FOCS 99 submission.

A tool for obtaining tighter security analyses of pseudorandom function based constructions, with applications to PRP→PRF conversion

Mihir Bellare* Russell Impagliazzo†

February 1999

Abstract

We present a general probabilistic lemma that can be applied to upper bound the advantage of an adversary in distinguishing between two families of functions. Our lemma reduces the task of upper bounding the advantage to that of upper bounding the ratio of two probabilities associated to the adversary, when this ratio is is viewed as a random variable. It enables us to obtain significantly tighter analyses than more conventional methods.

In this paper we apply the technique to the problem of PRP to PRF conversion. We present a simple, new construction of a PRF from a PRP that makes only two invocations of the PRP and has insecurity linear in the number of queries made by the adversary. We also improve the analysis of the truncation construction.

Keywords: Pseudorandom functions, pseudorandom permutations, provable security, birthday attacks.

*Dept. of Computer Science & Engineering, University of California at San Diego, 9500 Gilman Drive, La Jolla, CA 92093, USA. E-Mail: mihir@cs.ucsd.edu. URL: http://www-cse.ucsd.edu/users/mihir/. Supported in part by NSF CAREER Award CCR-9624439 and a 1996 Packard Foundation Fellowship in Science and Engineering.

†Dept. of Computer Science & Engineering, Mail Code 0114, University of California at San Diego, 9500 Gilman Drive, La Jolla, CA 92093, USA. E-mail: russell@cs.ucsd.edu. URL: http://www-cse.ucsd.edu/users/russell.

1 Introduction

In the "provable security" approach of Goldwasser and Micali [GM] one bases the security of a higher level cryptographic primitive on a lower level one via security reduction of polynomial complexity. This shows that the two securities are polynomially-related, which we interpret as saying that if the lower level primitive is secure, so is the higher level one. A major effort in modern cryptography has been to find designs of key cryptographic primitives with security polynomially related to as-weak-as-possible an underlying primitive. Instances of this are the basing of pseudorandom bit generators and signature schemes on one-way functions [HILL, Ro]. This line of work has been very successful, and at this point such relations are known for the central primitives.

More recently, attention has been turning to the tightness of security reductions, that is, their "concrete security". This natural second phase of research in complexity-based cryptography has several motivations. An important one of these is that the complexity of the security reduction has a direct impact on the efficiency of the scheme in applications. Let us see why.

Suppose we have two different schemes for solving a particular problem, for example, signature schemes. Furthermore, suppose they have the same asymptotic running time; for example, both schemes run in time cubic in the security parameter. That does not mean both are equally efficient in usage, because one must take into account the desired level of security. Suppose the first scheme has security loosely related to that of the underlying primitive, while the security of the second scheme is tightly related to that of the underlying primitive, whatever that primitive might be. This means that to get a desired level of security, we must use the first scheme with a higher value of the security parameter than we need for the second scheme. So although both algorithms are cubic, we may need to set the security parameter in the first case to the square of the one in the second case for the same level of security, so that the first algorithm effectively becomes $O(k^6)$ where k is the security paramter.

The efficiency of security reductions has been treated by Goldreich et. al. [GILVZ] in the context of one-way permutations. More recently, there has been much interest in this subject in the context of pseudorandom functions [BCK, BGK]. The reason is that pseudorandom functions, introduced by [GGM], are a useful tool in cryptographic protocol design for both theoretical and practical reasons.

However, experience has shown that given even quite a simple construction it is often not easy to get a security analysis, let alone a precise and tight one. Analyses for known constructions tend to be *ad hoc* and often there are significant gaps between known upper and lower bounds on the loss of security.

An instance is the problem of converting pseudorandom permutations (PRPs) [LR] into pseudorandom functions. (A PRP is a family of functions that is indistinguishable from the family of random permutations, while a PRF is a family of functions that is indistinguishable from the family of random functions.) This problem, first considered in [BKrR], at first sounds paradoxical, since there seems nothing to do; isn't any PRP itself a PRF? (Indeed, effort, beginning with [LR], usually goes into turning PRFs into PRPs, not PRPs into PRFs.) This is true if we are only interested in polynomial security relations, but not if we are interested in achieving optimal security. In that context the problem is both non-trivial and well-motivated.

The motivation for this problem is that it is reasonable to assume that block ciphers behave like PRPs. Yet, experience with scheme design has shown that PRFs are the more useful tool. Many designs using PRPs simply view them as PRFs, and hence incur a loss of security. For the efficiency and other reasons mentioned above, we would like transformations of PRPs to PRFs that are as optimal as possible, meaning losing very little security.

The loss of security in simply regarding a PRP as a PRF comes from the "birthday attack" which can distinguish a random permutation from a random function in time $2^{n/2}$ where n is the block-length of the function. To prevent such attacks it has been suggested to first convert the block cipher into a pseudorandom function [BKrR]. It is easy to think of a variety of transformations of pseudorandom

permutations (block ciphers) into pseudorandom functions. However for the intended application it is important that the transformation looses very little security even when the adversary is allowed time greater more than $2^{n/2}$ queries to the function. In particular a quantitative loss of security of $q^2/2^n$ after q queries is practically meaningless in this context, although considered reasonable in other contexts. Of course the smaller the loss in security, the better, even in constructions not subject to birthday attacks. Several transformations of PRPs into PRFs are known [BKrR, HWKS] but none achieving an optimal security.

In this paper we present a general tool for obtaining improved security analyses of constructions of pseudorandom functions of various kinds. The analysis we obtain via our technique is often a quadratic improvement over those obtained by more standard techniques. The technique is very general, reducing the task of computing the statistical distance between two function families to estimating the ratio of certain probabilities associated to the adversary. It is described in Section 3.

In this paper we limit ourselves to applying it to constructions of pseudorandom functions from pseudorandom permutations. We analyze two constructions. The first is a new construction which we show has a loss of security which is only linear in the number of queries. Thus our new construction is provably superior to the truncation one for large numbers of queries. The construction is quite simple and efficient: it consists simply of taking the value of the given PRP at two points and XORing the results. See Section 4. The second construction we analyze is the truncation construction of [HWKS]. We improve their analysis. Our analysis is tight for certain ranges of the parameters, in particular when the adversary makes a large number of queries.

Our main theorem is a very general analysis of the statistical distance of two distributions on q-element sequences. We show that if with high probability the conditional distributions on the i-th element given the previous elements are close in the sense of having a ratio whose difference from 1 is bounded by a quantity $\delta < 1$ then the statistical distance between the original distributions is bounded by roughly $\delta q^{1/2}$. The bound obtained by a more conventional argument would be δq.

We start by presenting some definitions. The main theorem is in Section 3. In that section we state the main lemmas and prove the theorem given them. The proofs of the lemmas are in the appendices. The applications to PRP to PRF conversion are in Section 4. The main theorem is the technical brunt of the paper and is proved in full. Proofs of the applications are sketched.

Bellare, Goldreich and Krawczyk [BGK] show that using the XOR of the values of PRF on a number of random points can increase security of several applications relative to using the value on a single point. Their motivation was to avoid the use of state information, like counters, in applications. However, they assumed they were in possession of PRFs with good security. An application of our techniques is to be able to start with PRPs, use them to implement high-security PRFs, and then plug these into the constructions of [BGK], achieving their goals from the more practical starting primitive of a PRP.

2 Some definitions

Let D: $R^q \to [0,1]$ be a probability distribution on the set R^q of all n-element sequences over some base set R. For any $i \in [q-1]$ we extend D to R^i in the obvious say, namely for any $b_1, \ldots, b_i \in R$ we let

$$D(b_1, \ldots, b_i) \;=\; \Pr_{(a_1,\ldots,a_q)\leftarrow D}\left[\,(a_1, \ldots, a_i) = (b_1, \ldots, b_i)\,\right].$$

For any $i \in [q]$ we then set $\mathrm{Supp}(D, i) = \{\ (b_1, \ldots, b_i) \in R^i \ :\ D(b_1, \ldots, b_i) > 0\ \}$. Note that $\mathrm{Supp}(D, i) \subseteq R^i$, and $\mathrm{Supp}(D, q) \subseteq R^q$ is the support of the distribution D. For any $i \in [q-1]$, any $b_1, \ldots, b_{i-1} \in R$ and any $b \in R$ we let

$$D(b_1, \ldots, b_{i-1}; b) \;=\; \Pr_{(a_1,\ldots,a_q)\leftarrow D}\left[\,a_i = b \;\mid\; (a_1, \ldots, a_{i-1}) = (b_1, \ldots, b_{i-1})\,\right]$$

be the probability of obtaining b upon drawing a sequence $(a_1, \ldots, a_q)$ randomly according to D conditional on the first $i-1$ elements of the drawn sequence equaling $b_1, \ldots, b_{i-1}$ respectively. Notice that this conditional probability is only well defined if $D(b_1, \ldots, b_{i-1}) > 0$. We adopt the convention that $D_i(b_1, \ldots, b_{i-1}; b) = \infty$ if $D(b_1, \ldots, b_{i-1}) = 0$.

FUNCTION FAMILIES. A family of functions is a map F: $\mathrm{Keys}(F) \times \mathrm{Dom}(F) \to \mathrm{Range}(F)$. Each key $K \in \mathrm{Keys}(F)$ specifies a specific function $F_K \stackrel{\text{def}}{=} F(K, \cdot)$: $\mathrm{Dom}(F) \to \mathrm{Range}(F)$ in the family. A family of permutations is a family of functions F in which $\mathrm{Dom}(F) = \mathrm{Range}(F)$ and each individual function is a permutation on this set. The key set is equipped with some underlying distribution, and the notation $f \leftarrow F$ is shorthand for $K \leftarrow \mathrm{Keys}(F)\,;\ f \leftarrow F_K$. If FT, FB are families of functions and A is an adversary that is given an oracle f, its advantage in distinguishing FT from FB is defined as per [GGM] by

$$\mathsf{Adv}(A, FT, FB) \;=\; \left| \Pr_{f \leftarrow FT}\left[A^f = 1\right] - \Pr_{f \leftarrow FB}\left[A^f = 1\right] \right| .$$

A PRP is a family of permutations, and a PRF is a family of functions. Concrete security definitions for these following [BKrR] will be given in Section 4.

3 Ratio based comparison theorem

In this section we state and prove the theorem which provides a general way of determining the statistical distance between two function families. Later, we will apply it to the problem of PRP to PRF conversion. Let us begin by motivating the theorem in terms of applications to estimating the distance between function families.

3.1 Application setting

Suppose we fix two families of functions, FT: $\mathrm{Keys}(FT) \times D \to R$ and FB: $\mathrm{Keys}(FB) \times D \to R$, having the same domain D and range R. Consider an adversary A that is given an oracle f. Suppose A makes q queries of its oracle. We are interested in upper bounding $\mathsf{Adv}(A, FT, FB)$ as a function of q and the parameters of the families.

This setting is purely information-theoretic; we put no restrictions on the running time of A, but only on the number of queries it makes. This is because in analyses of pseudorandom function/permutation based constructions, the main technical content is usually in such an information theoretic question. Translation to the computational setting is then a standard argument.

Typical analyses try to isolate some "bad event" such that conditioned on this not happening, the view of the adversary is the same in both games. Bounding the advantage then reduces to bounding the probability of the bad event. This kind of approach, however, will not always work, because in some constructions, there is no such bad event. Rather, the two games differ continuously.

Our first step is to consider two associated distributions, B and T. Given that the adversary is fixed and can be assumed deterministic, a given set of responses to oracle queries uniquely determines the next query. Now consider a sequence $a_1, \ldots, a_q$ of values in the range R. We let $T(a_1, \ldots, a_q)$ (respectively $B(a_1, \ldots, a_q)$) be the probability that $a_1, \ldots, a_q$ is the sequence of answers from the oracle when A is run with an oracle f chosen at random from FT (respectively FB). Our goal is now to upper bound the statistical distance between B and T. We consider the quantity

$$\frac{T(a_1, \ldots, a_{i-1}; a_i)}{B(a_1, \ldots, a_{i-1}; a_i)} ,$$

with the notation being as defined above in Section 2. Namely it is the ratio of two conditional probabilities. Our ratio based comparison theorem says, roughly, that bounding $\mathsf{Adv}(A, FT, FB)$ reduces to analyzing the above ratio as a random variable over the choices of $a_1, \ldots, a_{i-1}$ drawn

according to B. In the formalization below, we don't look explicitly at such a ratio because of technicalities like the fact of the denominator being zero and the ratio being undefined, but look instead at a difference measure. The theorem below gives a bound in terms of this measure. In applications, we are left with the task of estimating this measure.

3.2 Statement of the ratio based comparison theorem

The general statement does not refer to function families; it considers two arbitrary probability distributions. Namely let $B, T\colon\ R^q \to [0,1]$ be two probability distributions on the set R^q of q-element sequences over a base set R. We want to define a measure of how they compare to each other. For any real number $\delta \in [0,1]$ we set

$$\begin{aligned}&\mathsf{Diff}_{B,T}(a_1,\ldots,a_{i-1};\delta)\\ &\quad = \max_{a_i \in R}\{\ |\ T(a_1,\ldots,a_{i-1};a_i) - B(a_1,\ldots,a_{i-1};a_i)\ |\ -\ \delta \cdot B(a_1,\ldots,a_{i-1};a_i)\ \}\ .\end{aligned}$$

This quantity must be appropriately interpreted when some of the quantities involved are ∞. It turns out (below) that we will only be looking at this when $B(a_1,\ldots,a_{i-1};a_i)$ is well defined, so the only question is what happens when $T(a_1,\ldots,a_{i-1};a_i) = \infty$. In that case, the entire quantity $\mathsf{Diff}_{B,T}(a_1,\ldots,a_{i-1};\delta)$ has value ∞, and in particular is positive, which is what counts below.

We are interested in how this difference behaves when $a_1,\ldots,a_q$ are drawn according to B. We set

$$\mathsf{Dev}(B,T;\delta)\ =\ \mathrm{Pr}_{(a_1,\ldots,a_q)\leftarrow B}\left[\ \exists i \in [q]\ :\ \mathsf{Diff}_{B,T}(a_1,\ldots,a_{i-1};\delta) > 0\right]\ .$$

Let $\mathsf{Dist}(B,T)$ denote the statistical distance between distributions B and T.

Theorem 3.1 Let $B, T\colon\ R^q \to [0,1]$ be two probability distributions on the set R^q of all q-element sequences over a base set R. Suppose $0 \le \delta < 1$. Then for any $\lambda > 0$ we have

$$\mathsf{Dist}(B,T)\ \le\ \mathsf{Dev}(B,T;\delta) + \frac{e-1}{2}\cdot\left(\delta^2 q + 3\delta q^{1/2}\lambda\right) + 2e^{-\lambda^2/2}\ . \tag{1}$$

In other words, the statistical distance between B and T can be bounded in terms of the deviation $\mathsf{Dev}(B,T;\delta)$ and $O(\delta^2 q + \delta q^{1/2}\lambda)$. It will be enough that λ is logarithmic, so that the last term of the bound is small. The following corollary is obtained simply by optimizing, namely plug in an appropriate value of λ and simplify.

Corollary 3.2 Let $B, T\colon\ R^q \to [0,1]$ be two probability distributions on the set R^q of all q-element sequences over a base set R. Suppose $0 \le \delta < 1$ and define

$$\lambda(\delta,q)\ =\ \max\left[1,\ \sqrt{2\lg\frac{1}{\delta q^{1/2}}}\right]\ .$$

Then

$$\mathsf{Dist}(B,T)\ \le\ \mathsf{Dev}(B,T;\delta) + [4.3\cdot\lambda(\delta,q)]\cdot\delta q^{1/2}\ .$$

Proof: We set $\lambda = \lambda(\delta,q)$ and apply Theorem 3.1. The value of the last term in Equation (1) is

$$2e^{-\lambda^2/2}\ \le\ 2e^{2\lg(\delta q^{1/2})/2}\ =\ 2\delta q^{1/2}/2\ =\ \delta q^{1/2}\ .$$

Hence from Equation (1) we get

$$\begin{aligned}\mathsf{Dist}(B,T)\ &\le\ \mathsf{Dev}(B,T;\delta) + \frac{e-1}{2}\cdot\left(\delta^2 q + 3\delta q^{1/2}\lambda\right) + \delta q^{1/2}\\ &\le\ \mathsf{Dev}(B,T;\delta) + \left[\frac{5(e-1)}{2}\lambda\right]\cdot\delta q^{1/2}\ .\end{aligned}$$

However the constant above is less than 4.3. ▌

THE KEY IMPROVEMENT. It is trivial to obtain a bound of the form $\mathsf{Dist}(B,T) \leq \mathsf{Dev}(B,T;\delta)+O(\delta q)$. The key improvement, as evidenced by the above corollary, is that the additional term for us is proportional to $\delta q^{1/2}$ as opposed to δq. (The $\lambda(\delta, q)$ term is negligible, since it is logarithmic.)

3.3 Probabilistic lemmas

The proof of Theorem 3.1 will make use of three general probability lemmas which are summarized here, with the proofs in appendices.

STATISTICAL DISTANCE FROM RATIOS. We will reduce the problem of bounding the statistical distance between two distributions B and T to the problem of bounding the probability that B, T have a small comparative ratio, via the following lemma.

Lemma 3.3 Let B, T: $S \to [0,1]$ be probability distributions on a set S, and $\alpha, \gamma \in [0,1]$. Assume

$$\Pr_{x \leftarrow B}\left[1-\alpha \leq \frac{T(x)}{B(x)} \leq 1+\alpha\right] \;\geq\; 1-\gamma\,.$$

Then $\mathsf{Dist}(B,T) \leq \alpha+\gamma$.

The proof of this lemma is in Appendix A.1.

ALMOST MARTINGALE TAIL INEQUALITY. A standard version of Azuma's inequality considers a sequence of bounded random variables having zero conditional expectation, and provides a bound on the probability of their sum being large. We will need a generalization of this inequality.

Definition 3.4 Let μ: $R^n \to [0,1]$ be a probability distribution on the set R^n of all n-element sequences over a set R. Let L_i: $\mathrm{Supp}(\mu, i) \to \mathbf{R} \cup \{\pm\infty\}$ for $i = 1, \ldots, n$ be functions. Let $\delta_1, \delta_2 \in \mathbf{R}$ and let $m \in [n]$. We say that $(b_1, \ldots, b_n) \in R^n$ is (δ_1, δ_2, m)-*good for* $L_1, \ldots, L_m$ *over* μ if the following two conditions hold–

(1) *Bounded conditional expectation:* For all $i \in [m]$–

$$\left|\,\mathbf{E}_{(a_1,\ldots,a_n)\leftarrow\mu}\left[\,L_i(a_1,\ldots,a_i) \;\middle|\; (a_1,\ldots,a_{i-1}) = (b_1,\ldots,b_{i-1})\,\right]\,\right| \leq \delta_1$$

(2) *Bounded range:* For all $i \in [m]$–

$$\Pr_{(a_1,\ldots,a_n)\leftarrow\mu}\left[\,|\,L_i(a_1,\ldots,a_i)\,| > \delta_2 \;\middle|\; (a_1,\ldots,a_{i-1}) = (b_1,\ldots,b_{i-1})\,\right] \;=\; 0\,.$$

We say that $(b_1, \ldots, b_n) \in R^n$ is (δ_1, δ_2)-*good for* $L_1, \ldots, L_n$ *over* μ if it is (δ_1, δ_2, n)-good for $L_1, \ldots, L_n$ over μ.

The following lemma says that the sum $L_1 + \ldots + L_n$ can be bounded in absolute value as in a standard martingale tail inequality modulo a term accounting for the probability of a bad sequence.

Lemma 3.5 [**Almost Martingale Tail Inequality**] Let μ: $R^n \to [0,1]$ be a probability distribution on the set R^n of all n-element sequences over a set R. Let L_i: $\mathrm{Supp}(\mu, i) \to \mathbf{R} \cup \{\pm\infty\}$ for $i \in [n]$ be functions. Let $\delta_1, \delta_2, \beta \in \mathbf{R}_{\geq 0}$. Suppose

$$\Pr_{(b_1,\ldots,b_n)\leftarrow\mu}\left[\,(b_1,\ldots,b_n) \text{ is } (\delta_1,\delta_2)\text{-good for } L_1,\ldots,L_n \text{ over } \mu\,\right] \;\geq\; 1-\beta\,.$$

Then

$$\Pr_{(a_1,\ldots,a_n)\leftarrow\mu}\left[\,|\,\textstyle\sum_{i=1}^{n} L_i(a_1,\ldots,a_i)\,| \geq \delta_1 n + \delta_2 n^{1/2}\lambda\right] \;\leq\; \beta + 2e^{-\lambda^2/2}\,.$$

The proof of Lemma 3.5 is in Section A.2.

Notice that $L_1, \ldots, L_n$ are allowed to take the values $\pm\infty$. We adopt the convention that

$$| \sum_{i=1}^{n} L_i(a_1, \ldots, a_i) | \;=\; \infty$$

if any of the terms in the sum is $\pm\infty$. In particular when any term in the sum is $\pm\infty$ then the sum cannot be bounded in absolute value by $\delta_1 n + \delta_2 n^{1/2}\lambda$, so the lemma is implying that such terms are covered by the failure probability.

SQUARED EXPECTATION BOUND. Consider two distributions μ_1, μ_2 which are close in the sense that $|\mu_2(x)/\mu_1(x)| \leq (1+\delta)\mu_1(x)$ for all x. The next lemma says that if look at the expectation of the log of the ratio of the two probabilities, then it will be significantly smaller than δ, in fact around δ^2.

Lemma 3.6 Let μ_1, μ_2: $S \to [0,1]$ be two probability distributions on a set S, and $\delta \in [0,1]$, such that $| \mu_2(x) - \mu_1(x) | \leq \delta \cdot \mu_1(x)$ for all $x \in S$. Then

$$\left| \mathop{\mathbf{E}}_{x \leftarrow \mu_1} \left[\ln \frac{\mu_2(x)}{\mu_1(x)} \right] \right| \leq \frac{\delta^2}{2} \, .$$

The proof of Lemma 3.6 is in Appendix A.3.

3.4 Main lemma, and proof of theorem

We state the main lemma, and then prove the theorem assuming it. The proof of the lemma is given in Section 3.5.

Lemma 3.7 Let B, T: $R^q \to [0,1]$ be two probability distributions on the set R^q of all q-element sequences over a base set R. Suppose $0 \leq \delta < 1$. Let $\delta_1 = \delta^2/2$ and $\delta_2 = 3\delta/2$, and suppose $\lambda > 0$. Let $\epsilon = \delta_1 q + \delta_2 q^{1/2}\lambda$. Then

$$\Pr_{(a_1,\ldots,a_q)\leftarrow B} \left[\left| \ln \frac{T(a_1, \ldots, a_q)}{B(a_1, \ldots, a_q)} \right| \leq \epsilon \right] \;\geq\; 1 - \gamma \, , \tag{2}$$

where $\gamma = \mathsf{Dev}(B, T; \delta) + 2e^{-\lambda^2/2}$.

Put another way, say $(a_1, \ldots, a_q) \in \mathrm{Supp}(B, q)$ is good if $| \ln[T(a_1, \ldots, a_q)/B(a_1, \ldots, a_q)] | \leq \epsilon$. The lemma is showing how to lower bound the probability of drawing a good sequence under B.

Notice that $| \ln[T(a_1, \ldots, a_q)/B(a_1, \ldots, a_q)] |$ might be ∞; this happens when $T(a_1, \ldots, a_q) = 0$. (It is always the case that $B(a_1, \ldots, a_q) > 0$ since we are drawing $(a_1, \ldots, a_q)$ according to B.) That is not a problem; by definition such sequences are not good, so the lemma implies that their probability is covered by γ.

Proof of Theorem 3.1: We prove Theorem 3.1 assuming Lemma 3.7. Note that if $\epsilon \geq 1$ then the conclusion of Theorem 3.1 is trivially true, so we may assume $\epsilon \leq 1$. Let $\alpha = (e-1)\epsilon$. We claim that

$$\Pr_{(a_1,\ldots,a_q)\leftarrow B} \left[1 - \alpha \leq \frac{T(a_1, \ldots, a_q)}{B(a_1, \ldots, a_q)} \leq 1 + \alpha \right] \;\geq\; 1 - \gamma \, . \tag{3}$$

Given this, the theorem follows from Lemma 3.3. It remains to establish Equation (3). We start from Equation (2) and exponentiate the quantities in the probability expression to get

$$\Pr_{(a_1,\ldots,a_q)\leftarrow B} \left[e^{-\epsilon} \leq \frac{T(a_1, \ldots, a_q)}{B(a_1, \ldots, a_q)} \leq e^{\epsilon} \right] \;\geq\; 1 - \gamma \, .$$

Now we can use the inequalities $e^{-x} \geq 1 - x$ and $e^x \leq 1 + (e-1)x$ valid for all $x \in [0,1]$. Setting $x = \epsilon$ (recall $\epsilon \leq 1$ so the inequalities apply, and $\alpha = (e-1)\epsilon$) we get Equation (3) as desired. ∎

3.5 Proof of Lemma 3.7

For each $i = 1, \ldots, n$ we define a function L_i: $\mathrm{Supp}(B, i) \to \mathbf{R} \cup \{\pm\infty\}$ as follows: for any $(a_1, \ldots, a_i) \in \mathrm{Supp}(B, i)$ let

$$L_i(a_1, \ldots, a_i) \;=\; \ln \frac{T(a_1, \ldots, a_{i-1}; a_i)}{B(a_1, \ldots, a_{i-1}; a_i)} \; .$$

Notice that the denominator in the fraction is always non-zero due to the choice of the domain of the function. On the other hand the numerator might be ∞, in which case $L_i(a_1, \ldots, a_i)$ is itself ∞; and the numerator might be 0, in which case $L_i(a_1, \ldots, a_i)$ is $-\infty$.

We wish to apply Lemma 3.5 to $L_1, \ldots, L_q$ viewed as random variables over the distribution $\mu = B$ on R^q. The first step will be to lower bound the probability of (δ_1, δ_2)-good sequences for appropriate values of δ_1, δ_2. For this purpose we let

$$G \;=\; \{\, (b_1, \ldots, b_q) \in \mathrm{Supp}(B, q) \;:\; \forall i \in [q] \text{ we have } \mathsf{Diff}_{B,T}(b_1, \ldots, b_{i-1}; \delta) \leq 0 \,\} \; .$$

The following two claims together are saying that if $(b_1, \ldots, b_q) \in G$ then $(b_1, \ldots, b_q)$ is (δ_1, δ_2)-good for $L_1, \ldots, L_q$ over B, where $\delta_1 = \delta^2/2$ and $\delta_2 = 3\delta/2$.

Claim 3.8 If $(b_1, \ldots, b_q) \in G$ then for every $i \in [q]$ we have

$$\left| \mathbf{E}_{(a_1,\ldots,a_q) \leftarrow B} \left[\, L_i(a_1, \ldots, a_i) \;\mid\; (a_1, \ldots, a_{i-1}) = (b_1, \ldots, b_{i-1}) \,\right] \right| \;\leq\; \delta_1 \tag{4}$$

where $\delta_1 = \delta^2/2$.

Claim 3.9 If $(b_1, \ldots, b_q) \in G$ then for every $i \in [q]$ we have

$$\Pr_{(a_1,\ldots,a_q) \leftarrow \mu} \left[\, |\, L_i(a_1, \ldots, a_i)\, | > \delta_2 \;\mid\; (a_1, \ldots, a_{i-1}) = (b_1, \ldots, b_{i-1}) \,\right] \;=\; 0 \; ,$$

where $\delta_2 = 3\delta/2$.

Based on these claims we can use Lemma 3.5 to prove Lemma 3.7. Let us do that and then return to the proofs of the claims.

Proof of Lemma 3.7: We claim that for all $(a_1, \ldots, a_q) \in \mathrm{Supp}(B, q)$ we have

$$\left| \ln \frac{T(a_1, \ldots, a_q)}{B(a_1, \ldots, a_q)} \right| \;=\; |\, \textstyle\sum_{i=1}^{q} L_i(a_1, \ldots, a_i)\, | \; . \tag{5}$$

To establish this we consider cases for $(a_1, \ldots, a_q) \in \mathrm{Supp}(B, q)$. First if $(a_1, \ldots, a_q) \in \mathrm{Supp}(T, q)$ then by conditioning we have

$$\ln \frac{T(a_1, \ldots, a_q)}{B(a_1, \ldots, a_q)} \;=\; \ln \prod_{i=1}^{n} \frac{T(a_1, \ldots, a_{i-1}; a_i)}{B(a_1, \ldots, a_{i-1}; a_i)} \;=\; \sum_{i=1}^{n} L_i(a_1, \ldots, a_i) \; ,$$

so Equation (5) is true in this case. Now suppose $(a_1, \ldots, a_q) \in \mathrm{Supp}(B, q) - \mathrm{Supp}(T, q)$, meaning $T(a_1, \ldots, a_q) = 0$. Then $\ln[T(a_1, \ldots, a_q)/B(a_1, \ldots, a_q)]$ is $-\infty$. On the other hand, in the sum $\sum_{i=1}^{q} L_i(a_1, \ldots, a_i)$ there will be terms that are $-\infty$, and possibly also terms that are ∞. Under the convention that a sum of terms involving $\pm\infty$ has absolute value ∞, Equation (5) is again true.

The proof of Lemma 3.7 is now concluded by showing that

$$\Pr_{(a_1,\ldots,a_q) \leftarrow B} \left[\, |\, \textstyle\sum_{i=1}^{n} L_i(a_1, \ldots, a_i)\, | \geq \epsilon \,\right] \;\leq\; \gamma \; . \tag{6}$$

We establish Equation (6) by applying Lemma 3.5 to $L_1, \ldots, L_q$ with $\mu = B$. To do this, note

$$\begin{aligned} &\Pr_{(b_1,\ldots,b_q) \leftarrow B} \left[\, (b_1, \ldots, b_q) \text{ is } (\delta_1, \delta_2)\text{-good for } L_1, \ldots, L_q \text{ over } B \,\right] \\ &\qquad \geq \;\; \Pr_{(b_1,\ldots,b_q) \leftarrow B} \left[\, (b_1, \ldots, b_q) \in G \,\right] \\ &\qquad \geq \;\; 1 - \mathsf{Dev}(B, T; \delta) \; . \end{aligned}$$

The first inequality is by Claims 3.8 and 3.9. The second inequality is by the assumption in the statement of Lemma 3.7. Now apply Lemma 3.5 with $\beta = \mathsf{Dev}(B, T; \delta)$. That gives us Equation (6) and concludes the proof of Lemma 3.7. ▌

Proof of Claim 3.9: Let $(b_1, \ldots, b_q) \in G$. Let $i \in [q]$ and let $a \in R$ be arbitrary. We will show that

$$| L_i(b_1, \ldots, b_{i-1}, a) | \ \leq \ -\ln(1-\delta) \ . \tag{7}$$

Since $a \in R$ was arbitrary it follows that

$$\Pr_{(a_1,\ldots,a_q)\leftarrow\mu} [\, | L_i(a_1, \ldots, a_i) | > -\ln(1-\delta) \ \mid \ (a_1, \ldots, a_{i-1}) = (b_1, \ldots, b_{i-1}) \,] \ = \ 0 \ .$$

Note $\ln(1-\delta) \geq -\delta - \delta^2/2$, and this is at least $-3\delta/2$ since $\delta < 1$. So $-\ln(1-\delta) \leq 3\delta/2 = \delta_2$, which proves the claim.

It remains to establish Equation (7). The definition of G tells us that for every $i \in [q]$ and every $a \in R$ we have

$$| T(b_1, \ldots, b_{i-1}; a) - B(b_1, \ldots, b_{i-1}; a) | \leq \delta \cdot B(b_1, \ldots, b_{i-1}; a) \ .$$

Since $(b_1, \ldots, b_q) \in G$ we also know that $B(b_1, \ldots, b_{i-1}; a) \neq 0$. So we can divide to get

$$\left| \frac{T(b_1, \ldots, b_{i-1}; a)}{B(b_1, \ldots, b_{i-1}; a)} - 1 \right| \leq \delta \ ,$$

or, equivalently,

$$1 - \delta \ \leq \ \frac{T(b_1, \ldots, b_{i-1}; a)}{B(b_1, \ldots, b_{i-1}; a)} \ \leq \ 1 + \delta \ .$$

Taking logs we get

$$\ln(1-\delta) \ \leq \ \ln \frac{T(b_1, \ldots, b_{i-1}; a)}{B(b_1, \ldots, b_{i-1}; a)} \ \leq \ \ln(1+\delta) \ .$$

The middle term is $L_i(b_1, \ldots, b_{i-1}, a)$. Also $|\ln(1-\delta)| = -\ln(1-\delta) \geq \ln(1+\delta)$ so upon taking absolute values we get Equation (7) as desired. This concludes the proof of Claim 3.9. ▌

Proof of Claim 3.8: Let $(b_1, \ldots, b_q) \in G$ and let $i \in [q]$. To establish Equation (4) we will use Lemma 3.6. Let $S = R$. For $x \in R$ let

$$\begin{aligned} \mu_1(x) &= B(b_1, \ldots, b_{i-1}; x) \\ \mu_2(x) &= T(b_1, \ldots, b_{i-1}; x) \ . \end{aligned}$$

To apply Lemma 3.6 we need to check that

$$\forall x \in S \ : \ | \mu_2(x) - \mu_1(x) | \leq \delta \cdot \mu_1(x) \ . \tag{8}$$

However since $(b_1, \ldots, b_q) \in G$ we know that $\mathsf{Diff}_{B,T}(b_1, \ldots, b_{i-1}; \delta) \leq 0$. From the definition of $\mathsf{Diff}_{B,T}(b_1, \ldots, b_{i-1}; \delta)$ we directly get Equation (8). Now Equation (4) follows from the conclusion of Lemma 3.6. ▌

4 PRP to PRF transforms

Definitions. We follow the concrete security treatment of [BKrR] and in particular use their notation. If F is any family of functions we let $\mathbf{InSec}^{\mathrm{prf}}_F(q, t)$ be the insecurity of F as a PRF under q queries and time t. This is the maximum possible value of the advantage $\mathsf{Adv}(A, F, \mathcal{R})$, where $\mathcal{R}$ is the family of all functions mapping $\mathrm{Dom}(F)$ to $\mathrm{Range}(F)$, the maximum being taken over adversaries making only q queries to the given oracle and running in time at most t. Similarly if P is any family of

permutations we let $\mathbf{InSec}_P^{\mathrm{prf}}(q,t)$ be the insecurity of F as a PRP under q queries and time t. This is the maximum possible value of the advantage $\mathsf{Adv}(A,F,\mathcal{P})$, where $\mathcal{P}$ is the family of all permutations over $\mathrm{Dom}(P) = \mathrm{Range}(P)$, the maximum being taken over adversaries making only q queries to the given oracle and running in time at most t.

GIVEN. Let P: $\{0,1\}^k \times \{0,1\}^n \to \{0,1\}^n$ be a PRP with key length k and block length n. We want to build from it a PRF. We want to upper bound the insecurity of the PRF in terms of the given insecurity of P.

4.1 Sum construction

This is our new proposal, which is very efficient and achieves very high security. The simplest version is to define S: $\{0,1\}^{2k} \times \{0,1\}^n \to \{0,1\}^n$ by

$$S_{k1,k2}(x) \;=\; P_{k1}(x) \oplus P_{k2}(x)\;. \tag{9}$$

Namely the key consists of two keys for the PRP, which specify two individual permutations, whose results are XORed together. A variant on this construction that uses just one key is S: $\{0,1\}^k \times \{0,1\}^{n-1} \to \{0,1\}^n$ defined by

$$S_k(x) \;=\; P_k(x \,\|\, 0) \oplus P_k(x \,\|\, 1)\;. \tag{10}$$

Namely the input x is $n-1$ bits long. Form from it the two n bit inputs $x \,\|\, 0$ and $x \,\|\, 1$, feed them to the single permutation P_k involved, and XOR the results. Since the analysis of the second is a little more involved than that of the first, and it is also more practical, we concentrate on analyzing the second.

Theorem 4.1 Let P: $\{0,1\}^k \times \{0,1\}^n \to \{0,1\}^n$ be a PRP and define S as per Equation (10). Let $N = 2^n$. Then for any $0 < q, t < N/4$

$$\mathbf{InSec}_S^{\mathrm{prf}}(q,t) \le \mathbf{InSec}_P^{\mathrm{prp}}(2q, t + O(q(n+k))) + D(n,q)\;,$$

where

$$D(n,q) \;=\; \frac{q}{N} + O(n) \cdot \frac{q^{3/2}}{N^{3/2}}\;.$$

This bound is very good, in particular superior to those of any of the known constructions [BKrR, HWKS]. Roughly it says the insecurity drops as $q/2^n$ modulo the insecurity of the given PRP.

PROOF SKETCH. We concentrate on the information theoretic case; going from there to the computational statement of the theorem is a standard simulation argument. Thus we consider the construction $g(x) = f(x \,\|\, 0) \oplus f(x \,\|\, 1)$ where f is a truly random permutation, and want to estimate how close to a truly random function is g. We use the framework of Section 3.1 to reduce this question to one for which we can apply our main theorem. Namely, define the distributions B, T as there, with FB being the family of random functions, meaning the family of all functions of n bits to $n-1$ bits, and FT being the family given by our construction when f is chosen as a random permutation of n bits to n bits. To apply Theorem 3.2, we must compute a bound on $\mathsf{Dev}(B,T;\delta)$ for an appropriate value of δ. Thus our task reduces to bounding the ratio of probabilities that is implicit in the definition of $\mathsf{Dev}(B,T;\delta)$.

We give only a very brief summary of how this works. First, we give to the adversary even more information that it would normally receive, and show the bound holds even then. Namely, rather than the XOR $f(x \,\|\, 0) \oplus f(x \,\|\, 1)$, the values $f(x \,\|\, 0)$ and $f(x \,\|\, 1)$ are individually given. The advantage of the adversary can only increase. We also remove the value 0 from the range of the functions. This can change the advantage by at most q/N and accounts for the q/N term in the bound. (Note that the construction of Equation (10) can never output the value 0^n. So the adversary can always get an

advantage of q/N. On the other hand comparing random functions to random functions with range $\{0,1\}^n - \{0^n\}$ using the traditional analysis method of conditioning on some bad event (in this case, outputting 0^n) shows that q/N is also an upper bound on the statistical difference.) The definitions of the distributions B, T are adapted to take these changes into account. In particular B corresponds to the experiment of running the adversary with a random function from the class of functions with this restricted range. We now apply Theorem 3.2 to this setting. We will show that it yields as bound the second term in the expression $D(n,q)$ in the theorem.

In estimating the ratio of conditional probabilities, define S_i to be the set of all elements of the form $f(x_j \,\|\, b)$ for $j = 1, \ldots, i$ and $b \in \{0,1\}$. Note for f a random permutation, S_i is a random set of $2i$ distinct elements, no matter what is the adversary strategy. Let $v \in \{0,1\}^n - \{0^n\}$ and let I be the number of pairs $\{u, u \oplus v\}$ which are both in S_i. Let $E = 2q^2/N$, which approximates the expected size of I. Let $D = I - E$. Then the number of ways that we could output v is $N/2 - 2i + I$. A calculation then shows that the ratio between the probability of outputting v under T and that under B is $1 + O(D/N)$. Using Chernoff bounds, we obtain that with probability $1 - 1/N$, for every i and every v, we have $|D| \leq \max(n^2, O(\sqrt{nE})) = O(n^2 + \sqrt{n}q/\sqrt{N})$. This gives a value of $\delta = O(nq/N^{3/2})$ if $q > N^{1/2}n^2$ and $\delta = n^2/N^{3/2}$ otherwise. Theorem 4.1 follows from Corollary 3.2 with this value of δ.

4.2 Truncation construction

The truncation construction [HWKS] associated to each integer $1 \leq m \leq n$ a family T^m: $\{0,1\}^k \times \{0,1\}^n \to \{0,1\}^m$ defined as follows:

$$T_k^l(x) \;=\; [P_k(x)]_{1\ldots m} \,. \tag{11}$$

Namely, drop the last $n - m$ bits of the output of P_k and leave the rest intact.

(**Note:** In comparing with [HWKS, Section 2.4] note that the notation is different. There, m is the number of bits dropped, for us it is the number of bits kept. Furthermore they drop high order bits, not low order, but that makes no difference.)

Theorem 4.2 Let P: $\{0,1\}^k \times \{0,1\}^n \to \{0,1\}^n$ be a PRP and define T^m as per Equation (11). Let $N = 2^n$ and $M = 2^m$. Then for any $0 < q, t < N/2$

$$\mathbf{InSec}_S^{\mathrm{prf}}(q,t) \leq \mathbf{InSec}_P^{\mathrm{prp}}(q, t + O(q(n+k))) + D(n,q) \,,$$

where

$$D(n,q) \;=\; \begin{cases} O(n) \cdot \dfrac{M \cdot \sqrt{q}}{N} & \text{if } 0 < q \leq M \\ \\ O(n) \cdot \dfrac{q \cdot \sqrt{M}}{N} & \text{if } M < q < N/\sqrt{M} \end{cases}$$

Our bounds are better than those of [HWKS] in the region $q > M^{2/3}$. The best overall bound is obtained by combining the results of our paper with theirs. We omit the proof of Theorem 4.2 from this abstract. It uses ideas similar to the proof of Theorem 4.1. Briefly, the probability of an output v depends on the number of times its been output before. We bound this using Chernoff bounds. For $q > M$ this reasoning can also yield an attack that basically matches our upper bound. The adversary counts for each output v whether it has previously occurred more frequently or less frequently than expected. If the majority of times it has occurred less frequently, the adversary predicts that the function is from the truncation construction.

References

[AV] W. Aiello, and R. Venkatesan. Foiling birthday attacks in length-doubling transformations. *Advances in Cryptology – Eurocrypt* 96 *Proceedings*, Lecture Notes in Computer Science Vol. 1070, U. Maurer ed., Springer-Verlag, 1996.

[BCK] M. Bellare, R. Canetti and H. Krawczyk. seudorandom functions revisted: the cascade construction and its concrete security. *Proceedings of the* 37th *Symposium on Foundations of Computer Science*, IEEE, 1996.

[BGK] M. Bellare, O. Goldreich and H. Krawczyk. Beyond the birthday barrier, without counters. *Advances in Cryptology – Crypto* 99 *Proceedings*, Lecture Notes in Computer Science Vol. ??, M. Weiner ed., Springer-Verlag, 1999.

[BKrR] M. Bellare, T. Krovetz and P. Rogaway. Luby-Rackoff backwards: Increasing security by making block ciphers non-invertible. *Advances in Cryptology – Eurocrypt* 97 *Proceedings*, Lecture Notes in Computer Science Vol. 1233, W. Fumy ed., Springer-Verlag, 1997.

[GGM] O. Goldreich, S. Goldwasser and S. Micali, How to construct random functions. *Journal of the ACM*, Vol. 33, No. 4, 1986, pp. 792–807.

[GILVZ] O. Goldreich, R. Impagliazzo, L. Levin, R. Venkatesan and D. Zuckerman. Security preserving amplification of hardness. *Proceedings of the* 31st *Symposium on Foundations of Computer Science*, IEEE, 1990.

[GM] S. Goldwasser and S. Micali, Probabilistic encryption. *Journal of Computer and System Sciences*, Vol. 28, 1984, pp. 270–299.

[HWKS] C. Hall, D. Wagner, J. Kelsey and B. Schneier. Building PRFs from PRPs. *Advances in Cryptology – Crypto* 98 *Proceedings*, Lecture Notes in Computer Science Vol. 1462, H. Krawczyk ed., Springer-Verlag, 1998.

[HILL] J. Håstad, R. Impagliazzo, L. Levin and M. Luby. Pseudo-random generation from one-way functions. Manuscript. Earlier versions in STOC 89 and STOC 90.

[LR] M. Luby and C. Rackoff. How to construct pseudorandom permutations from pseudorandom functions. *SIAM J. Comput*, Vol. 17, No. 2, April 1988.

[MR] R. Motwani and P. Raghavan, Randomized algorithms, Cambridge University Press, 1995.

[Ro] J. Rompel. One-way functions are necessary and sufficient for secure signatures. *Proceedings of the* 22nd *Annual Symposium on Theory of Computing*, ACM, 1990.

A Proofs of Probabilistic Lemmas

We prove the three lemmas of Section 3.3.

A.1 Proof of Lemma 3.3

Let

$$\begin{aligned} G &= \{\, x \in S : |\, T(x) - B(x) \,| \leq \alpha \cdot B(x) \,\} \\ N &= \{\, x \in S : |\, T(x) - B(x) \,| > \alpha \cdot B(x) \,\}\,. \end{aligned}$$

Claim A.1 $\Pr_{x \leftarrow B}[\, x \in G \,] \geq 1 - \gamma$.

Proof: When x is drawn randomly according to B it will be true that $B(x) > 0$, and thus dividing by $B(x)$ is OK. This means that

$$\Pr_{x \leftarrow B}[\, x \in G \,] \;=\; \Pr_{x \leftarrow B}\left[\, 1 - \alpha \leq \frac{T(x)}{B(x)} \leq 1 + \alpha \,\right],$$

and so the claim follows from the assumption in the lemma statement. ∎

Claim A.2 $\Pr_{x \leftarrow T}[\, x \in G \,] \geq 1 - (\alpha + \gamma)$.

Proof: $T(x) \geq (1-\alpha)B(x)$ for all $x \in G$ so

$$\Pr_{x \leftarrow T}[\, x \in G \,] \;=\; \textstyle\sum_{x \in G} T(x) \;\geq\; \sum_{x \in G}(1-\alpha)B(x) \;=\; (1-\alpha)\Pr_{x \leftarrow B}[\, x \in G \,] \,.$$

Now by Claim A.1 the above is at least $(1-\alpha)(1-\gamma) \geq 1 - \alpha - \gamma$, which concludes the proof. ▌

We now complete the proof of the lemma. We have

$$\begin{aligned}
\mathsf{Dist}(B,T) &= \frac{1}{2}\sum_{x \in G} |\, B(x) - T(x) \,| \;+\; \frac{1}{2}\sum_{x \in N} |\, B(x) - T(x) \,| \\
&\leq \frac{1}{2}\sum_{x \in G} \alpha \cdot B(x) \;+\; \frac{1}{2}\sum_{x \in N} (B(x) + T(x)) \\
&\leq \frac{\alpha}{2}\Pr_{x \leftarrow B}[\, x \in G \,] + \frac{1}{2}\Pr_{x \leftarrow B}[\, x \in N \,] + \frac{1}{2}\Pr_{x \leftarrow T}[\, x \in N \,] \\
&\leq \frac{\alpha}{2} \cdot 1 + \frac{1}{2} \cdot \gamma + \frac{1}{2} \cdot (\alpha + \gamma) \\
&= \alpha + \gamma \,.
\end{aligned}$$

Above we used the claims to do the bounding. This concludes the proof of Lemma 3.3.

A.2 Proof of Lemma 3.5

The proof is based on the standard Azuma inequality. The version of the latter stated below is taken from Motwani and Raghavan [MR, Corollary 4.17].

Lemma A.3 [Azuma's inequality] Let μ: $S \to [0,1]$ be a probability distribution over a set S, and let $X_0, X_1, \ldots, X_n$: $S \to \mathbf{R}$ be functions. For $i \in [n]$ and $b \in S$ let

$$S(i,b) \;=\; \{\, a \in S \;:\; (X_0(a), \ldots, X_{i-1}(a)) = (X_0(b), \ldots, X_{i-1}(b)) \,\} \,.$$

Suppose the following two conditions hold–

(1) *Martingale:* For every $i \in [n]$ and every $b \in S$–

$$\mathbf{E}_{a \leftarrow \mu}[\, X_i(a) \;\mid\; a \in S(i,b) \,] \;=\; X_{i-1}(b) \,.$$

(2) *Bounded differences:* $|\, X_i(b) - X_{i-1}(b) \,| \leq c$ for all $i \in [n]$ and all $b \in S$.

Then for any $\lambda > 0$

$$\Pr_{a \leftarrow \mu}\left[\, |\, X_n(a) - X_0(a) \,| \geq \lambda c n^{1/2} \,\right] \leq 2e^{-\lambda^2/2} \,.$$

We now prove Lemma 3.5 assuming Lemma A.3. Let $S = R^n$. Let $i \in [n]$ and $b \in S$. We define the following sets:

$$\begin{aligned}
P(i,b) &= \{\, a \in S \;:\; (a_1, \ldots, a_{i-1}) = (b_1, \ldots, b_{i-1}) \,\} \\
G(i) &= \{\, a \in S \;:\; a \text{ is } (\delta_1, \delta_2, i)\text{-good for } L_1, \ldots, L_i \text{ over } \mu \,\}
\end{aligned}$$

We now define $X_0, \ldots, X_n$ as follows. First we set $X_0(b) = 0$ for all $b \in S$. Then for each $i \in [n]$ and each $b \in S$ set

$$X_i(b) \;=\; \begin{cases} X_{i-1}(b) + L_i(b_1, \ldots, b_i) - \mathbf{E}_{a \leftarrow \mu}[\, L_i(a_1, \ldots, a_i) \;\mid\; a \in P(i,b) \,] & \text{if } b \in G(i) \\ X_{i-1}(b) & \text{otherwise} \,. \end{cases}$$

The three lemmas stated below say that $X_0, \ldots, X_n$ satisfy the conditions necessary to apply Lemma A.3. The first lemma notes that $X_0, \ldots, X_n$ are finite, meaning do not take values $\pm\infty$, even though $L_1, \ldots, L_n$ could take the values $\pm\infty$.

Lemma A.4 $X_0, \ldots, X_n$ take values in $\mathbf{R}$.

The next lemma verifies that $X_0, \ldots, X_n$ satisfy the martingale condition.

Lemma A.5 For every $i \in [n]$ and every $b \in S$–

$$\mathbf{E}_{a \leftarrow \mu}\left[\, X_i(a) \mid a \in S(i,b) \,\right] \;=\; X_{i-1}(b)\,. \tag{12}$$

The last lemma says the bounded difference condition holds with $c = \delta_1 + \delta_2$–

Lemma A.6 $|\, X_i(b) - X_{i-1}(b) \,| \leq c$ for all $i \in [n]$ and all $b \in S$, where $c = \delta_1 + \delta_2$.

The proofs of these lemmas will be given below. Let us first use them to conclude the proof of Lemma 3.5. We wish to upper bound

$$A \;=\; \mathrm{Pr}_{b \leftarrow \mu}\left[\, |\, \textstyle\sum_{i=1}^n L_i(b_1, \ldots, b_i) \,| \geq \lambda c n^{1/2} + \delta_2 n \,\right].$$

We let

$$A_1 \;=\; \mathrm{Pr}_{b \leftarrow \mu}\left[\, |\, \textstyle\sum_{i=1}^n L_i(b_1, \ldots, b_i) \,| \geq \lambda c n^{1/2} + \delta_2 n \mid b \in G(n) \,\right].$$

We know that $\mathrm{Pr}_{b \leftarrow \mu}[\, b \notin G(n) \,] \leq \beta$ so

$$\begin{aligned} A &\leq A_1 \cdot \mathrm{Pr}_{b \leftarrow \mu}[\, b \in G(n) \,] + \mathrm{Pr}_{b \leftarrow \mu}[\, b \notin G(n) \,] \\ &\leq A_1 \cdot \mathrm{Pr}_{b \leftarrow \mu}[\, b \in G(n) \,] + \beta\,. \end{aligned}$$

To conclude the proof we need to show that $A_1 \cdot \mathrm{Pr}_{b \leftarrow \mu}[\, b \in G(n) \,] \leq 2e^{-\lambda^2/2}$. For $b \in G(n)$ let $M_i(b) = \mathbf{E}_{a \leftarrow \mu}[\, L_i(a_1, \ldots, a_i) \mid a \in P(i,b) \,]$. We know that $|\, M_i(b) \,| \leq \delta_2$ when $b \in G(n)$, so

$$\begin{aligned} A_1 &\leq \mathrm{Pr}_{b \leftarrow \mu}\left[\, |\, \textstyle\sum_{i=1}^n L_i(b_1, \ldots, b_i) \,| \geq \lambda c n^{1/2} + |\, \sum_{i=1}^n M_i(b) \,| \mid b \in G(n) \,\right] \\ &\leq \mathrm{Pr}_{b \leftarrow \mu}\left[\, |\, \textstyle\sum_{i=1}^n L_i(b_1, \ldots, b_i) - M_i(b) \,| \geq \lambda c n^{1/2} \mid b \in G(n) \,\right]. \end{aligned}$$

Now notice that for $b \in G(n)$

$$X_n(b) - X_0(b) \;=\; \sum_{i=1}^n X_i(b) - X_{i-1}(b) \;=\; \sum_{i=1}^n L_i(b_1, \ldots, b_i) - M_i(b)\,.$$

So we get

$$\begin{aligned} A_1 \cdot \mathrm{Pr}_{b \leftarrow \mu}[\, b \in G(n) \,] &\leq \mathrm{Pr}_{b \leftarrow \mu}\left[\, |\, X_n(b) - X_0(b) \,| \geq \lambda c n^{1/2} \mid b \in G(n) \,\right] \cdot \mathrm{Pr}_{b \leftarrow \mu}[\, b \in G(n) \,] \\ &\leq \mathrm{Pr}_{b \leftarrow \mu}\left[\, |\, X_n(b) - X_0(b) \,| \geq \lambda c n^{1/2} \,\right] \\ &\leq 2e^{-\lambda^2/2}\,. \end{aligned}$$

This concludes the proof of Lemma 3.5. It remains to establish the above lemmas.

Proof of Lemma A.4: Suppose $b \in G(i)$, meaning b is (δ_1, δ_2, i)-good for $L_1, \ldots, L_i$ over μ. Then

$$\mathbf{E}_{a \leftarrow \mu}\left[\, L_i(a_1, \ldots, a_i) \mid a \in P(i,b) \,\right] < \infty \quad \text{and} \quad L_i(b_1, \ldots, b_i) < \infty\,.$$

The first is true by the first condition in Definition 3.4. The second is true because otherwise we would contradict the second condition in Definition 3.4. The claim that $X_0, \ldots, X_n < \infty$ follows from the definition of the random variables; a full proof would be by induction on i. ▌

Notice that $P(i,b) \subseteq S(i,b)$ for any $b \in S$. We will prove Lemma A.5 by showing something stronger:

Claim A.7 For any $i \in [n]$ and any $b \in S$ we have

$$\mathbf{E}_{a \leftarrow \mu}[\, X_i(a) \mid a \in P(i,b)\,] \;=\; X_{i-1}(b)\,. \tag{13}$$

We now prove Lemma A.5 given this claim, and will then prove the claim.

Proof of Lemma A.5: Let $i \in [n]$ and $b \in S$. We break up $S(i,b)$ as $S(i,b) = P(i,z_1) \cup \cdots \cup P(i,z_m)$ where $z_1, \ldots, z_m \in S(i,b)$ are some values such that the sets $P(i,z_1), \ldots, P(i,z_m)$ are mutually disjoint. Then

$$\begin{aligned}
&\mathbf{E}_{a \leftarrow \mu}[\, X_i(a) \mid a \in S(i,b)\,] \\
&\quad = \sum_{j=1}^{m} \mathbf{E}_{a \leftarrow \mu}[\, X_i(a) \mid a \in P(i,z_j)\,] \cdot \Pr_{x \leftarrow \mu}[\, x \in P(i,z_j) \mid x \in S(i,b)\,]\,.
\end{aligned}$$

By Claim A.7 the conditional expectation in the above equation equals $X_{i-1}(z_j)$. However $z_j \in S(i,b)$ so $X_{i-1}(z_j) = X_{i-1}(b)$, so the above equals

$$\sum_{z_j \in S(i,b)} X_{i-1}(b) \cdot \Pr_{x \leftarrow \mu}[\, x \in P(i,z_j) \mid x \in S(i,b)\,]\,.$$

However we can now factor out $X_{i-1}(b)$, and the remaining sum equals one, so we have established Equation (12) as desired. ▌

Proof of Claim A.7: We begin by defining some sets. For any $i \in [n]$ and any $a \in S$ we set

$$\begin{aligned}
N(i) &= S - G(i) \\
G(i,a) &= P(i,a) \cap G(i) \\
N(i,a) &= P(i,a) \cap N(i)
\end{aligned}$$

We need some claims about these sets.

Claim 1. If $a \in G(i)$ then $P(i,a) = G(i,a)$, and if $a \in N(i)$ then $P(i,a) = N(i,a)$.

Proof. We observe that whether some $y \in S$ is in $G(i)$ depends only on the values $y_1, \ldots, y_{i-1}$, by definition of being (δ_1, δ_2, i)-good for $L_1, \ldots, L_i$ over μ, and the same is true for $N(i)$. The claim follows. □

Claim 2. If $a \in G(i,b)$ then $P(i,a) = G(i,b)$.

Proof. By assumption $a \in P(i,b)$, so $P(i,a) = P(i,b)$. Then $G(i,b) = G(i) \cap P(i,b) = G(i) \cap P(i,a) = G(i,a)$. □

Claim 3. If $a \in N(i,b)$ then $X_i(a) = X_{i-1}(b)$.

Proof. By assumption $a \in N(i)$ so $X_i(a) = X_{i-1}(a)$ by definition of X_i. On the other hand $a \in P(i,b)$, whence $a \in S(i,b)$. But if $a \in S(i,b)$ then $X_{i-1}(a) = X_{i-1}(b)$. So $X_i(a) = X_{i-1}(b)$ as desired. □

We establish Equation (13) by considering separately the case of $b \in G(i)$ and $b \in N(i)$.

First suppose $b \in N(i)$. Then

$$\begin{aligned}
\mathbf{E}_{a \leftarrow \mu}[\, X_i(a) \mid a \in P(i,b)\,] &= \mathbf{E}_{a \leftarrow \mu}[\, X_i(a) \mid a \in N(i,b)\,] \\
&= \mathbf{E}_{a \leftarrow \mu}[\, X_{i-1}(b) \mid a \in N(i,b)\,] \\
&= X_{i-1}(b)\,,
\end{aligned}$$

as desired. Here we first used Claim 1 to say that $P(i,b) = N(i,b)$ and then used Claim 3 to say that $X_i(a) = X_{i-1}(b)$ in the conditional expectation.

Next we establish Equation (13) in the case where $b \in G(i)$. By Claim 1 we have $P(i,b) = G(i,b)$ so we wish to show that

$$\mathbf{E}_{a \leftarrow \mu}[\, X_i(a) \mid a \in G(i,b)\,] \;=\; X_{i-1}(b)\,. \tag{14}$$

By linearity of expectation and the definition of X_i we have

$$\mathbf{E}_{a \leftarrow \mu}[\, X_i(a) \mid a \in G(i,b)\,] \;=\; A + B - C$$

where

$$\begin{aligned} A &= \mathbf{E}_{a \leftarrow \mu}[\, L_i(a_1, \ldots, a_i) \mid a \in G(i,b)\,] \\ B &= \mathbf{E}_{a \leftarrow \mu}[\, X_{i-1}(a) \mid a \in G(i,b)\,] \\ C &= \mathbf{E}_{a \leftarrow \mu}[\, \mathbf{E}_{x \leftarrow \mu}[\, L_i(x_1, \ldots, x_i) \mid x \in P(i,a)\,] \mid a \in G(i,b)\,] \end{aligned}$$

We first claim that $B = X_{i-1}(b)$. This is true because in the expectation we always have $a \in G(i,b) \subseteq P(i,b) \subseteq S(i,b)$, so $X_{i-1}(a) = X_{i-1}(b)$. Now we claim that $C = A$. Given this and the above, the proof of Equation (14) will be complete.

Note that in the expression for C we have $a \in G(i,b)$. By Claim 2 we have $P(i,a) = G(i,b)$ and hence

$$\begin{aligned} C &= \mathbf{E}_{a \leftarrow \mu}[\, \mathbf{E}_{x \leftarrow \mu}[\, L_i(x_1, \ldots, x_i) \mid x \in G(i,b)\,] \mid a \in G(i,b)\,] \\ &= \mathbf{E}_{x \leftarrow \mu}[\, L_i(x_1, \ldots, x_i) \mid x \in G(i,b)\,] \\ &= A\,. \end{aligned}$$

This concludes the proof of Claim A.7. ▌

Proof of Lemma A.6: Consider separately the cases of $b \in G(i)$ and $b \in N(i)$. If $b \in N(i)$ then $X_i(b) = X_{i-1}(b)$ so $|\, X_i(b) - X_{i-1}(b)\,| = 0$. If $b \in G(i)$ then

$$\begin{aligned} |\, X_i(b) - X_{i-1}(b)\,| &= |\, L_i(b_1, \ldots, b_i) - \mathbf{E}_{a \leftarrow \mu}[\, L_i(a_1, \ldots, a_i) \mid a \in P(i,b)\,]\,| \\ &\le |\, L_i(b_1, \ldots, b_i)\,| + |\, \mathbf{E}_{a \leftarrow \mu}[\, L_i(a_1, \ldots, a_i) \mid a \in P(i,b)\,]\,| \\ &\le \delta_2 + \delta_1\,. \end{aligned}$$

The last inequality is true because b is (δ_1, δ_2, i)-good for $L_1, \ldots, L_i$ over μ. ▌

A.3 Proof of Lemma 3.6

Let $S_1 = \{\, x \in S \,:\, \mu_1(x) > 0\,\}$ and $S_2 = \{\, x \in S \,:\, \mu_2(x) > 0\,\}$.

Claim A.8 $S_1 = S_2$.

Proof: It suffices to show that for any $x \in S$ it is the case that $\mu_1(x) = 0$ iff $\mu_2(x) = 0$. So first suppose $\mu_1(x) = 0$. From the condition $|\mu_2(x) - \mu_1(x)| \le \delta\mu_1(x)$ we directly get $\mu_2(x) = 0$. Now suppose $\mu_2(x) = 0$. From the condition $|\mu_2(x) - \mu_1(x)| \le \delta\mu_1(x)$ we get $|\mu_1(x)| \le \delta\mu_1(x)$. However we assumed $\delta < 1$ so this possible only if $\mu_1(x) = 0$. ▌

Given the above we let $P = S_1 = S_2$ stand for the common support of the two distributions. Now

$$\begin{aligned} A &\stackrel{\text{def}}{=} \mathop{\mathbf{E}}_{x \leftarrow \mu_1}\left[\, \ln \frac{\mu_2(x)}{\mu_1(x)}\,\right] \\ &= \sum_{x \in P} \mu_1(x) \cdot \ln \frac{\mu_2(x)}{\mu_1(x)} \\ &= \sum_{x \in P} \mu_1(x) \cdot \ln\left(1 + \frac{\mu_2(x) - \mu_1(x)}{\mu_1(x)}\right)\,. \end{aligned}$$

To upper bound A we use the inequality $\ln(1+y) \leq y$ and get

$$\begin{aligned} A &\leq \sum_{x \in P} \mu_1(x) \cdot \frac{\mu_2(x) - \mu_1(x)}{\mu_1(x)} \\ &= \sum_{x \in P} \mu_2(x) - \mu_1(x) \\ &= 0 \,, \end{aligned}$$

the last equality being true by Claim A.8. To lower bound A we use the inequality $\ln(1+y) \geq y - y^2/2$. Setting $\delta_x = \mu_2(x) - \mu_1(x)$ we get

$$\begin{aligned} A &\geq \sum_{x \in P} \mu_1(x) \cdot \left[\frac{\delta_x}{\mu_1(x)} - \frac{\delta_x^2}{2\mu_1(x)^2} \right] \\ &= \sum_{x \in P} \delta_x - \sum_{x \in P} \frac{\delta_x^2}{2\mu_1(x)} \\ &= -\frac{1}{2} \cdot \sum_{x \in P} \frac{\delta_x^2}{\mu_1(x)} \\ &\geq -\frac{1}{2} \cdot \sum_{x \in P} \frac{\delta^2 \mu_1(x)^2}{\mu_1(x)} \\ &= -\frac{\delta^2}{2} \cdot \sum_{x \in P} \mu_1(x) \\ &= -\frac{\delta^2}{2} \,. \end{aligned}$$

That is, we have shown $-\delta^2/2 \leq A \leq 0$ whence $|A| \leq \delta^2/2$.