E}_{\mathsf{aff}}()$ denote the procedure which draws $w, r \sim \mathbb{F}$, and outputs the value $w + r \cdot \sigma + wr \cdot \sigma^2 \in \mathbb{K}$; let $\mathsf{D}_{\mathsf{aff}}$ be the decoding algorithm from the previous section. Let $m \in \mathbb{K}$ be a message.

- $\underline{\mathsf{Enc}(m)}$**:** Draw $v \sim \mathsf{E}_{\mathsf{aff}}()$; and $\Phi \sim \Gamma$ such that $\Phi(\mathbf{0}) = m$ and $\Phi(\mathsf{p}) = v$; output $\big\{(\mathsf{a}, \Phi_{\mathsf{a}})\big\}_{\mathsf{a} \in \mathsf{A}}$.
- $\underline{\mathsf{Dec}\big(\{(\mathsf{a}, \alpha)\}_{\mathsf{a} \in \mathsf{A}}\big)}$**:** Find $\Phi \in \Gamma$ such that $(\mathsf{a}, \alpha) = (\mathsf{a}, \Phi|_{\mathsf{a}})$ for all $\mathsf{a} \in \mathsf{A}$. If such Φ exists, and if $\mathsf{D}_{\mathsf{aff}}\big(\Phi(\mathsf{p})\big) \neq \perp$, output $m = \Phi(\mathbf{0})$, otherwise output $\perp$.
- $\underline{\mathsf{Test}\big(\{(\mathsf{a}, \alpha)\}_{\mathsf{a} \in \mathsf{A}}\big)}$**:** Draw $\mathsf{b} \sim \mathsf{B}(\mathsf{p})$, $\mathsf{c}_1, \mathsf{c}_2, \mathsf{c}_3 \sim \mathsf{C}(\mathsf{b})$, $\mathsf{c}, \mathsf{c}' \sim \mathsf{C}$, $\mathsf{a}_1 \sim \mathsf{A}(\mathsf{c}, \mathsf{c}_1)$, $\mathsf{a}_2 \sim \mathsf{A}(\mathsf{c}, \mathsf{c}', \mathsf{c}_2)$, $\mathsf{a}_3 \sim \mathsf{A}(\mathsf{c}', \mathsf{c}_3)$. Read (a_1, α_1), (a_2, α_2), (a_3, α_3) and do the following.
  1) Check that $\alpha_1|_{\mathsf{c}} = \alpha_2|_{\mathsf{c}}$ and $\alpha_2|_{\mathsf{c}'} = \alpha_3|_{\mathsf{c}'}$; if not output 0; if so use interpolation to recover $\beta \in \Gamma_{\mathsf{B}}$, the unique degree 2 polynomial such that $\beta|_{\mathsf{c}_i} = \alpha_i|_{\mathsf{c}_i}$ for $i = 1, 2, 3$; let $v = \beta|_{\mathsf{p}}$.
  2) If $\mathsf{D}_{\mathsf{aff}}(v) \neq \perp$, output 1; otherwise output 0.

Theorem 5. *Let ℓ, ε as in theorem 1. Then the code $(\mathsf{Enc}, \mathsf{Dec}, \mathsf{Test})$ above is a $(\ell, \varepsilon')-$locally testable, non-malleable code against $\mathcal{F}$, the family of coordinate-wise tampering functions where $\varepsilon' = \mathcal{O}(\varepsilon^{1/2})$.*

Proof. Fix a tampering function $\mathsf{f} = \{\mathsf{f}_{\mathsf{a}}\}_{\mathsf{a}} \in \mathcal{F}$. We prove that $(\mathsf{Enc}, \mathsf{Dec}, \mathsf{Test})$ is LTNMC using the sufficient conditions of Claim 1. The first condition is trivial. For all distinct $\{\mathsf{h}_{\mathsf{a}}\}_{\mathsf{a}}, \{\mathsf{h}'_{\mathsf{a}}\}_{\mathsf{a}} \in \mathcal{H}$, $\Pr_{\Phi, \mathsf{a}}\big[\mathsf{h}_{\mathsf{a}}(\Phi|_{\mathsf{a}}) = \mathsf{h}'_{\mathsf{a}}(\Phi|_{\mathsf{a}})\big] = \mathcal{O}\big(|\mathbb{F}|^{-1}\big)$, as before. Therefore, it remains to exhibbit a list $\mathsf{L}_{\mathsf{f}} \subset \mathcal{H}$ of size at most $|\mathsf{L}_{\mathsf{f}}| \leq \ell$ such that $\mathsf{val} \leq \mathcal{O}(\varepsilon^{1/2})$ where

$$\mathsf{val} := \Pr_{\Phi, \mathsf{rand}}\Big[\mathsf{Test} \text{ passes } \& \ (\tilde{\alpha}_1, \tilde{\alpha}_2, \tilde{\alpha}_3) \notin \big\{\big(\mathsf{h}_{\mathsf{a}_1}(\Phi|_{\mathsf{a}_1}), \mathsf{h}_{\mathsf{a}_2}(\Phi|_{\mathsf{a}_2}), \mathsf{h}_{\mathsf{a}_3}(\Phi|_{\mathsf{a}_3})\big) : \{\mathsf{h}_{\mathsf{a}}\}_{\mathsf{a}} \in \mathsf{L}_{\mathsf{f}}\big\}\Big],$$

where $\tilde{\alpha}_i = \mathsf{f}_{\mathsf{a}_i}(\Phi|_{\mathsf{a}_i})$. In the course of the proof of Theorem 1 from Section 4, a similar list $\mathsf{L}'_{\mathsf{f}} \subset \mathcal{G}$ of size at most $|\mathsf{L}'_{\mathsf{f}}| \leq \ell$ was constructed such that

$$\Pr_{\Phi, (\mathsf{c}, \mathsf{a}_1, \mathsf{a}_2)}\Big[\tilde{\alpha}_1|_{\mathsf{c}} = \tilde{\alpha}_2|_{\mathsf{c}} \ \& \ (\tilde{\alpha}_1, \tilde{\alpha}_2) \notin \big\{\big(\mathsf{g}_{\mathsf{a}_1}(\Phi|_{\mathsf{a}_1}), \mathsf{g}_{\mathsf{a}_2}(\Phi|_{\mathsf{a}_2})\big) : \{\mathsf{g}_{\mathsf{a}}\}_{\mathsf{a}} \in \mathsf{L}'_{\mathsf{f}}\big\}\Big] \leq \varepsilon,$$

where this probability is over $\Phi \sim \Gamma$ and $\mathsf{c} \sim \mathsf{C}$, $\mathsf{a}_1, \mathsf{a}_2 \sim \mathsf{A}(\mathsf{c})$. Our list $\mathsf{L}_{\mathsf{f}} \subset \mathcal{H}$ is the set of trivial (*i.e.*, constant or affine) $\{\mathsf{g}_{\mathsf{a}}\}_{\mathsf{a}} \in \mathsf{L}'_{\mathsf{f}}$. The quantity val can now be bounded

$$\mathsf{val} \leq \Pr_{\Phi, \mathsf{rand}}\big[\mathbf{E}_1 \vee \mathbf{E}'_1 \vee \mathbf{E}_2 \vee \mathbf{E}_3\big]$$

for the following events:

$\mathbf{E}_1$**:** $\tilde{\alpha}_1|_{\mathsf{c}} = \tilde{\alpha}_2|_{\mathsf{c}}$ & $(\tilde{\alpha}_1, \tilde{\alpha}_2) \notin \big\{\big(\mathsf{g}_{\mathsf{a}_1}(\Phi|_{\mathsf{a}_1}), \mathsf{g}_{\mathsf{a}_2}(\Phi|_{\mathsf{a}_2})\big) : \{\mathsf{g}_{\mathsf{a}}\}_{\mathsf{a}} \in \mathsf{L}'_{\mathsf{f}}\big\}$;

$\mathbf{E}'_1$**:** $\tilde{\alpha}_2|_{\mathsf{c}'} = \tilde{\alpha}_3|_{\mathsf{c}'}$ & $(\tilde{\alpha}_2, \tilde{\alpha}_3) \notin \big\{\big(\mathsf{g}'_{\mathsf{a}_2}(\Phi|_{\mathsf{a}_2}), \mathsf{g}'_{\mathsf{a}_3}(\Phi|_{\mathsf{a}_3})\big) : \{\mathsf{g}'_{\mathsf{a}}\}_{\mathsf{a}} \in \mathsf{L}'_{\mathsf{f}}\big\}$;

$\mathbf{E}_2$**:** the $\{\mathsf{g}_{\mathsf{a}}\}_{\mathsf{a}}, \{\mathsf{g}'_{\mathsf{a}}\}_{\mathsf{a}} \in \mathcal{G}$ which agree with f from $\mathbf{E}_1$ and $\mathbf{E}'_1$ are distinct and such that $\mathsf{g}_{\mathsf{a}_2}(\Phi|_{\mathsf{a}_2}) = \mathsf{g}'_{\mathsf{a}_2}(\Phi|_{\mathsf{a}_2})$;

$\mathbf{E}_3$**:** the same $\{\mathsf{g}_{\mathsf{a}}\}_{\mathsf{a}} \in \mathcal{G}$ results from $\mathbf{E}_1$ and $\mathbf{E}'_1$; this $\{\mathsf{g}_{\mathsf{a}}\}_{\mathsf{a}} \in \mathcal{G}$ is non-trivial, but the affine check passes: $\mathsf{D}_{\mathsf{aff}}(\tilde{v}) \neq \perp$.

The marginal distribution on a_2 from rand is uniform, so $\Pr_{\Phi,\mathsf{rand}}\big[\mathbf{E}_2\big] = \mathcal{O}\big(|\mathbb{F}|^{-1}\big)$. By Claim 10, $\Pr_{\Phi,\mathsf{rand}}\big[\mathbf{E}_3\big] \leq 4/|\mathbb{F}|$. We prove $\Pr_{\Phi,\mathsf{rand}}\big[\mathbf{E}_1\big] \leq \varepsilon^{1/2} + \mathcal{O}\big(|\mathbb{F}|^{-1}\big)$. The same holds for $\mathbf{E}'_1$, and the result follows. Towards bounding $\Pr_{\Phi,\mathsf{rand}}\big[\mathbf{E}_1\big]$, note that drawing $\Phi \sim \Gamma$ uniformly, rather than uniformly subject to $\Phi(\mathbf{0}) = m$ and $\Phi(\mathsf{p}) = v$ changes the probability by at most $\mathcal{O}\big(|\mathbb{F}|^{-1}\big)$. Therefore, in the calculation below, we assume $\Phi \sim \Gamma$. We have

$$\begin{aligned}\Pr_{\Phi,\mathsf{rand}} \quad \big[\mathbf{E}_1\big]^2 &= \mathbb{E}_{\Phi,\mathsf{c}\sim\mathsf{C},\mathsf{a}_2\sim\mathsf{A}(\mathsf{c})}\Big[\Pr_{\mathsf{a}_1\sim\mathsf{rand}(\mathsf{c},\mathsf{a}_2)}\big[\mathbf{E}_1\big]\Big]^2 \leq \mathbb{E}_{\Phi,\mathsf{c},\mathsf{a}_2}\Big[\Pr_{\mathsf{a}_1\sim\mathsf{rand}(\mathsf{c},\mathsf{a}_2)}\big[\mathbf{E}_1\big]^2\Big] \\ \leq \quad & \mathbb{E}_{\Phi,\mathsf{c},\mathsf{a}_2}\Big[\Pr_{\mathsf{a}_1,\mathsf{a}_3\sim\mathsf{rand}(\mathsf{c},\mathsf{a}_2)}\big[\tilde{\alpha}_1|_{\mathsf{c}} = \tilde{\alpha}_2|_{\mathsf{c}} = \tilde{\alpha}_3|_{\mathsf{c}} \;\&\; (\tilde{\alpha}_1, \tilde{\alpha}_2, \tilde{\alpha}_3) \notin \mathsf{L}'_{\mathsf{f}}\big]\Big] + \mathcal{O}(|\mathbb{F}|^{-1}),\end{aligned}$$

where "$(\tilde{\alpha}_1, \tilde{\alpha}_2, \tilde{\alpha}_3) \notin \mathsf{L}'_{\mathsf{f}}$" is shorthand for

$$(\tilde{\alpha}_1, \tilde{\alpha}_2, \tilde{\alpha}_3) \notin \big\{\big(\mathsf{g}_{\mathsf{a}_1}(\Phi|_{\mathsf{a}_1}), \mathsf{g}_{\mathsf{a}_2}(\Phi|_{\mathsf{a}_2}), \mathsf{g}_{\mathsf{a}_3}(\Phi|_{\mathsf{a}_3})\big) : \{\mathsf{g}_{\mathsf{a}}\}_{\mathsf{a}} \in \mathsf{L}'_{\mathsf{f}}\big\}$$

and the $\mathcal{O}(|\mathbb{F}|^{-1})$ term in the second line accounts for the case when there are $\{\mathsf{g}_{\mathsf{a}}\}_{\mathsf{a}}, \{\mathsf{g}'_{\mathsf{a}}\}_{\mathsf{a}} \in \mathsf{L}'_f$ such that $\mathsf{g}_{\mathsf{a}_2}(\Phi|_{\mathsf{a}_2}) = \mathsf{g}'_{\mathsf{a}_2}(\Phi|_{\mathsf{a}_2})$ holds. Note that if $\tilde{\alpha}_1 = \mathsf{g}_{\mathsf{a}_1}(\Phi|_{\mathsf{a}_1})$, and $\tilde{\alpha}_2 \neq \mathsf{g}_{\mathsf{a}_2}(\Phi|_{\mathsf{a}_2})$, then $\tilde{\alpha}_1|_{\mathsf{c}} = \tilde{\alpha}_2|_{\mathsf{c}}$ occurs with probability $\mathcal{O}\big(|\mathbb{F}|^{-1}\big)$. It follows that

$$\Pr_{\Phi,\mathsf{rand}}\big[\mathbf{E}_1\big]^2 \leq \Pr_{\substack{\Phi,\mathsf{c},\mathsf{a}_2 \\ \mathsf{a}_1,\mathsf{a}_3\sim\mathsf{rand}(\mathsf{c},\mathsf{a}_2)}}\Big[\tilde{\alpha}_1|_{\mathsf{c}} = \tilde{\alpha}_3|_{\mathsf{c}} \;\&\; (\tilde{\alpha}_1, \tilde{\alpha}_3) \notin \mathsf{L}'_{\mathsf{f}}\Big] + \mathcal{O}\big(|\mathbb{F}|^{-1}\big).$$

Therefore, it suffices to show that for all $\mathsf{c} \in \mathsf{C}$, the distribution which draws $\mathsf{a}_2 \sim \mathsf{A}(\mathsf{c})$, $\mathsf{a}_1, \mathsf{a}_3 \sim \mathsf{rand}(\mathsf{c}, \mathsf{a}_2)$ and outputs $(\mathsf{a}_1, \mathsf{a}_3)$ is within statistical distance $\mathcal{O}\big(|\mathbb{F}|^{-1}\big)$ of uniform on $\mathsf{A}(\mathsf{c})^2$. The distribution $\mathsf{rand}(\mathsf{c}, \mathsf{a}_2)$ draws $\mathsf{c}_2 \sim \mathsf{C}(\mathsf{a}_2)$, $\mathsf{c}_1 \sim \mathsf{C}(\mathsf{b})$, where b is the line through p and c_2, and outputs $\mathsf{a}_1 \sim \mathsf{A}(\mathsf{c}, \mathsf{c}_1)$. This is equivalent to drawing $\mathsf{c}_1 \sim \mathsf{C}(\bar{\mathsf{a}}_2)$ and outputting $\mathsf{a}_1 \sim \mathsf{A}(\mathsf{c}, \mathsf{c}_1)$, where $\bar{\mathsf{a}}_2$ is the $4-$plane containing a_2 and p. Thus the distribution which draws $\mathsf{a}_2 \sim \mathsf{A}(\mathsf{c})$ and then $\mathsf{a}_1, \mathsf{a}_3 \sim \mathsf{rand}(\mathsf{c}, \mathsf{a}_2)$, outputting $(\mathsf{a}_1, \mathsf{a}_2, \mathsf{a}_3)$ can be equivalently described by drawing $\mathsf{a}_1, \mathsf{a}_3 \sim \mathsf{A}(\mathsf{c})$, $\mathsf{c}_i \sim \mathsf{C}(\mathsf{a}_i)$ for $i = 1, 3$, $\bar{\mathsf{a}}_2 \sim \overline{\mathsf{A}}(\mathsf{c}, \mathsf{p}, \mathsf{c}_1, \mathsf{c}_3)$ (*i.e.*, a random $4-$plane containing $\mathsf{c}, \mathsf{p}, \mathsf{c}_1, \mathsf{c}_3$), $\mathsf{a}_2 \sim \mathsf{A}(\mathsf{c}, \bar{\mathsf{a}}_2)$ and outputting $(\mathsf{a}_1, \mathsf{a}_2, \mathsf{a}_3)$. In the previous calculation we have ignored error terms of size $\mathcal{O}\big(|\mathbb{F}|^{-1}\big)$. Thus the marginal distribution on $(\mathsf{a}_1, \mathsf{a}_3)$ is $\mathcal{O}\big(|\mathbb{F}|^{-1}\big)-$close to uniform on $\mathsf{A}(\mathsf{c})$, and the result follows. □

References

[ADKO15] Divesh Aggarwal, Yevgeniy Dodis, Tomasz Kazana, and Maciej Obremski. Non-malleable reductions and applications. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 459–468, 2015.

[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.

[ADN+19] Divesh Aggarwal, Nico Döttling, Jesper Buus Nielsen, Maciej Obremski, and Erick Purwanto. Continuous non-malleable codes in the 8-split-state model. In *Advances in Cryptology - EUROCRYPT 2019 - 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Darmstadt, Germany, May 19-23, 2019, Proceedings, Part I*, pages 531–561, 2019.

[Agg15] Divesh Aggarwal. Affine-evasive sets modulo a prime. *Inf. Process. Lett.*, 115(2):382–385, 2015.

[AGM+15] Shashank Agrawal, Divya Gupta, Hemanta K. Maji, Omkant Pandey, and Manoj Prabhakaran. Explicit non-malleable codes against bit-wise tampering and permutations. In *Advances in Cryptology - CRYPTO 2015 - 35th Annual Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2015, Proceedings, Part I*, pages 538–557, 2015.

[AKO17] Divesh Aggarwal, Tomasz Kazana, and Maciej Obremski. Inception makes non-malleable codes stronger. In *Theory of Cryptography - 15th International Conference, TCC 2017, Baltimore, MD, USA, November 12-15, 2017, Proceedings, Part II*, pages 319–343, 2017.

[ALM+98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, 1998.

[AO19] Divesh Aggarwal and Maciej Obremski. Inception makes non-malleable codes shorter as well! *IACR Cryptology ePrint Archive*, 2019:399, 2019.

[AS97] Sanjeev Arora and Madhu Sudan. Improved low-degree testing and its applications. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing, El Paso, Texas, USA, May 4-6, 1997*, pages 485–495, 1997.

[AS98] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of NP. *J. ACM*, 45(1):70–122, 1998.

[BDG+18] Marshall Ball, Dana Dachman-Soled, Siyao Guo, Tal Malkin, and Li-Yang Tan. Non-malleable codes for small-depth circuits. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 826–837, 2018.

[BDN17] Amey Bhangale, Irit Dinur, and Inbal Livni Navon. Cube vs. cube low degree test. In *8th Innovations in Theoretical Computer Science Conference, ITCS 2017, January 9-11, 2017, Berkeley, CA, USA*, pages 40:1–40:31, 2017.

[BGW19] Marshall Ball, Siyao Guo, and Daniel Wichs. Non-malleable codes for decision trees. In *Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part I*, pages 413–434, 2019.

[BLR93] Manuel Blum, Michael Luby, and Ronitt Rubinfeld. Self-testing/correcting with applications to numerical problems. *J. Comput. Syst. Sci.*, 47(3):549–595, 1993.

[CDTV16] Sandro Coretti, Yevgeniy Dodis, Björn Tackmann, and Daniele Venturi. Non-malleable encryption: Simpler, shorter, stronger. In *Theory of Cryptography - 13th International Conference, TCC 2016-A, Tel Aviv, Israel, January 10-13, 2016, Proceedings, Part I*, pages 306–335, 2016.

[CGL16] Eshan Chattopadhyay, Vipul Goyal, and Xin Li. Non-malleable extractors and codes, with their many tampered extensions. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 285–298, 2016.

[CKR15] Nishanth Chandran, Bhavana Kanukurthi, and Srinivasan Raghuraman. Information-theoretic local non-malleable codes and their applications. *IACR Cryptology ePrint Archive*, 2015:1056, 2015.

[CL17] Eshan Chattopadhyay and Xin Li. Non-malleable codes and extractors for small-depth circuits, and affine functions. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1171–1184. ACM, 2017.

[CZ14] Eshan Chattopadhyay and David Zuckerman. Non-malleable codes against constant split-state tampering. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 306–315, 2014.

[CZ16] Eshan Chattopadhyay and David Zuckerman. Explicit two-source extractors and resilient functions. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 670–683, 2016.

[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, May 5-8, 1991, New Orleans, Louisiana, USA*, pages 542–552, 1991.

[DHK+19] Irit Dinur, Prahladh Harsha, Tali Kaufman, Inbal Livni Navon, and Amnon Ta-Shma. List decoding with double samplers. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 2134–2153, 2019.

[DHKR19] Irit Dinur, Prahladh Harsha, Tali Kaufman, and Noga Ron-Zewi. From local to robust testing via agreement testing. In *10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10-12, 2019, San Diego, California, USA*, pages 29:1–29:18, 2019.

[DJMW12] Yevgeniy Dodis, Abhishek Jain, Tal Moran, and Daniel Wichs. Counterexamples to hardness amplification beyond negligible. In *Theory of Cryptography - 9th Theory of Cryptography Conference, TCC 2012, Taormina, Sicily, Italy, March 19-21, 2012. Proceedings*, pages 476–493, 2012.

[DK17] Irit Dinur and Tali Kaufman. High dimensional expanders imply agreement expanders. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 974–985, 2017.

[DKO13] Stefan Dziembowski, Tomasz Kazana, and Maciej Obremski. Non-malleable codes from two-source extractors. In *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part II*, pages 239–257, 2013.

[DLSZ15] Dana Dachman-Soled, Feng-Hao Liu, Elaine Shi, and Hong-Sheng Zhou. Locally decodable and updatable non-malleable codes and their applications. In *Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015, Warsaw, Poland, March 23-25, 2015, Proceedings, Part I*, pages 427–450, 2015.

[DLSZ20] Dana Dachman-Soled, Feng-Hao Liu, Elaine Shi, and Hong-Sheng Zhou. Locally decodable and updatable non-malleable codes and their applications. *J. Cryptology*, 33(1):319–355, 2020.

[DP07] Stefan Dziembowski and Krzysztof Pietrzak. Intrusion-resilient secret sharing. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2007), October 20-23, 2007, Providence, RI, USA, Proceedings*, pages 227–237, 2007.

[DPW18] Stefan Dziembowski, Krzysztof Pietrzak, and Daniel Wichs. Non-malleable codes. *J. ACM*, 65(4):20:1–20:32, 2018.

[FS95] Katalin Friedl and Madhu Sudan. Some improvements to total degree tests. In *Third Israel Symposium on Theory of Computing and Systems, ISTCS 1995, Tel Aviv, Israel, January 4-6, 1995, Proceedings*, pages 190–198, 1995.

[GPR16] Vipul Goyal, Omkant Pandey, and Silas Richelson. Textbook non-malleable commitments. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 1128–1141, 2016.

[GR19] Vipul Goyal and Silas Richelson. Non-malleable commitments using goldreich-levin list decoding. In *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, pages 686–699, 2019.

[GS06] Oded Goldreich and Madhu Sudan. Locally testable codes and pcps of almost-linear length. *J. ACM*, 53(4):558–655, 2006.

[Ham50] Richard Hamming. Error detecting and error correcting codes. In *Bell System Technical Journal*, pages 147–160, 1950.

[IKW12] Russell Impagliazzo, Valentine Kabanets, and Avi Wigderson. New direct-product testers and 2-query pcps. *SIAM J. Comput.*, 41(6):1722–1768, 2012.

[Jus72] Jørn Justesen. Class of constructive asymptotically good algebraic codes. *IEEE Trans. Information Theory*, 18(5):652–656, 1972.

[Li16] Xin Li. Improved two-source extractors, and affine extractors for polylogarithmic entropy. In *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Hyatt Regency, New Brunswick, New Jersey, USA*, pages 168–177, 2016.

[Li17] Xin Li. Improved non-malleable extractors, non-malleable codes and independent source extractors. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 1144–1156, 2017.

[Li19] Xin Li. Non-malleable extractors and non-malleable codes: Partially optimal constructions. In *34th Computational Complexity Conference, CCC 2019, July 18-20, 2019, New Brunswick, NJ, USA.*, pages 28:1–28:49, 2019.

[Mos17] Dana Moshkovitz. Low-degree test with polynomially small error. *Computational Complexity*, 26(3):531–582, 2017.

[RS60] Irving Reed and Gustave Solomon. Polynomial codes over certain finite fields. *Journal for the Society for Industrial and Applied Mathematics*, 8:300–304, 1960.

[RS96] Ronitt Rubinfeld and Madhu Sudan. Robust characterizations of polynomials with applications to program testing. *SIAM J. Comput.*, 25(2):252–271, 1996.

[RS97] Ran Raz and Shmuel Safra. A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing, El Paso, Texas, USA, May 4-6, 1997*, pages 475–484, 1997.

[Sha49] Claude Shannon. A mathematical theory of communication. In *Urbana, IL: University of Illinois Press*, 1949.

[Sud97] Madhu Sudan. Decoding of reed solomon codes beyond the error-correction bound. *J. Complexity*, 13(1):180–193, 1997.

[WZ93] Avi Wigderson and David Zuckerman. Expanders that beat the eigenvalue bound: explicit construction and applications. In *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, May 16-18, 1993, San Diego, CA, USA*, pages 245–251, 1993.

[Zuc97] David Zuckerman. Randomness-optimal oblivious sampling. *Random Struct. Algorithms*, 11(4):345–367, 1997.

A Sampler Replacement

In the body we used the following fact with $(\varepsilon', \delta') = (\varepsilon, \delta)$ and $\rho = \zeta = \varepsilon$.

Fact 3 (Restated). *Let $\varepsilon, \delta, \varepsilon', \delta', \varepsilon^*, \delta^*, \rho, \zeta > 0$ be such that $\delta^*(\varepsilon^* - \varepsilon - \varepsilon' - 2\rho - \zeta) \geq \delta'/\zeta + \delta/\rho$. Suppose $A/B/C$ is such that:*

- *A/C, B/C and $B(a)/C(a)$ are $0-$biregular for all $a \in A$; and*
- *A/C is $(\varepsilon, \delta)-$sampling and $A(c)/B(c)$ is $(\varepsilon', \delta')-$sampling for all $c \in C$.*

Then A/B is $(\varepsilon^*, \delta^*)-$sampling.

Proof. Fix $\varepsilon, \delta, \varepsilon', \delta', \varepsilon^*, \rho, \zeta > 0$ and $A/B/C$ as in the statement. Let $B' \subset B$ be a set of size $|B'| = \lambda \cdot |B|$, and let $A' \subset A$ be the set of $a \in A$ such that $\left|\Pr_{b\sim B(a)}(b \in B') - \lambda\right| > \varepsilon^*$, let $\nu = |A'|/|A|$. We must show that $\nu \leq \big(\delta'/\zeta + \delta/\rho\big)/(\varepsilon^* - \varepsilon - \varepsilon' - 2\rho - \zeta)$. We have

$$\begin{aligned}
\varepsilon^* &< \mathbb{E}_{a\sim A'}\Big[\left|\Pr_{b\sim B(a)}(b \in B') - \lambda\right|\Big] \leq \mathbb{E}_{a\sim A'}\left[\left|\mathbb{E}_{c\sim C(a)}\big[\Pr_{b\sim B(a,c)}(b \in B')\big] - \lambda\right|\right] \\
&\leq \mathbb{E}_{\substack{a\sim A'\\ c\sim C(a)}}\left[\left|\Pr_{b\sim B(a,c)}(b \in B') - \lambda(c)\right|\right] + \mathbb{E}_{a\sim A'}\left[\left|\mathbb{E}_{c\sim C(a)}\big[\lambda(c)\big] - \mathbb{E}_{c\sim C}\big[\lambda(c)\big]\right|\right],
\end{aligned}$$

where for $c \in C$, $\lambda(c) := \Pr_{b\sim B(c)}(b \in B')$. We have used the biregularity of $B(a)/C(a)$ for all $a \in A$ and that $\mathbb{E}_{c\sim C}\big[\lambda(c)\big] = \lambda$, which follows from biregularity of B/C. Let RHS_1 and RHS_2 be the two expectations on the right hand side of the equation above. We bound RHS_1 and RHS_2 separately. Note,

$$\mathsf{RHS}_2 \leq \varepsilon + 2\rho + \nu^{-1} \cdot \Pr_{a\sim A}\left[\left|\mathbb{E}_{c\sim C(a)}\big[\lambda(c)\big] - \mathbb{E}_{c\sim C}\big[\lambda(c)\big]\right| > \varepsilon + 2\rho\right] \leq \varepsilon + 2\rho + \nu^{-1} \cdot \delta/\rho.$$

Thus, it suffices to show that $\mathsf{RHS}_1 \leq \zeta + \varepsilon' + \nu^{-1} \cdot \delta'/\zeta$. Let $C' \subset C$ be the set of $c \in C$ such that $\Pr_{\substack{a\sim A'\\ c'\sim C(a)}}(c' = c) < \zeta/|C|$. Clearly, $\Pr_{\substack{a\sim A'\\ c\sim C(a)}}(c \in C') < \zeta$. Also, whenever $c \notin C'$, we have

$$\nu \cdot \zeta \leq \nu \cdot |C| \cdot \Pr_{\substack{a\sim A\\ c'\sim C(a)}}\big[c' = c \big| a \in A'\big] = \Pr_{\substack{c'\sim C\\ a\sim A(c')}}\big[a \in A' \big| c' = c\big] = \Pr_{a\sim A(c)}\big[a \in A'\big].$$

We have used the biregularity of A/C. This gives

$$\begin{aligned}
\mathsf{RHS}_1 &< \zeta + \varepsilon' + \max_{c\notin C'}\left\{\Pr_{a\sim A(c)}\left[\left|\Pr_{b\sim B(a,c)}(b \in B') - \lambda(c)\right| > \varepsilon'\right] \Big/ \Pr_{a\sim A(c)}(a \in A')\right\} \\
&\leq \zeta + \varepsilon' + \nu^{-1} \cdot \delta'/\zeta,
\end{aligned}$$

and the result follows. □

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il