NIST Interagency Report
NIST IR 8427

Discussion on the Full Entropy Assumption of the SP 800-90 Series

Darryl Buller
Aaron Kaufer
Allen Roginsky
Meltem Sönmez Turan

This publication is available free of charge from:
https://doi.org/10.6028/NIST.IR.8427.ipd

NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY
U.S. DEPARTMENT OF COMMERCE

NIST Interagency Report
NIST IR 8427

Discussion on the Full Entropy Assumption of the SP 800-90 Series

Darryl Buller
Aaron Kaufer
Cybersecurity Directorate
National Security Agency

Allen Roginsky
Meltem Sönmez Turan
Computer Security Division
Information Technology Laboratory

This publication is available free of charge from:
https://doi.org/10.6028/NIST.IR.8427

April 2023

U.S. Department of Commerce
Gina M. Raimondo, Secretary

National Institute of Standards and Technology
Laurie E. Locascio, NIST Director and Under Secretary of Commerce for Standards and Technology

NIST Technical Series Policies

Publication History

Approved by the NIST Editorial Review Board on 2023-03-30

How to Cite this NIST Technical Series Publication:

Buller D, Kaufer A, Roginsky A, Sönmez Turan M (2023) Discussion on the Full Entropy Assumption of the SP 800-90 Series. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Interagency or Internal Report (IR) NIST IR 8427. https://doi.org/10.6028/NIST.IR.8427

Author ORCID iDs

Roginsky, Allen: 0000-0003-2684-6736
Meltem Sönmez Turan: 0000-0002-1950-7130

Submit Comments

rbg_comments@nist.gov

All comments are subject to release under the Freedom of Information Act (FOIA).

Abstract

The NIST Special Publication (SP) 800-90 series supports the generation of high-quality random bits for cryptographic and non-cryptographic use. The security strength of a random number generator depends on the *unpredictability* of its outputs. This unpredictability can be measured in terms of entropy, which the NIST SP 800-90 series measures using *min-entropy*. A full-entropy bitstring has an amount of entropy equal to its length. Full-entropy bitstrings are important for cryptographic applications, as these bitstrings have ideal randomness properties and may be used for any cryptographic purpose. Due to the difficulty of generating and testing full-entropy bitstrings, the SP 800-90 series assumes that a bitstring has full entropy if the amount of entropy per bit is at least $1 - \varepsilon$, where ε is at most 2^{-32}. This report provides a justification for the selection of this value of ε.

Keywords

entropy; min-entropy; random number generation.

Reports on Computer Systems Technology

The Information Technology Laboratory (ITL) at the National Institute of Standards and Technology (NIST) promotes the U.S. economy and public welfare by providing technical leadership for the Nation's measurement and standards infrastructure. ITL develops tests, test methods, reference data, proof of concept implementations, and technical analyses to advance the development and productive use of information technology. ITL's responsibilities include the development of management, administrative, technical, and physical standards and guidelines for the cost-effective security and privacy of other than national security-related information in federal information systems.

Table of Contents

1. Introduction

The NIST Special Publication (SP) 800-90 series [1][2][3] supports the generation of high-quality random bits for cryptographic and non-cryptographic use. The security strength of a random number generator depends on the *unpredictability* of its outputs. This unpredictability can be measured in terms of entropy, which the NIST SP 800-90 series measures using *min-entropy*. The SP 800-90 series refers to certain bitstrings as having *full entropy*, meaning that they have an amount of entropy equal to their length. Such bitstrings are important for cryptographic applications since they have the highest level of unpredictability and may therefore be used for any cryptographic purpose. The SP 800-90 series assumes that a bitstring has full entropy if the amount of entropy per bit is at least $1 - \varepsilon$, where ε is at most 2^{-32}. This report provides a justification for the selection of this value of ε. Due to the difficulty of ensuring full entropy through the analysis of random bit generators and sample data sequences, this is accomplished by using an alternative definition of full entropy and calculating the resulting entropy level when this definition is satisfied.

2. Problem Statement

The SP 800-90 series uses a definition of full entropy that prescribes a numerical threshold on the entropy per bit (at least $1 - \varepsilon$, where ε is at most 2^{-32}). However, although this is an intuitive way to define full entropy, it is generally not possible to ensure sufficient entropy within this tolerance by analyzing a Random Bit Generator (RBG) or estimating entropy from a data sample. It is, therefore, necessary to use a different definition from which a practical approach to ensuring full entropy can be derived. The report begins by defining full entropy in terms of a hypothetical distinguishing game. The report then presents and proves two results following from this definition. First, it is shown how output satisfying this definition can be generated using a vetted conditioning function (see [2]) acting on data having an entropy level that meets or exceeds a certain value. Second, it is shown that the entropy level of output produced by such a process satisfies the full entropy threshold used in the SP 800-90 series, thereby demonstrating a connection between these two definitions.

3. Full Entropy Definition

Full entropy will be defined as follows. Consider a distinguishing game where an adversary attempts to distinguish between two cases – REAL and IDEAL. Assume that the adversary is provided with W n-bit outputs $b_{1,}, b_2, \dots, b_W$. In the REAL case, each output is generated by a cryptographic conditioning function operating on an output from a real-life entropy source. In the IDEAL case, the outputs are generated by an idealized randomness source that produces independent n-bit outputs with each output value having a probability of 2^{-n}. One of the two cases is chosen at random with each being equally likely. Outputs of length n bits generated in the REAL case are defined as having *full entropy* with respect to W and δ (where $\delta > 0$) if a computationally unlimited adversary cannot correctly distinguish between the REAL and IDEAL cases with probability higher than $\frac{1}{2} + \delta$.

4. Claims

4.1. Claim 1

Output from a vetted conditioning function with sufficient input entropy satisfies the full entropy definition.

Suppose that values for the W and δ parameters used in the definition of full entropy are given and that a vetted conditioning function generates an n-bit output by processing an entropy sequence having min-entropy H. Then if $H \geq n + \log_2\left(\frac{W}{\delta^2}\right) - log_2\pi - 3$, the n-bit output satisfies the definition of full entropy with respect to W and δ. The proof of this claim is given in A.1.

From the above inequality, $H - n$ represents the minimum required amount by which the input entropy must exceed n in order to ensure full entropy. The following table shows this minimum value for various values of W and δ.

Table 1. Minimum value of additional entropy $\boldsymbol{H - n}$ required for various values of $\boldsymbol{W}$ and $\boldsymbol{\delta}$

	$\boldsymbol{\delta}$						
$\boldsymbol{W}$	2^{-20}	2^{-18}	2^{-16}	2^{-14}	2^{-12}	2^{-10}	2^{-8}
2^{32}	67.3	63.3	59.3	55.3	51.3	47.3	43.3
2^{40}	75.3	71.3	67.3	63.3	59.3	55.3	51.3
2^{48}	83.3	79.3	75.3	71.3	67.3	63.3	59.3
2^{56}	91.3	87.3	83.3	79.3	75.3	71.3	67.3

It is assumed in SP 800-90C that there is an upper bound of 2^{64} bits on the amount of output that an adversary attempting a distinguishing attack can obtain. Consider the combination $W = 2^{48}$ and $\delta = 2^{-10}$. Given $W = 2^{48}$ n-bit RBG outputs, each output can be up to $2^{16} = 65536$ bits long without exceeding the 2^{64} data-quantity bound. Note that 10000 RBGs –each producing 1000 outputs per second– would require nearly a year to produce $W = 2^{48}$ outputs. According to the table above, an adversary who obtains $W = 2^{48}$ n-bit outputs has a distinguishing probability no greater than $\frac{1}{2} + \delta = \frac{1}{2} + 2^{-10} \cong 0.501$ when H, the conditioning function input min-entropy for each n-bit output, is at least $n + 63.3$. This minimum value, rounded up to $n + 64$, is used in this document and in SP 800-90C as the condition for satisfying the full entropy definition.

4.2. Claim 2

Outputs that satisfy the definition of full entropy in Section 3 also satisfy the full entropy threshold in SP 800-90C.

Suppose that $W = 2^{48}$ and $\delta = 2^{-10}$. According to Claim 1, if the input min-entropy satisfies $H \geq n + 64$, then the conditioning function output satisfies the definition of full entropy. In this case, it can also be shown that the average per-bit min-entropy of the resulting n-bit output is at least $1 - 2^{-32}$. The proof of this claim is given in A.3. It is interesting to note that if a process generates outputs satisfying the definition of full entropy based on the distinguishing game, the outputs also satisfy the more intuitive definition in terms of an actual numerical entropy threshold. The resulting value $\varepsilon = 2^{-32}$ is the value given in SP 800-90C as the threshold for full entropy.

5. Assumptions

The proofs of the two claims made in this paper require some assumptions. These assumptions are presented and explained below.

5.1. Assumption 1: The Test Statistic is Normally Distributed

The proof of Claim 1 uses a test statistic X computed from data values generated in either the REAL or the IDEAL case. It will be assumed that X is approximately normally distributed. This statistic is of the form $X = \sum_{i=1}^{W} x_i$, where the variables x_i are independent and identically distributed, and depend on the observed values b_i. If μ_x and ${\sigma_x}^2$ are the true mean and variance of the individual variables x_i, respectively (μ_x and ${\sigma_x}^2$ are derived in A.1), then $\frac{\frac{X}{W}-\mu_x}{\left({\sigma_x}/{\sqrt{W}}\right)} = \frac{\frac{\sum_{i=1}^{W} x_i}{W}-\mu_x}{\left({\sigma_x}/{\sqrt{W}}\right)}$ is approximately normally distributed by the Central Limit Theorem when W is large (we assume $W = 2^{48}$). Since X can be obtained from $\frac{\frac{X}{W}-\mu_x}{\left({\sigma_x}/{\sqrt{W}}\right)}$ by a linear transformation, X is also approximately normally distributed.

5.2. Assumption 2: The Conditioning Function Output Probabilities are Random Variables

The discussion of full entropy in this paper is in the context of the SP 800-90 series, where bit sequences having full entropy are to be generated by processing entropy source output sequences (that generally do not have full entropy) with a vetted conditioning function or a Deterministic Random Bit Generator (DRBG). The assumption below is made in that context.

Suppose that a conditioning function generates an n-bit output by processing a sequence from an entropy source. Consider the probability p_j for the j^{th} possible output value from the conditioning function. These probabilities are determined by the interaction between the specific conditioning function used and the space of possible inputs to that function. In the distinguishing game used in the full entropy definition, the adversary has complete knowledge of the conditioning function and its input space, and—being computationally unlimited—can determine the probabilities p_j. For the purposes of this paper, determining the values of these probabilities is infeasible. However, it is useful to consider the p_j as random variables rather than fixed values and use statistics associated with these random variables to find the probability distribution of the test statistic X. This can be justified as follows. Consider the application of the conditioning function to the input space of entropy source sequences to obtain the conditioning function output. The resulting mapping depends on both the selected input domain and the details of the conditioning function. The domain of this function is determined by the characteristics of the entropy source and the length of the entropy source sequences that are input to the conditioning function. These details effectively select from a large number of possible domains for the function. A consequence of the cryptographic properties of conditioning functions used in RBGs is that there is no simple pattern in the mapping from elements of the domain to elements of the range of the function, and there is no discernible bias in the assignment of outputs to inputs. Therefore, although the conditioning function is deterministic, this report will assume that it is effectively a mapping that assigns an n-bit output to each value in the domain with a uniform distribution so that each possible output value has a probability of 2^{-n} of being the output value assigned to a given input value. (Note that multiple input values can be assigned a given output value.) Under this assumption, since probability p_j is the sum of the probabilities of the inputs mapped to the j^{th} output value, the probability p_j can be treated as a random variable.

5.3. Assumption 3: The Conditioning Function Output Probabilities are Normally Distributed

Consider that p_j is treated as a random variable as explained in Assumption 2. Suppose that there are M possible inputs to the conditioning function, with probabilities $\{q_1, q_2, \dots, q_M\}$. (Note that no assumptions will be made on the input probability distribution.) The output probability p_j can then be written as $p_j = \sum_{i=1}^{M} q_i I_{i,j}$, where $I_{i,j} = 1$ if the conditioning function maps the i^{th} input to the j^{th} output, and $I_{i,j} = 0$ otherwise. Under Assumption 2, the assignment of inputs to outputs is done independently for each input, so for each value of j, the variables $q_i I_{i,j}$ are independent. The value of M is dependent on the characteristics of the entropy source outputs and the bit length of the conditioning function input. However, it is reasonable to assume that M is at least 2^n, the number of possible conditioning function outputs (it would be difficult to achieve a level of output entropy near n otherwise). Given the above statements, p_j is of a similar form to that of X in Assumption 1 and, according to a version of the Central Limit Theorem applying to independent random variables not having the same distribution, can be assumed to be approximately normally distributed.

References

[1] Barker EB, Kelsey JM (2015) *Recommendation for Random Number Generation Using Deterministic Random Bit Generato*rs. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-90Ar1. https://doi.org/10.6028/NIST.SP.800-90Ar1

[2] Sönmez Turan M, Barker EB, Kelsey JM, McKay KA, Baish ML, Boyle M (2018*) Recommendation for the Entropy Sources Used for Random Bit Generati*on. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-90B. https://doi.org/10.6028/NIST.SP.800-90B

[3] Barker EB, Kelsey JM, McKay K, Roginsky A, Sönmez Turan M (2022) *Recommendation for Random Bit Generator (RBG) Constructions.* (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-90C 3pd, Third public draft. https://doi.org/10.6028/NIST.SP.800-90C.3pd

[4] Handbook of Mathematical Functions, with Formulas, Graphs, and Mathematical Tables, Milton Abramowitz and Irene A. Stegun, editors, Dover Publications, New York, 1972.

Appendix A. Proofs

A.1. Proof of Claim 1

Suppose that random output is generated by processing a quantity of entropy data using a conditioning function. The claim is that given values of $W$and δ, it is possible to find a threshold such that if the min-entropy of the input to the conditioning function meets or exceeds that threshold, the conditioning function output will satisfy the definition of full entropy.

Let $B = \{b_{1,}, b_2, \dots, b_W\}$ be the set of observed n-bit outputs, p_j be the probability of the j^{th} possible output from the conditioning function applied to the specified quantity of entropy data, and p_{b_i} be the probability of the i^{th} observed output in the REAL case. Note that each p_{b_i} corresponds to some p_j.

Now consider the likelihood ratio $\frac{Pr[REAL|B]}{Pr[IDEAL|B]}$. Clearly, the adversary will conclude that B was produced by the REAL case if this likelihood ratio is greater than one and by the IDEAL case otherwise (there would be no reason to guess the less likely case). Since the REAL and IDEAL cases are equally likely, this likelihood ratio can be rewritten as $\frac{Pr[B|\text{REAL}]}{Pr[B|\text{IDEAL}]}$ using Bayes Theorem. For ease of computation, compute the base-2 log of the likelihood ratio and denote the resulting statistic as X. The adversary will conclude that B was produced by the REAL case if $X > 0$ and by the IDEAL case otherwise. Then the following is true:

$$\begin{aligned} X &= \log_2\left(\frac{Pr[B|\text{REAL}]}{Pr[B|\text{IDEAL}]}\right) \\ &= \log_2(Pr[B|\text{REAL}]) - \log_2(Pr[B|\text{IDEAL}]) \\ &= \log_2\left(\prod_{i=1}^{W} p_{b_i}\right) - \log_2(2^{-nW}) \\ &= \sum_{i=1}^{W}\left(n + \log_2 p_{b_i}\right) \end{aligned}$$

The statistic X is a random variable that depends on the set B of observed n-bit outputs b_i and the probabilities p_{b_i} of those outputs in the REAL case. To assess the adversary's distinguishing success probability, the probability distribution of X in both the REAL and IDEAL cases is required. Note that X is the sum of W individual random variables $x_i = n + \log_2 p_{b_i}$. These variables, being determined by the generation of independent outputs b_i, are independent and identically distributed. (In the IDEAL case, this is clearly true. In the REAL case, it follows from the assumed properties of the conditioning function.) As noted in Assumption 1, it is assumed that X is approximately normally distributed.

Consider that p_j is treated as a random variable as explained in Assumption 2. Then $E[p_j] = \sum_{i=1}^{M} q_i E[I_{i,j}] = \sum_{i=1}^{M} 2^{-n} q_i = 2^{-n}$. Similarly,

$$\begin{aligned}
VAR[p_j] &= \sum_{i=1}^{M} VAR[q_i I_{i,j}] \\
&= \sum_{i=1}^{M} \left(E\left[(q_i I_{i,j})^2 \right] - \left(E[q_i I_{i,j}] \right)^2 \right) \\
&= \sum_{i=1}^{M} (2^{-n} q_i^2 - 2^{-2n} q_i^2) \\
&= (2^{-n} - 2^{-2n}) \sum_{i=1}^{M} q_i^2
\end{aligned}$$

Now write p_j as $p_j = 2^{-n}(1 + \theta_j)$. Then $\theta_j = 2^n p_j - 1$, so $E[\theta_j] = 2^n E[p_j] - 1 = 0$ and $VAR[\theta_j] = 2^{2n} VAR[p_j] = (2^n - 1) \sum_{i=1}^{M} q_i^2$. Since the input collision entropy $H_2 = -\log_2 \sum_{i=1}^{M} q_i^2$, $VAR[\theta_j] = (2^n - 1) 2^{-H_2}$. Note that since p_j is assumed to be normally distributed as explained in Assumption 3, $\theta_j = 2^n p_j - 1$ is also normally distributed.

The mean and variance of X depend on whether the source is REAL or IDEAL. Let $\mu_R = E[x_i | REAL]$, $\mu_I = E[x_i|\text{IDEAL}]$, $\sigma_R^2 = VAR[x_i|\text{REAL}]$, and $\sigma_I^2 = VAR[x_i|\text{IDEAL}]$.

Now derive μ_R, μ_I, σ_R^2, and σ_I^2. Each of these values is computed by summing over the relevant expression using 2^{-n} or p_j as the probability weighting factors for the IDEAL and REAL cases, respectively. Thus,

$$\begin{aligned}
E[x_i | IDEAL] &= E\left[n + \log_2 p_{b_i} | \text{IDEAL} \right] \\
&= \sum_{j=1}^{2^n} \left(n + \log_2 p_j \right) 2^{-n} \\
&= \sum_{j=1}^{2^n} \left(n + \frac{\ln\left(2^{-n}(1 + \theta_j) \right)}{\ln 2} \right) 2^{-n} \\
&= \sum_{j=1}^{2^n} \frac{\ln(1 + \theta_j)}{\ln 2} 2^{-n}
\end{aligned}$$

The Taylor series for $\ln(1 + \theta_j)$ is $\theta_j - \frac{{\theta_j}^2}{2} + \frac{{\theta_j}^3}{3} - \frac{{\theta_j}^4}{4} + \cdots$. It is shown in A.2 that for cases of interest, $|\theta_j|$ is on the order of 10^{-8} or smaller. For such values of θ_j, $\ln(1 + \theta_j) \cong \theta_j - \frac{{\theta_j}^2}{2}$, and omitting the terms beyond ${\theta_j}^2$ results in a relative error in $\ln(1 + \theta_j)$ on the order of 10^{-16}. The sum above is therefore approximately

$$\sum_{j=1}^{2^n} \frac{\theta_j - \frac{\theta_j^2}{2}}{\ln 2} 2^{-n} = \frac{1}{\ln 2} \frac{\sum_{j=1}^{2^n} \theta_j}{2^n} - \frac{1}{2\ln 2} \frac{\sum_{j=1}^{2^n} \theta_j^2}{2^n}.$$

The first sum in this expression is zero, by the definition of θ_j. To evaluate the second sum, note that the sum is computed over the 2^n values of θ_j. Each of these 2^n values can be considered as a specific value of the corresponding random variable. Since these random variables have the same distribution, the 2^n values can also be treated as a sample of size 2^n from any one of these random variables. By definition, $VAR[\theta_j] = E[\theta_j^2] - E[\theta_j]^2$. The term $\frac{\sum_{j=1}^{2^n} \theta_j^2}{2^n}$ is the sample mean of θ_j^2 and is, therefore, approximately $VAR[\theta_j] + E[\theta_j]^2$. Substituting the values of $E[\theta_j]$ and $VAR[\theta_j]$ found above, the following is obtained:

$$\begin{aligned} E[x_i|\mathrm{IDEAL}] &\cong -\frac{1}{2\ln 2}\left(VAR[\theta_j] + E[\theta_j]^2\right) \\ &= -\frac{1}{2\ln 2}(2^n - 1)2^{-H_2} \end{aligned}$$

The derivation of $E[x_i|\mathrm{REAL}]$ is similar and is as follows (again omitting powers of θ_j beyond ${\theta_j}^2$).

$$\begin{aligned} E[x_i|\mathrm{REAL}] &= E\left[n + \log_2 p_{b_i}|\mathrm{REAL}\right] \\ &= \sum_{j=1}^{2^n}\left(n + \log_2 p_j\right)p_j \\ &= \sum_{j=1}^{2^n}\left(n + \frac{\ln\left(2^{-n}(1+\theta_j)\right)}{\ln 2}\right)p_j \\ &= \sum_{j=1}^{2^n}\frac{\ln(1+\theta_j)}{\ln 2}p_j \\ &= \sum_{j=1}^{2^n}\frac{\ln(1+\theta_j)}{\ln 2}2^{-n}(1+\theta_j) \\ &\cong \sum_{j=1}^{2^n}\frac{\theta_j - \frac{\theta_j^2}{2}}{\ln 2}2^{-n}(1+\theta_j) \\ &\cong \frac{1}{\ln 2}\frac{\sum_{j=1}^{2^n}\theta_j}{2^n} + \frac{1}{2\ln 2}\frac{\sum_{j=1}^{2^n}\theta_j^2}{2^n} \\ &\cong \frac{1}{2\ln 2}\left(VAR[\theta_j] + E[\theta_j]^2\right) \\ &= \frac{1}{2\ln 2}(2^n - 1)2^{-H_2} \end{aligned}$$

Reusing portions of these calculations, the variance of x_i in the IDEAL case is obtained as follows:

$$VAR[x_i|\mathrm{IDEAL}] = E\left[\left(n + \log_2 p_{b_i}\right)^2|\mathrm{IDEAL}\right] - E\left[n + \log_2 p_{b_i}|\mathrm{IDEAL}\right]^2$$

$$
\begin{aligned}
&\cong \sum_{j=1}^{2^n} \left(\frac{\theta_j - \frac{\theta_j^2}{2}}{\ln 2} \right)^2 2^{-n} - \left(-\frac{1}{2 \ln 2} (2^n - 1) 2^{-H_2} \right)^2 \\
&\cong \frac{1}{(\ln 2)^2} \frac{\sum_{j=1}^{2^n} \theta_j^2}{2^n} - \left(\frac{1}{2 \ln 2} (2^n - 1) 2^{-H_2} \right)^2 \\
&= \frac{1}{(\ln 2)^2} (2^n - 1) 2^{-H_2} - \left(\frac{1}{2 \ln 2} (2^n - 1) 2^{-H_2} \right)^2 \\
&= \frac{1}{(\ln 2)^2} (2^n - 1) 2^{-H_2} \left(1 - \frac{1}{4} (2^n - 1) 2^{-H_2} \right)
\end{aligned}
$$

Similarly, the variance of x_i in the REAL case is obtained as follows:

$$
\begin{aligned}
VAR[x_i|\text{REAL}] &= E\left[\left(n + \log_2 p_{b_i} \right)^2 |\text{REAL} \right] - E\left[n + \log_2 p_{b_i} |\text{REAL} \right]^2 \\
&\cong \sum_{j=1}^{2^n} \left(\frac{\theta_j - \frac{\theta_j^2}{2}}{\ln 2} \right)^2 2^{-n} \left(1 + \theta_j \right) - \left(\frac{1}{2 \ln 2} (2^n - 1) 2^{-H_2} \right)^2 \\
&\cong \frac{1}{(\ln 2)^2} \frac{\sum_{j=1}^{2^n} \theta_j^2}{2^n} - \left(\frac{1}{2 \ln 2} (2^n - 1) 2^{-H_2} \right)^2 \\
&= \frac{1}{(\ln 2)^2} (2^n - 1) 2^{-H_2} - \left(\frac{1}{2 \ln 2} (2^n - 1) 2^{-H_2} \right)^2 \\
&= \frac{1}{(\ln 2)^2} (2^n - 1) 2^{-H_2} \left(1 - \frac{1}{4} (2^n - 1) 2^{-H_2} \right)
\end{aligned}
$$

Note that for typical values of n, μ_I and μ_R are closely approximated as $-\frac{1}{2 \ln 2} 2^{n-H_2}$ and $\frac{1}{2 \ln 2} 2^{n-H_2}$, respectively. Also, assuming that H_2 will need to exceed n by at least a moderate amount in order to satisfy the definition of full entropy, $\sigma_I^2 = \sigma_R^2$ can be closely approximated as $\sigma^2 = \frac{1}{(\ln 2)^2} 2^{n-H_2}$. The statistic X is therefore approximately normally distributed with means and variance as follows:

$$
E[X|\text{REAL}] = -E[X|\text{IDEAL}] = -W\mu_I \cong \frac{W}{2 \ln 2} 2^{n-H_2}
$$

$$
VAR[X|\text{REAL}] = VAR[X|\text{IDEAL}] = W\sigma^2 \cong \frac{W}{(\ln 2)^2} 2^{n-H_2}
$$

Now consider the probability that the adversary correctly determines whether the REAL or IDEAL case produced the observed sample B. This probability is as follows:

$$
\begin{aligned}
Pr[\text{Correct}] &= Pr[\text{IDEAL}]Pr[\text{Correct}|\text{IDEAL}] + Pr[\text{REAL}]Pr[\text{Correct}|\text{REAL}] \\
&= \frac{1}{2} Pr[X < 0|\text{IDEAL}] + \frac{1}{2} Pr[X > 0|\text{REAL}]
\end{aligned}
$$

Note that because of the symmetry resulting from X having a normal distribution with variance $W\sigma^2$ in both the REAL and IDEAL cases and expected values that are negatives of each other in these two cases, $Pr[X < 0|\text{IDEAL}] = Pr[X > 0|\text{REAL}]$, which gives the following:

$$\begin{aligned} Pr[\text{Correct}] &= Pr[X < 0|\text{IDEAL}] \\ &= Pr\left[\frac{X - W\mu_I}{\sqrt{W\sigma^2}} < \frac{0 - W\mu_I}{\sqrt{W\sigma^2}}|\text{IDEAL}\right] \end{aligned}$$

Since in the IDEAL case, X is normally distributed with mean $W\mu_I$ and variance $W\sigma^2$, the value $z = \frac{X - W\mu_I}{\sqrt{W\sigma^2}}$ is a standard normal random variable, so this probability is $F\left(\frac{-W\mu_I}{\sqrt{W\sigma^2}}\right)$, where F is the cumulative distribution function of the standard normal distribution. $F(x) \leq \frac{1}{2} + \frac{1}{2}\sqrt{1 - e^{-2x^2/\pi}}$ when $x > 0$ (see Section 26.2.24 of [4]). Thus, $Pr[\text{Correct}] = F\left(\frac{-W\mu_I}{\sqrt{W\sigma^2}}\right) \leq \frac{1}{2} + \delta$ if the following inequality is satisfied:

$$\frac{1}{2} + \frac{1}{2}\sqrt{1 - e^{-2\left(\frac{-W\mu_I}{\sqrt{W\sigma^2}}\right)^2/\pi}} \leq \frac{1}{2} + \delta$$

From the derivations above, $\frac{-W\mu_I}{\sqrt{W\sigma^2}} = \frac{1}{2}\sqrt{W} \cdot 2^{\frac{n-H_2}{2}}$, giving the following sequence of inequalities:

$$\frac{1}{2}\sqrt{1 - e^{-2\left(\frac{1}{4}W \cdot 2^{n-H_2}\right)/\pi}} \leq \delta$$

$$1 - e^{-\frac{1}{2}W \cdot 2^{n-H_2}/\pi} \leq 4\delta^2$$

$$1 - 4\delta^2 \leq e^{-\frac{1}{2}W \cdot 2^{n-H_2}/\pi}$$

$$\ln(1 - 4\delta^2) \leq -\frac{1}{2}W \cdot 2^{n-H_2}/\pi$$

$$-2\pi \ln(1 - 4\delta^2) \geq W \cdot 2^{n-H_2}$$

$$\log_2(2\pi) + \log_2(-\ln(1 - 4\delta^2)) \geq \log_2 W + n - H_2$$

$$H_2 \geq n + \log_2 W - \log_2(2\pi) - \log_2(-\ln(1 - 4\delta^2))$$

Since $4\delta^2 \cong 0$ when $\delta \cong 0$, $-\ln(1 - 4\delta^2)$ is closely approximated by $4\delta^2$, so the inequality can be written as:

$$H_2 \geq n + \log_2\left(\frac{W}{\delta^2}\right) - log_2\pi - 3$$

This derivation has shown that if the above inequality is satisfied by a sufficiently high value of H_2, then the distinguishing probability $Pr[\text{Correct}] \leq \frac{1}{2} + \delta$. Now note that min-entropy H is a lower bound on collision-entropy H_2 and consider the following inequality where collision-entropy H_2 is replaced with min-entropy H:

$$H \geq n + \log_2\left(\frac{W}{\delta^2}\right) - log_2\pi - 3$$

If this inequality is satisfied by a sufficiently high value of min-entropy H, then the previous inequality involving collision-entropy H_2 is also satisfied, and $Pr[\text{Correct}] \leq \frac{1}{2} + \delta$.

Q.E.D.

A.2. Justification of the Claim on Higher Powers of θ_j

In the proof of Claim 1, sums of the powers of θ_j higher than θ_j^2 were omitted. This does not affect the validity of the conclusion if θ_j is sufficiently near zero. This is established as follows. Recall that there are 2^n values of θ_j, each of which is approximately normally distributed with mean zero and variance approximately 2^{n-H_2}. Consider the largest θ_j: $\theta_{max} = max_j\{\theta_j\}$. θ_{max} is $z = \frac{\theta_{max}}{2^{\frac{n-H_2}{2}}}$ standard deviations away from zero (which was shown in A.1 to be the mean of θ_j). The value of z is expected to be such that in a collection of 2^n standard normal random variates, approximately one of the variates is greater than or equal to this value of z. If $f(z)$ and $F(z)$ are the density function and the cumulative distribution function of the standard normal distribution, respectively, then for large z, $1 - F(z) \cong \frac{f(z)}{z}$ (see Section 26.2.12 of [4]). The desired value of z, therefore, gives $(1 - F(z))2^n \cong 1$, which leads to $\frac{2^n}{z\sqrt{2\pi}}e^{-\frac{z^2}{2}} = 1$, or $z^2 + 2\ln z = 2n\ln 2 - \ln(2\pi)$. Since z^2 dominates the left side of this equation for moderately large values of z, the desired value of z is approximately $\sqrt{2n\ln 2 - \ln\,(2\pi)}$. The value of θ_{max} is then expected to be approximately $2^{\frac{n-H_2}{2}}\sqrt{2n\ln 2 - \ln(2\pi)}$. For any of the typical values of n and a value of H_2 given by the lower bound established in Claim 1, $H_2 \geq n + 64$, so $2^{\frac{n-H_2}{2}} \leq 2^{-32}$, and it can be calculated that θ_{max} is a positive value that is likely to be less than 10^{-8}. A similar argument leads to θ_{min} being approximately $-\theta_{max}$, so it is expected that $|\theta_j| \leq 10^{-8}$ for all j. Therefore, it is safe to omit powers of θ_j higher than θ_j^2, since it is shown in A.1 that doing so has a negligible effect.

A.3. Proof of Claim 2

It is shown in A.2 that $\theta_{max} \cong 2^{\frac{n-H_2}{2}}\sqrt{2n\ln 2 - \ln(2\pi)}$, which implies that the corresponding value $p_{max} = max_j\{p_j\}$ is approximately $2^{-n}\left(1 + 2^{\frac{n-H_2}{2}}\sqrt{2n\ln 2 - \ln(2\pi)}\right)$. If the min-entropy of the input to the conditioning function is H, then $H_2 \geq H$, so

$$p_{max} \leq 2^{-n}\left(1 + 2^{\frac{n-H}{2}}\sqrt{2n\ln 2 - \ln(2\pi)}\right).$$

The min-entropy corresponding to this value of p_{max} is:

$$-\log_2 p_{max} \geq n - \log_2\left(1 + 2^{\frac{n-H}{2}}\sqrt{2n\ln 2 - \ln(2\pi)}\right)$$

$$= n - \frac{ln\left(1 + 2^{\frac{n-H}{2}}\sqrt{2n\ln 2 - \ln(2\pi)}\right)}{ln2}$$

Since $H \geq n + 64$, $2^{\frac{n-H}{2}}\sqrt{2n\ln 2 - \ln(2\pi)}$ is a very small positive number, so $ln\left(1 + 2^{\frac{n-H}{2}}\sqrt{2n\ln 2 - \ln(2\pi)}\right) \cong 2^{\frac{n-H}{2}}\sqrt{2n\ln 2 - \ln(2\pi)}$, giving

$$-\log_2 p_{max} \geq n - \frac{2^{\frac{n-H}{2}}\sqrt{2n\ln 2 - \ln(2\pi)}}{\ln 2}$$

Dividing this value by n gives an average per-bit min-entropy of at least

$$1 - \frac{2^{\frac{n-H}{2}}\sqrt{2n\ln 2 - \ln(2\pi)}}{nln2}$$

When $H \geq n + 64$, a per-bit min-entropy of at least $1 - 2^{-32}c$ is obtained, where $0 < c < 1$ for all the values of n of interest. Therefore, when $H \geq n + 64$, the average per-bit min-entropy of the n-bit conditioning function output is at least $1 - 2^{-32}$. Q.E.D.

Appendix B. List of Symbols, Abbreviations, and Acronyms

DRBG
Deterministic Random Bit Generator

NIST
National Institute of Standards and Technology

RBG
Random Bit Generator

SP
(NIST) Special Publication

ε
A positive constant that is assumed to be no greater than 2^{-32}

E(*X*)
The expected value of the random variable X

log$_2$(x)
Base-2 logarithm of x

ln(x)
Natural logarithm of x

Var(x)
Variance of random variable x

Appendix C. Glossary

adversary

A malicious entity whose goal is to determine, to guess, or to influence the output of an RBG.

bitstring

An ordered sequence (string) of 0s and 1s. The leftmost bit is the most significant bit.

conditioning function

A deterministic function used to reduce bias and/or improve the entropy per bit.

cryptographic boundary

An explicitly defined physical or conceptual perimeter that establishes the physical and/or logical bounds of a cryptographic module and contains all of the hardware, software, and/or firmware components of a cryptographic module.

entropy

A measure of the randomness or uncertainty of a random variable.

entropy source

The combination of a noise source, health tests, and optional conditioning component that produce bitstrings containing entropy.

full-entropy bitstring

A bitstring with ideal randomness (i.e., the amount of entropy per bit is equal to 1). This publication proves that a bitstring satisfying a certain definition of *full entropy* has an entropy rate of at least $1 - \varepsilon$, where ε is at most 2^{-32}.

ideal randomness source

The source of an ideal random sequence of bits. Each bit of an ideal random sequence is unpredictable and unbiased, with a value that is independent of the values of the other bits in the sequence. Prior to an observation of the sequence, the value of each bit is equally likely to be 0 or 1, and the probability that a particular bit will have a particular value is unaffected by knowledge of the values of any or all of the other bits. An ideal random sequence of *n* bits contains *n* bits of entropy.

likelihood ratio test

A statistical test aimed at distinguishing between two competing models that could have produced an observed event based on a comparison of the likelihoods of the observed event, given the two models.

min-entropy

A lower bound on the entropy of a random variable. The precise formulation for min-entropy is $(-\log_2 \max p_i)$ for a discrete distribution having probabilities $p_1, ..., p_k$. Min-entropy is often used as a measure of the unpredictability of a random variable.