Viktoria Ronge, Christoph Egger, Russell W. F. Lai, Dominique Schröder, and Hoover H. F. Yin

Foundations of Ring Sampling

Abstract: A ring signature scheme allows the signer to sign on behalf of an ad hoc set of users, called a ring. The verifier can be convinced that a ring member signs, but cannot point to the exact signer. Ring signatures have become increasingly important today with their deployment in anonymous cryptocurrencies. Conventionally, it is implicitly assumed that all ring members are equally likely to be the signer. This assumption is generally false in reality, leading to various practical and devastating deanonymizing attacks in Monero, one of the largest anonymous cryptocurrencies. These attacks highlight the unsatisfactory situation that how a ring should be chosen is poorly understood.

We propose an analytical model of ring samplers towards a deeper understanding of them through systematic studies. Our model helps to describe how anonymous a ring sampler is with respect to a given signer distribution as an information-theoretic measure. We show that this measure is robust – it only varies slightly when the signer distribution varies slightly. We then analyze three natural samplers – uniform, mimicking, and partitioning – under our model with respect to a family of signer distributions modeled after empirical Bitcoin data. We hope that our work paves the way towards researching ring samplers from a theoretical point of view.

Keywords: Ring Signature, Anonymity, Monero

DOI 10.2478/popets-2021-0047

Received 2020-11-30; revised 2021-03-15; accepted 2021-03-16.

1 Introduction

A ring signature scheme [16] allows the signer to sign on behalf of an ad hoc chosen set of users, called a ring. The verifier can be convinced that a ring member signed, but cannot tell who it was exactly. Initially motivated by anonymous disclosure of secrets, the concept of ring signatures has subsequently been studied extensively, and has been extended to many variants such as linkable [10] and accountable [24] ring signatures. A notable extension of linkable ring signatures, known as ring confidential transactions (RingCT) [14], is the foundation of some privacy-preserving cryptocurrencies such as Monero. An overall market capitalization of more than two billion USD[1] makes Monero a high-value target of deanonymization attacks. Understanding the concrete anonymity of RingCT, or ring signatures in general, is thus unprecedentedly important.

In most applications of ring signatures and its extensions, it is implicitly assumed that all honest members of a ring are equally likely to be the actual signer(s). This assumption could be justified in applications of ring signatures where there is a natural choice of ring from the context. For example, if a high-rank member of an organization wishes to disclose a secret, a natural choice of the ring consists of all high-rank members of the organization. In other applications, such as in anonymous cryptocurrencies where ring members are picked from a universe of seemingly indifferent anonymous accounts, the signer distribution is not at all obvious. For example, it is shown that the signer distribution of (an old version of) Monero is highly correlated with the "age" distribution of the signers [12].

Picking a ring whose members have highly uneven signing probabilities could provide a false sense of anonymity. To illustrate the problem with a simple example, consider that Alice chooses to form a ring with Bob and Charlie and issues a ring signature. Suppose that an adversary somehow knows that Bob and Charlie (*e.g.*, by social engineering) are very unlikely to issue such a signature, then Alice would not have much anonymity despite using a ring signature.

In practice, imperfect rings were exploited by the devastating attacks against Monero (see Section 2.1), which sometimes completely deanonymized the signers. Although countermeasures were proposed, to the best of our knowledge, all proposals are based on the intuition derived from known exploits and are tailored to solve those specific issues.

Viktoria Ronge, Christoph Egger, Russell W. F. Lai, Dominique Schröder: Friedrich-Alexander-Universität Erlangen-Nürnberg

Hoover H. F. Yin: The Chinese University of Hong Kong

1 https://coinmarketcap.com/currencies/monero/ 28 Nov. 2020

1.1 Problem Statement

As of today, no analytical model is proposed for ring samplers, which prohibits a systematic study. For example, without such a model, it is difficult to make sense of the following questions, not to mention answering them: Is ring sampler Π better than Π'? Can we provably say that Π is good? Is Alice more anonymous than Bob when using Π?[2] These questions call for a framework for quantifying and hence comparing the anonymity of (users of) ring samplers.

1.2 Our Methodology

Unlike existing bottom-up (concrete, attack-driven) approaches, we use a top-down (abstract) approach towards understanding the anonymity of ring samplers.

1.2.1 Model of Ring Samplers

In Section 4 we model ring samplers as an oracle machine Π which optionally gets oracle access to an (estimated) signer distribution, inputs the identifiers of the signers and outputs a ring. We then propose an information-theoretic measure of the anonymity of a ring sampler with respect to a signer distribution. More concretely, given a random variable S following a signer distribution $\mathcal{S}$, we define the anonymity $\alpha[\mathsf{S}, \Pi]$ of the ring sampler Π to be the conditional min-entropy $H_\infty(\mathsf{S}|\Pi(\mathsf{S}))$. In the presence of a side-channel Λ, the anonymity is defined as

$$\alpha[\mathsf{S}, \Pi, \Lambda] := H_\infty(\mathsf{S}|\Pi(\mathsf{S}), \Lambda(\mathsf{S})).$$

We discuss potential extensions, and the difficulty thereof, to capture anonymity "over time" in Section 4.2.3.

Furthermore, in Section 5 we show that the definition is robust, in the sense that the anonymity changes only slightly when the signer distribution changes slightly. In particular, if a ring sampler is shown to be good with respect to a close estimate $\hat{\mathcal{S}}$ of the real signer distribution $\mathcal{S}$, then it should also be good with respect to the real signer distribution $\mathcal{S}$.

1.2.2 Attacks against Ring Samplers

Our definition covers all (possibly computationally unbounded) deanonymization attacks against ring samplers in which the goal of the attacker is to guess the real signer. Such attacks could be classified in two orthogonal dimensions: passive v.s. active and direct v.s. side-channel.

Passive and Active Attacks Passive adversaries have no influence on the signer distribution $\mathcal{S}$ or the execution of the ring sampler Π. This captures a wide range of "after-the-fact" attacks which rely on publicly available information such as transaction times, transaction graphs, account correlations, *etc.*

Active adversaries influence or specify the signer distribution $\mathcal{S}$, or subvert the ring sampler Π, *e.g.*, by manipulating the input randomness or its implementation. This captures attacks which are powerful but rely on stronger setup assumptions.

Direct and Side-Channel Attacks In a direct attack, the only information available to the adversary about the signer distribution $\mathcal{S}$ is a sample from $\Pi(\mathcal{S})$ output by the ring sampler. In a sense, the adversary is attempting to deanonymize the signer by directly attacking the ring sampler. In case more side-channel information is be available to the adversary, this extra information is abstracted as $\Lambda(\mathsf{S})$, where Λ is some leakage function.

1.2.3 Signer Distributions for Comparison

To understand the anonymity of different samplers, we analyze them with respect to various distributions, with a focus on the cryptocurrency context due to its high real-world impact.

A cryptocurrency consists of a history of transactions each encoding a set of spenders and a set of receivers (possibly in a hidden manner). Each set of spenders can be thought of as a sample of the real signer distribution at that particular point in time. The real signer distributions at different points in time could be correlated arbitrarily.

Due to the anonymous nature of anonymous cryptocurrencies, it is (supposedly) infeasible to learn the real signer distributions. Möser et al. [12] empirically analyzed the transaction graph of Monero in the pre-RingCT (*i.e.*, non-anonymous) era, and "heuristically determined" that the signer distribution of Monero matches a gamma distribution over the logarithm of the

2 Section 4.2.3 discusses the difficulty of formalizing this.

age[3] of accounts, which we simply call the log-gamma distribution hereinafter. While their heuristic lacks a physical interpretation, the distribution is nevertheless later used in the ring sampler of Monero [12], which can be seen as an instantiation of the mimicking sampler that we introduce in Section 6.2. As the graph analysis tools and results of [12] are not publicly available, we could not replicate their procedures of identifying the log-gamma distribution. We remark that there is no guarantee that the Monero distributions in the pre- and post-RingCT (*i.e.*, current) eras are similar.

Another reference is the signer distribution of a non-anonymous cryptocurrency such as Bitcoin, despite the potential differences in spending behavior in a non-anonymous cryptocurrency compared to anonymous ones. To this end, we analyze the 300,000 to 400,000-th block of Bitcoin as in [12], and found that the age of a transaction output matches a (shifted) Pareto distribution. We thus propose to use the (shifted) Pareto distributions as a baseline for evaluating ring samplers.

For the appropriate parameters, the probability density functions (PDFs) of the (shifted) Pareto distribution and the log-gamma distribution have very similar asymptotic behavior. Indeed, their PDFs only differ by a poly-logarithmic factor. In terms of physical interpretation, we found the modeling by the (shifted) Pareto distribution more convincing as it is classically used to model a wide range of human-related phenomenon, whereas the log-gamma distribution seems somewhat arbitrary. We emphasize, however, that the proposed baseline distribution, or any other non-application-specific ones, should only be treated as reference points. Even if a ring sampler is good with respect to the baseline distribution, it does not necessarily mean that it is good for a particular anonymous cryptocurrency, since the signer distributions of anonymous and non-anonymous cryptocurrencies could be very different.

1.2.4 Analysis of Natural Samplers

We analyze the anonymity of three natural families of samplers – uniform, mimicking, and partitioning – with respect to general (possibly adversarially influenced) signer distribution $\mathcal{S}$, assuming that the samplers are not subverted.

Uniform Samplers The uniform sampler samples uniformly random rings. It is shown to be bad with respect to signer distributions which are far from uniform, such as the (shifted) Pareto distributions. The fact that the uniform sampler is generally bad is expected. In a signer distribution which is very far away from uniform, the majority of users are unlikely to be the signer. Therefore the real signer would end up with a ring in which most members other than himself are unlikely to be the signer, and thus stand out from the crowd.

Mimicking Samplers The $\hat{\mathcal{S}}$-mimicking sampler is an abstraction of the (current) Monero sampler. It is given oracle access to some signer distribution $\hat{\mathcal{S}}$, which supposedly estimates the real distribution $\mathcal{S}$, and aims to output rings which "mimic" $\hat{\mathcal{S}}$ in some sense. For the special case where $\hat{\mathcal{S}} = \mathcal{S}$, which we call *the* mimicking sampler, we prove that its anonymity is lower bounded by half of the optimal value. The tightness of the lower bound is limited by the use of an intermediate distribution, which has lower conditional min-entropy but is easier to analyze, and the available bounding techniques. We believe that the exact anonymity should be considerably closer to the optimal value. Suppose that is the case, due to robustness, the $\hat{\mathcal{S}}$-mimicking sampler is also good when $\hat{\mathcal{S}} \approx \mathcal{S}$. This can be seen as a theoretical confirmation of the approach used in Monero, albeit conditioned on the strong assumption that Monero chose a good $\hat{\mathcal{S}}$.

The major drawback of the $\hat{\mathcal{S}}$-mimicking sampler is the requirement of the knowledge of an estimation $\hat{\mathcal{S}}$ of $\mathcal{S}$. Indeed, as $\mathcal{S}$ could depend on the economic situation and the free will of signers, it is arguably unknowable and inapproximable. Even if an estimation $\hat{\mathcal{S}}$ is known, the description of an $\hat{\mathcal{S}}$-mimicking sampler could be quite complicated depending on the description of the distribution $\hat{\mathcal{S}}$, and its anonymity depends on how well $\hat{\mathcal{S}}$ estimates $\mathcal{S}$. It could also be difficult to write down its exact anonymity if $\hat{\mathcal{S}}$ lacks a close form.

Partitioning Samplers Partitioning samplers are based on another natural strategy of grouping users with similar signing probabilities together. More concretely, a partitioning sampler is defined by a distribution over a public family of partitions and optionally a ring size n. A ring is sampled by sampling a partition from the distribution, and then, if the ring size n is given, outputting a uniformly random n-subset of the chunk (an element of the partition) which contains the signer. If n is not given, the sampler simply outputs the unique chunk containing the signer. A special case of the partitioning sampler was also suggested in [25].

3 The difference between spent and creation time.

We show that the anonymity of a partitioning sampler is at most $\lg \varepsilon$ away from the optimal value, where ε measures the non-uniformity of signing probabilities within chunks of the partitions. For the variant where n is given, if the partitions are chosen in such a way that the signing probabilities are constant within each chunk (which can be done naturally for the baseline distributions), the partitioning sampler is optimal.

Partitioning samplers are easy to describe and preferable in practice. Depending on how partitions are chosen, they could also have other nice properties which are not captured by our model. We refer to Section 6.3 for details.

1.2.5 Implication to Ring Signatures

To help grasp the meaning of our work more concretely, in Appendix B we define a generalized notion of ring signatures which captures extended variants such as linkable ring signatures. We also define a simulation-based notion of anonymity which, although being equivalent to the usual indistinguishability-based notion, synergizes better with our anonymity notion of ring samplers. Finally, we define the concrete anonymity of the composition of ring signatures and ring samplers, and relate the anonymity of the composed system to those of the components.

2 Related Work

To better position our work in the literature, we overview the existing attacks against Monero, other ring sampler formalizations, and the formalizations of other anonymous systems.

2.1 Attacks against Monero

Our formal study of ring samplers is motivated by the deanonymization attacks against Monero. To understand them, we first explain how ring samplers are involved.

In Monero, a spender can spend funds from possibly multiple source accounts to possibly multiple receivers as follows. First, the spender samples a ring of potential source accounts which is a superset of the actual source accounts. It then creates a proof that it knows the secret keys of the actual source accounts and that it wishes to transfer funds to some specified target accounts. Such a proof can be seen as a generalization of a ring signature.

While ring signatures provably guarantee anonymity, the concrete anonymity of signers can only be as high as what is offered from the sampled ring. In the following we overview existing deanonymization attacks against Monero which are mostly based on ill-chosen rings.

2.1.1 Passive Direct Attacks

Passive direct attacks against Monero exploit public information available after the target transaction is made. These attacks are particularly well-captured by our anonymity definition and the analysis of the natural samplers. They also constitute the majority of existing attacks, and are more realistic due to the minimal assumption on the attacker.

Exploiting Transaction Times The age of an account influences the likelihood of it being an actual source of a transaction [12]: Old accounts become increasingly less likely to still be unspent and therefore be an actual source account of a transaction. On the other hand, freshly created target accounts are highly likely to be used as source accounts in transactions soon. Using the above observation, the ring sampling strategy which selects accounts uniformly at random over the set of all accounts is not a good idea, as younger accounts are less likely to function as decoys in the ring. These attacks have been deployed in [9, 12].

In particular, for over 95% of existing transactions in an older pre-RingCT version of Monero, the newest account in the ring is the signer [9]. This makes the simple attack of guessing the newest account to be the signer devastating (with 95% success rate), highlighting the importance of using a good ring sampler.

Exploiting Graph Structures When rings selected for different transactions overlap, by analyzing the graph induced by the relation between the rings, one can infer non-trivial information about the actual source ac-

4 But already has many ideas of permanent-based metrics
5 In the form of information gain
6 Global is mentioned but considered infeasible
7 A metric measuring the loss of anonymity computed from the conditional distributions of the observable events
8 In the form of relative entropy

Work	Venue	Domain	Metric	Scope	Remarks
[5, 17]	PETS'02	Communication	Shannon-Entropy	Local	(Co-)introduced information-theoretic metrics
[13]	PETS'03	Communication	Shannon-Entropy[4]	Global	Extended information-theoretic metrics to global scope
[4]	FAST'06	Communication	Shannon-Entropy[5]	Local	Integration with process calculi
[7]	ISI'07	Communication	Permanent	Global	Permanent-based metric to better capture global scope
[20]	PETS'04	Mixnet	Posterior Probability	Local[6]	Handles cover traffic
[8]	PETS'13	Tor	Shannon-Entropy	Local	Evaluating performance trade-offs using the metrics
[2]	PETS'13	Tor	"Impact"[7]	Local	Evaluating real-world attacks
[18, 19]	PETS, SP'11	Location	Shannon-Entropy[8]	Both	New metric & tool applicable to location privacy
[25]	CSF'19	Ring Sampling	Permanent	Global	Formal analysis of ring sampling, orthogonal work
Ours		Ring Sampling	Min-Entropy	Global	Introduced information-theoretic metrics to ring sampling

Table 1. Comparison of Entropy/Probability-Based Anonymity Metrics

counts of a transaction (*e.g.*, [9, 12]). In the extreme case, which is considered in the "zero-mixin attack", some transactions use rings of size one. If such a zero-mixin account is used in other rings, it does not add any anonymity as an observer can clearly rule out this account as a possible source (*c.f.*, the illustrative example in Section 1).

Exploiting Correlated Accounts Transactions with multiple source accounts expose an additional problem [9]. For example, let more than one target account be the output of the same previous transaction. If these accounts are included in a ring of a subsequent transaction with multiple inputs, then it is quite likely that they are the actual source of the subsequent transaction. This attack is based on the implicit assumption that output accounts in one transaction have a significant chance to belong to the same receiver and that both output accounts being chosen as decoys is low.

2.1.2 Active Direct Attacks

The best known attack of this kind against Monero is the so called "black marble attack", proposed in [11, 15, 23], of which there are two variants.

In the first variant, the attacker compromises existing accounts or spawns new accounts in the system, and hopes that some of them (the black marbles) will be included in future rings. This can be modelled by considering an adversarially influenced signer distribution $\mathcal{S}$. If it happens that a ring chosen by the victim consists of mostly black marbles, then the anonymity of the victim is severely limited. While this attack seems reasonable in theory, its practicality is unclear. Even for the provably bad uniform sampler, the probability of randomly picking a black marble as a ring member is low, assuming the universe of signers is large. To increase this probability, the attacker could spawn an overwhelming number of accounts, which however requires substantial transaction fees.

In the second variant, the attacker subverts the victim's ring sampler, so that black marbles are injected into the rings chosen by the signer. While this attack is in no doubt devastating, the assumption on the attacker's ability to subvert ring samplers is very strong. Indeed, if subversion is allowed, the attacker might as well directly embed the signer's identity in the chosen ring in an undetectable manner[9], without going through all the trouble of spawning black marbles. No ring sampler could defend against this.

We remark that although our analysis of the partitioning sampler does not consider subversion attacks, the fact that its output must be a subset of a chunk of a publicly defined partition limits the flexibility of the attacker in planting black marbles.[10] The signer could easily detect the subversion if the sampled ring is "illegal" (*e.g.*, if it contains black marbles chosen from chunks where the signer does not belong to).

2.1.3 Side-Channel Attacks

In existing implementations of Monero, a client consists of two parts known as the node and the wallet, which may or may not be co-located in the same device. Some

9 The subverted ring sampler could for example only output rings whose hash value equals to a one-time pad of the i-th bit of the signer's identity or secret key. From a steganographic point of view [26], this is provably undetectable. If the setting allows to repeat the procedure many times for different i, the attacker could recover the exact identity of the signer.

10 Depending on the choice of the partition(s), the subset and the chunk that the signer belongs to could even be unique.

side-channel attacks (*e.g.*, [21]) exploit the communication patterns between the two parts, and/or the reaction of the client to inbound communication, to deduce whether the client is the intended receiver of a transaction. Such attacks are purely side-channel and are independent of the ring sampler. Moreover they aim to deanonymize the receiver but not the signer, which is out of the scope of our model.

2.2 Other Formalization Attempts

Yu, Au, and Esteves-Verissimo [25] analyzed ring samplers from another perspective: They studied the anonymity of a group of signers after all of them have chosen particular instances of rings, while implicitly assuming that the signer distribution is uniform. The core technique of their study is modelling signer identities using graphs to rule out impossible signers. This can be seen as an instantiation of the work by Edman, Sivrikaya, and Yener [7] who studied matrix permanents to understand potential message flows in an anonymous communication system.

In contrast, we model ring samplers as probabilistic algorithms, and focus on analyzing the (loss of) anonymity of signers when given individual rings. The two approaches are orthogonal and complementary [7] since they give different insights about the anonymity of ring samplers. We envision the unification of them towards a more comprehensive theory of ring samplers.

Although the approach taken by Yu, Au, and Esteves-Verissimo [25] is very different, they arrived at a similar conclusion to ours that a partitioning sampler (using our terminology) is optimal (in both our and their sense). Indeed, their partitioning sampler can be seen as a specific instantiation of our generic one. Furthermore, we analyzed the robustness of our definition, and the anonymity of the uniform and the mimicking samplers. Similar results were not in [25].

2.3 Other Anonymous Systems

It is common to quantify the anonymity of anonymous systems, in particular anonymous communication systems. Proposed metrics include information-theoretic measures, permanents (of induced bipartite graphs), or other metrics derived from probabilities. Table 1 provides an overview of these quantification efforts, but due to the volume of the literature it is necessarily incomplete. For a comprehensive survey, see [22].

Historically it is not uncommon that concrete tasks are guided by anonymity metrics after the latter are sufficiently well studied. These include for example the analysis of performance tradeoffs [8] and attacks [2] against the Tor system. It has also been noted that other areas like location privacy [18] profit from the guidance of information-theoretic measures. However a information-theoretic treatment of ring (or generally speaking, decoy) selection is, thus far, missing in literature.

Unique to ring samplers is the use of decoys (ring members), which means that the set of parties whose anonymity is to be preserved is always a (proper) subset of the the observed ring – a property that we will use repeatedly in proofs. This aspect makes ring samplers different from, *e.g.*, mixnet-style anonymous communication where all inputs to a mixer are "real".

3 Preliminaries

Denote by λ the security parameter. For $M, N \in \mathbb{N}$, we denote $[N] := \{1, \ldots, N\}$ and $[M : N] := \{M, M+1, \ldots, N\}$. Logarithms are either with base 2, denoted by lg, or natural, denoted by ln. The sets of polynomials and negligible functions in λ are denoted by $\mathsf{poly}(\lambda)$ and $\mathsf{negl}(\lambda)$ respectively. Probabilistic polynomial time is abbreviated as PPT. If $\mathcal{A}$ is a PPT algorithm, $y \leftarrow \mathcal{A}(x)$ means assigning the result of running $\mathcal{A}$ on x (with implicit randomness) to y. Sets are denoted by capital letters. For a finite set S, $x \leftarrow\!\!\$\, S$ means that a random x is chosen uniformly from S. An algorithm $\mathcal{A}$ with oracle access to a subroutine $\mathcal{R}$ is written as $\mathcal{A}^{\mathcal{R}}$.

Let f and g be real-valued functions. If f is proportional to g, *i.e.*, $f(x) = k \cdot g(x)$ for all x for some constant k, we write $f \propto g$. We use this primarily to express probability density functions (PDFs) without specify the normalizing constant.

We denote the power set of S by 2^S. If $A \subseteq B$ and $|A| = n$, we write $A \subseteq_n B$. If $|A| \leq n$, we write $A \subseteq_{\leq n} B$.

3.1 Random Variables and Min-Entropy

(Discrete) random variables are written in sans-serif, *e.g.*, X, and distributions of random variables are written in calligraphic, *e.g.*, $\mathcal{X}$. When the connection is obvious

we only denote the distribution and use the same letter in sans-serif for the underlying random variable.

The support of X is denoted by $\mathsf{Supp}(\mathsf{X}) := \{X : \Pr[\mathsf{X} = X] > 0\}$. If it is obvious that X has support S, we write $X \in S$ instead of $X \in \mathsf{Supp}(\mathsf{X})$. When summing over all X in the support of X, *i.e.*, $\sum_{X \in \mathsf{Supp}(\mathsf{X})}$, we usually omit $\mathsf{Supp}(\mathsf{X})$ and simply write $\sum_X$ unless there is an ambiguity. The same holds when taking minimum or maximum. For a function f, we write $f(\mathsf{X}) \equiv Y$ if $\Pr[f(\mathsf{X}) = Y] = 1$.

Definition 3.1. *The guessing probability of* X *is*

$$\mathsf{Guess}(\mathsf{X}) := \max_X \Pr[\mathsf{X} = X].$$

This gives an upper bound on the probability that an (unbounded) adversary can "guess" the value of the random variable X correctly.

Definition 3.2. *The (average) conditional guessing probability of* X *given* Y *is defined as*

$$\mathsf{Guess}(\mathsf{X}|\mathsf{Y}) := \sum_Y \Pr[\mathsf{Y} = Y] \max_X \Pr[\mathsf{X} = X | \mathsf{Y} = Y].$$

This gives an upper bound on the probability that an (unbounded) adversary "guesses" the value of the random variable X correctly when given a sample of Y.

Definition 3.3. *The min-entropy of* X *is defined as*

$$H_\infty(\mathsf{X}) := -\lg(\mathsf{Guess}(\mathsf{X})).$$

The min-entropy of X is in a sense the most pessimistic measure of information given in X. It has significance in randomness extraction in the sense that nearly $H_\infty(\mathsf{X})$ random bits can be extracted from the source X [6]. Therefore a higher value of $H_\infty(\mathsf{X})$ is desirable. As is common in cryptography, in this work we consider only min-entropy. All definitions however naturally extend to other measures of entropy.

Definition 3.4. *The (average) conditional min-entropy of* X *given* Y *is defined as*

$$H_\infty(\mathsf{X}|\mathsf{Y}) := -\lg(\mathsf{Guess}(\mathsf{X}|\mathsf{Y})).$$

The conditional min-entropy has similar interpretations as that of min-entropy. There are several equivalent expressions of the conditional min-entropy useful for different occasions.

Lemma 3.1. $H_\infty(\mathsf{X}|\mathsf{Y})$ *can be expressed as:*

$$\begin{aligned} & H_\infty(\mathsf{X}|\mathsf{Y}) \\ &= -\lg\left(\sum_Y \Pr[\mathsf{Y} = Y] \max_X \Pr[\mathsf{X} = X | \mathsf{Y} = Y]\right) \\ &= -\lg\left(\sum_Y \max_X \Pr[\mathsf{Y} = Y | \mathsf{X} = X] \Pr[\mathsf{X} = X]\right) \\ &= -\lg\left(\sum_Y \max_X \Pr[\mathsf{X} = X \wedge \mathsf{Y} = Y]\right). \end{aligned}$$

The following properties about min-entropy and conditional min-entropy are well-known.

Lemma 3.2 (Non-Negativity, Monotonicity)**.** *For any random variables* X *and* Y*,*

$$0 \le H_\infty(\mathsf{X}|\mathsf{Y}) \le H_\infty(\mathsf{X}).$$

Lemma 3.3 (Data Processing Inequality)**.** *Let* $\mathsf{S}, \mathsf{R}, \mathsf{X}$ *be random variables where* $\mathsf{R} = f(\mathsf{X})$ *for some function* f*. Then* $H_\infty(\mathsf{S}|\mathsf{R}) \ge H_\infty(\mathsf{S}|\mathsf{X})$*.*

Proof. Note that $H_\infty(\mathsf{S}|\mathsf{R},\mathsf{X}) = H_\infty(\mathsf{S}|\mathsf{X})$ (trivial) and $H_\infty(\mathsf{S}|\mathsf{R},\mathsf{X}) - H_\infty(\mathsf{S}|\mathsf{R}) \le 0$ (monotonicity). Therefore $H_\infty(\mathsf{S}|\mathsf{X}) \le H_\infty(\mathsf{S}|\mathsf{R})$. □

We recall the Rényi divergence (of ∞-order) to measure the closeness of two distributions.

Definition 3.5. *Let* S *and* S' *be such that* $\mathsf{Supp}(\mathsf{S}) \subseteq \mathsf{Supp}(\mathsf{S}')$*. Their Rényi divergence of order* ∞ *is*

$$D_\infty(\mathsf{S}\|\mathsf{S}') := \lg \max_{S \in \mathsf{Supp}(\mathsf{S}')} \frac{\Pr[\mathsf{S} = S]}{\Pr[\mathsf{S}' = S]}.$$

4 Modeling

In this section, we devise a formal model of ring samplers which in particular includes an information-theoretic measure of anonymity. We also derive general lower and upper bounds of anonymity, and discuss its extensions.

4.1 Syntax

Throughout this work, we consider a universe of users indexed by the set $[N]$ where a subset S of them wish to hide themselves among a ring of users.

Definition 4.1 (Signer Distributions)**.** *A signer distribution* $\mathcal{S}$ *is a distribution over the set* $2^{[N]} \setminus \{\emptyset\}$*. Let* $k \in [N]$*.* $\mathcal{S}$ *is said to be a* k*-signer distribution if* $\Pr[|\mathsf{S}| \le k] = 1$*.*

In this work, we focus mostly on 1-signer distribution, *i.e.*, there is only one signer. Next, we state a minimalistic syntax of ring samplers.

Definition 4.2 (Ring Samplers)**.** *A ring sampler* Π *is a* PPT *(oracle) machine which inputs a set of signers* $S \subseteq [N]$ *and outputs a ring* R *satisfying* $S \subseteq R \subseteq [N]$*. Let* $k \leq n \in [N]$*.* Π *is said to an* n*-ring sampler if it always holds that* $|R| \leq n$*. If additionally* Π *only takes* S *with* $|S| \leq k$ *as input, then it is a* (k, n)*-ring sampler.*

For concreteness, think of k and n to be small constants, *e.g.*, $k = 1$ or 2, and $n = 10$, and $\lg N = \mathsf{poly}(\lambda)$. In such cases, the input and output of Π can each be represented by $\mathsf{poly}(\lambda)$ bits, and Π should run in time $\mathsf{poly}(\lambda)$.

4.2 Anonymity

We measure the quality of a ring sampler in terms of its anonymity, *i.e.*, the difficulty to guess who the signer(s) are when given a ring. From an adversary's point of view, before seeing any information about the signing/transaction event, *e.g.*, a ring $\mathsf{R} = \Pi(\mathsf{S})$, the anonymity of all participants is considered to be maximum, and can be measured by the value $H_\infty(\mathsf{S})$.

The knowledge of R or other side-channel leakage of S can only reduce the anonymity from $H_\infty(\mathsf{S})$ towards zero. From this viewpoint, let Λ be a leakage function capturing the side-channel. We define the anonymity of a ring sampler Π with respect to S in presence of the side-channel Λ as the min-entropy of S conditioning on the ring and the leakage.

Definition 4.3 (Anonymity)**.** *Let* Π *be a ring sampler and* $\Lambda : \{0,1\}^* \to \{0,1\}^*$ *be a leakage function. The anonymity of* Π *with respect to* S *in presence of* Λ *is defined as*

$$\alpha(\mathsf{S}, \Pi, \Lambda) := H_\infty(\mathsf{S}|\Pi(\mathsf{S}), \Lambda(\mathsf{S}))$$

where $\Pi(\mathsf{S})$ *is the random variable induced by applying* Π *on* S *with uniform randomness, and* $\Lambda(\mathsf{S})$ *is the leaked side-channel information about* S *due to* Λ*. If* Λ *is a constant function (i.e., there is no leakage), then we simply write*

$$\alpha(\mathsf{S}, \Pi) := H_\infty(\mathsf{S}|\Pi(\mathsf{S}))$$

and regard it as the anonymity of Π *with respect to* S*.*

4.2.1 Scope and Implications

Our approach of defining anonymity is natural and general in the sense that it can be adopted to any anonymous system. Abstractly, if S is a distribution over a set of objects whose anonymity is to be protected by an anonymous system Π, and Λ is a leakage function capturing any side-channel information leakage external to Π, then the anonymity of Π with respect to S in presence of Λ can be measured by $H_\infty(\mathsf{S}|\Pi(\mathsf{S}), \Lambda(\mathsf{S}))$, exactly like how we measure the anonymity of a ring sampler.

Our definition captures all deanonymization attacks: passive, active, direct, and side-channel, by any computationally unbounded adversary. Given a sample of the induced ring distribution $\Pi(\mathsf{S})$ and a sample of the leakage $\Lambda(\mathsf{S})$, the goal of a deanonymizing adversary is to output a guess of the signer S.

Remark 1. *While our definition captures active and side-channel attacks, it is somewhat unnatural. A more convenient and expressive way of capturing those is through security experiments akin to those used in (computational) cryptography, such as the ones defined in Appendix B for the anonymity of the composition of ring samplers and ring signatures.*

For a more concrete feeling of the definition, we state the following immediate implication on *any* deanonymization attacks from *any* computationally unbounded adversaries. The proof is obvious and is omitted.

Theorem 4.4. *Let* $\mathcal{A}$ *be any computationally unbounded adversary, who inputs a ring* $\Pi(S)$ *(where* Π *is possibly subverted by* $\mathcal{A}$*) and some leakage* $\Lambda(S)$*, where* S *is sampled from the distribution* $\mathcal{S}$ *(possibly influenced or specified by* $\mathcal{A}$*), and outputs a guess* S'*. The probability of* $\mathcal{A}$ *correctly guessing the signer* S*, i.e.,* $S' = S$*, is upper bounded by*

$$\mathsf{Guess}\left(\mathsf{S}|\Pi(\mathsf{S}), \Lambda(\mathsf{S})\right) = 2^{-\alpha(\mathsf{S}, \Pi, \Lambda)}.$$

Thus, α directly relates to the probability an adversary is guessing the actual spender correctly.

4.2.2 Basic Properties

Intuitively, a higher value of $\alpha(\mathsf{S}, \Pi)$ (or $\alpha(\mathsf{S}, \Pi, \Lambda)$) means a higher anonymity, or rather, the amount of anonymity lost due to the use of the ring sampler (and the leakage) is smaller. Due to monotonicity

(Lemma 3.2), $\alpha(\mathsf{S}, \Pi, \Lambda)$ lies between zero and $H_\infty(\mathsf{S})$ for any Π and Λ, which aligns with our subtractive view of anonymity.

General Bounds When analyzing the fundamentals of a ring sampler Π, it is instrumental to focus on the value $\alpha(\mathsf{S}, \Pi)$ even if there might be some external leakage Λ which is not the "fault" of the sampler Π. The following lemma relates the anonymity definitions with and without leakage. Its proof follows immediately from the chain rule and the monotonicity of min-entropy.

Lemma 4.1. *For any* S*,* Π*, and* Λ*, it holds that*

$$\alpha(\mathsf{S}, \Pi) - \lg |\mathsf{Supp}(\Lambda(\mathsf{S}))| \leq \alpha(\mathsf{S}, \Pi, \Lambda) \leq \alpha(\mathsf{S}, \Pi).$$

Although the above bound is loose (compared to the Shannon entropy counterpart), it suggests that the anonymity of the ring sampler without leakage, *i.e.*, $\alpha(\mathsf{S}, \Pi)$, is the dominating component of $\alpha(\mathsf{S}, \Pi, \Lambda)$ when the max-entropy $\lg |\mathsf{Supp}(\Lambda(\mathsf{S}))|$ of the leakage is small. For the typical size leakage where $\Lambda(\mathsf{S}) = |\mathsf{S}|$, where $\mathcal{S}$ is a k-signer distribution and Π is a (k, n)-ring sampler with $k \ll n$, this is indeed the case since $\lg |\mathsf{Supp}(\Lambda(\mathsf{S}))| = \lg k \ll \lg n$.

Let Π_{All} be the "all sampler" which always outputs $[N]$. Let Π_{Id} be the "identity sampler" which on input S outputs S. We first state some trivial bounds of anonymity, and how they can be achieved. The proof is obvious and omitted.

Lemma 4.2. *For any* S*,* Π*, and* Λ*, it holds that*

$$\alpha(\mathsf{S}, \Pi_{\mathrm{Id}}, \Lambda) = 0 \leq \alpha(\mathsf{S}, \Pi, \Lambda) \leq H_\infty(\mathsf{S}) = \alpha(\mathsf{S}, \Pi_{\mathrm{All}}).$$

The upper and lower bounds above are too far apart to tell us anything useful about $\alpha(\mathsf{S}, \Pi, \Lambda)$. The main reason is that the "all sampler" and the "identity sampler" have extreme ring sizes ($n = N$ and k respectively), while in practice we are interested in n-ring samplers for (small) fixed $n > k$. We therefore state another upper bound of anonymity of n-ring samplers, whose proof can be found in Appendix C.

Lemma 4.1. *For any* k*-signer distribution* S*, any* n*-ring sampler* Π*, and any leakage function* Λ*,*

$$\alpha(\mathsf{S}, \Pi, \Lambda) \leq \lg \sum_{i=1}^{k} \binom{n}{i}.$$

In particular, for $k = 1$ *we have*

$$\alpha(\mathsf{S}, \Pi, \Lambda) \leq \lg n.$$

Optimality Our definition of anonymity in a sense describes the anonymity of the system employing the ring sampler as a whole. In other words, the (conditional) signing probabilities of individual signers are collapsed into a single value. In Lemma 4.1, we showed that the anonymity of a ring sampler with ring size n for a 1-signer distribution is at most $\lg n$, which is also the entropy of the uniform distribution over a set of size n. If the anonymity is (significantly) below $\lg n$, then not much about the individual signing probabilities can be inferred. However, if the anonymity reaches $\lg n$, then the signing probability of each signer in the ring is exactly $1/n$. The optimality of the anonymity is in this sense informative. Interestingly, in the formulation of [25] optimal global anonymity also implies optimal local anonymity.

Later in this work, we will show that for 1-signer distributions $\mathcal{S}$ the optimal anonymity is always almost achievable. More concretely, in Section 6.2 we show that there exists a "mimicking" sampler Π_{Mimic} which achieves anonymity $\alpha(\mathsf{S}, \Pi_{\mathsf{Mimic}}) \gtrapprox \frac{1}{2} \lg n$, which is only a constant fraction away from the optimum, assuming minimally that S has at least $\lg n$ bits of min-entropy.

Although the mimicking sampler achieves near-optimal anonymity, it is mostly theoretical as it requires the knowledge of the distribution $\mathcal{S}$. More realistically, with a mild assumption that the support of S can be partitioned into chunks of size at least n, such that the signing probabilities of the signers within a chunk are similar, then the partitioning sampler presented in Section 6.3 also achieves near-optimal anonymity.

4.2.3 Extensions

In the following we discuss natural extensions of our anonymity definition, and why we decide not to incorporate them into our main definition.

"Local" Anonymity In some sense, the value $\alpha[\mathsf{S}, \Pi, \Lambda] = H_\infty(\mathsf{S}|\Pi(\mathsf{S}), \Lambda(\mathsf{S}))$ captures the "global" anonymity of all participants as a whole. To capture the "local" anonymity of a certain subset $I \subseteq [N]$ of users, one might want to consider the value $H_\infty(\mathsf{S}_I|\Pi(\mathsf{S}), \Lambda(\mathsf{S}))$, where $\mathsf{S}_I := \mathsf{S} \cap I$. We argue that however the value $H_\infty(\mathsf{S}_I|\Pi(\mathsf{S}), \Lambda(\mathsf{S}))$ does not capture the intuitive anonymity enjoyed by the subset I of users.

For a counter-argument, it suffices to consider the case where Λ is constant, $|\mathsf{S}| \equiv 1$ and $I = \{i\}$ for some $i \in [N]$. Note that S_I is a Boolean random variable (with

support $\{\emptyset, \{i\}\}$). Recall that

$$H_\infty(\mathsf{S}_I|\Pi(\mathsf{S}), \Lambda(\mathsf{S})) = H_\infty(\mathsf{S}_I|\Pi(\mathsf{S}))$$
$$= -\lg\left(\sum_R \mathsf{Pr}\left[\Pi(\mathsf{S}) = R\right] \max_{S_I} \mathsf{Pr}\left[\mathsf{S}_I = S_I|\Pi(\mathsf{S}) = R\right]\right).$$

Note that for any $R \not\ni i$, which are the majority,

$$\mathsf{Pr}\left[\mathsf{S}_I = \emptyset|\Pi(\mathsf{S}) = R\right] = 1,$$

and therefore

$$\max_{S_I} \mathsf{Pr}\left[\mathsf{S}_I = S_I|\Pi(\mathsf{S}) = R\right] = 1.$$

Therefore, intuitively, the expected value of $\max_{S_I} \mathsf{Pr}\left[\mathsf{S}_I = S_I|\Pi(\mathsf{S})\right]$ is close to 1, which means the conditional min-entropy $H_\infty(\mathsf{S}_I|\Pi(\mathsf{S}))$ is close to 0, even for the "best" samplers.

The above issue was due to the fact that user i is almost always not in the ring, and therefore an adversary could be successful by always guessing that user i is not a signer, *i.e.*, $\mathsf{S}_I = \emptyset$. An attempt to avoid this issue is to consider the entropy of S_I conditioning on R_I, where the latter is distributed as $\Pi(\mathsf{S})$ conditioned on that $I \subseteq \Pi(\mathsf{S})$. We examine the value

$$H_\infty(\mathsf{S}_I|\mathsf{R}_I)$$
$$= -\lg\left(\sum_R \mathsf{Pr}\left[\mathsf{R}_I = R\right] \max_{S_I} \mathsf{Pr}\left[\mathsf{S}_I = S_I|\mathsf{R}_I = R\right]\right)$$

for a hypothetical "best" sampler with a fixed ring size n, where for every ring $R \in \mathsf{Supp}(\mathsf{R})$ such that $S \subseteq R$, it holds that $\mathsf{Pr}\left[\mathsf{S} = S|\Pi(\mathsf{S}) = R\right] = 1/n$ (note that we are assuming that $|\mathsf{S}| \equiv 1$).

One would have hoped that the value is close to or exactly 1, which is the highest (min-)entropy that a Boolean random variable can have. However, note that in particular $\mathsf{Pr}\left[\mathsf{S} = \{i\}|\Pi(\mathsf{S}) = R\right] = 1/n$ in the case $I = \{i\} \subseteq R$, and hence $\mathsf{Pr}\left[\mathsf{S}_I = \{i\}|\mathsf{R}_I = R\right] = \frac{1}{n}$ for any $R \in \mathsf{Supp}(\mathsf{R}_I)$. We therefore have

$$\max_{S_I} \mathsf{Pr}\left[\mathsf{S}_I = S_I|\mathsf{R}_I = R\right] = \mathsf{Pr}\left[\mathsf{S}_I = \emptyset|\mathsf{R}_I = R\right] = \frac{n-1}{n}$$

for all $R \in \mathsf{Supp}(\mathsf{R}_I)$, and hence $H_\infty(\mathsf{S}_I|\mathsf{R}_I) = \lg n - \lg(n-1)$ (≈ 0 for large n), which is still counter-intuitive.

Anonymity "Over Time" Our main definition captures the remaining anonymity of the users in the view of an adversary after seeing a single ring. In reality, however, multiple rings would be sampled throughout the lifetime of the system, possibly even via different ring samplers, which might collectively leak more information about the signers (behind each ring) than any single ring does. For the ease of exposition we omit the leakage Λ in the discussion below.

Formally, suppose the system has been run for t time steps, *i.e.*, t rings have been sampled. For time step $i \in [t]$, let $N_i \in \mathbb{N}$ be the universe size, $\mathcal{S}_i$ be the signer distribution over the universe $[N_i]$, Π_i be the ring sampler, and $\mathsf{R}_i = \Pi_i(\mathsf{S}_i)$ be the random variable denoting the sampled ring. Then, for any subset $\{i_1, \ldots, i_\ell\} \subseteq [t]$, we might want to consider the value $H_\infty(\mathsf{S}_{i_1}, \ldots, \mathsf{S}_{i_\ell}|\mathsf{R}_1, \ldots, \mathsf{R}_t)$ which captures the anonymity of the signers at time steps $i \in \{i_1, \ldots, i_\ell\}$, after seeing the rings from all time steps.[11] In particular, the extreme values $H_\infty(\mathsf{S}_1, \ldots, \mathsf{S}_t|\mathsf{R}_1, \ldots, \mathsf{R}_t)$ and $H_\infty(\mathsf{S}_j|\mathsf{R}_1, \ldots, \mathsf{R}_t)$ for $j \in [t]$ might be of interest.

It is not difficult to show that

$$\max_{j\in[t]} H_\infty(\mathsf{S}_j|\mathsf{R}_1, \ldots, \mathsf{R}_t) \le H_\infty(\mathsf{S}_1, \ldots, \mathsf{S}_t|\mathsf{R}_1, \ldots, \mathsf{R}_t)$$
$$\le \sum_{j\in[t]} H_\infty(\mathsf{S}_j|\mathsf{R}_1, \ldots, \mathsf{R}_t)$$
$$\le \sum_{j\in[t]} H_\infty(\mathsf{S}_j|\mathsf{R}_j),$$

which relates the aforementioned extreme values with our definition of anonymity. Unfortunately, not much more can be said about these values in general since, for any $i \neq j$, (S_i, Π_i) and (S_j, Π_j) can be arbitrarily correlated depending on the application and user behavior.

For example, if (S_i, Π_i) and (S_j, Π_j) are independent for all $i \neq j$, then $t \cdot \max_{j\in[t]} H_\infty(\mathsf{S}_j|\mathsf{R}_1, \ldots, \mathsf{R}_t) = H_\infty(\mathsf{S}_1, \ldots, \mathsf{S}_t|\mathsf{R}_1, \ldots, \mathsf{R}_t)$, and the last two inequalities become equalities. On the other extreme, if (S_i, Π_i) and (S_j, Π_j) are identical and dependent for all i, j, then the first inequality becomes an equality, while $t \cdot H_\infty(\mathsf{S}_1, \ldots, \mathsf{S}_t|\mathsf{R}_1, \ldots, \mathsf{R}_t) = \sum_{j\in[t]} H_\infty(\mathsf{S}_j|\mathsf{R}_1, \ldots, \mathsf{R}_t)$.

In summary, the values $H_\infty(\mathsf{S}_{i_1}, \ldots, \mathsf{S}_{i_\ell}|\mathsf{R}_1, \ldots, \mathsf{R}_t)$ are extremely sensitive to the correlations between (S_i, Π_i) and (S_j, Π_j) for $i \neq j$, which highly depend on the real-world application and user behavior. Therefore, in a general theory about ring samplers where minimal assumptions about the signer distributions and user behavior are made, not much can be said about the "anonymity over time" meaningfully.

11 To account for leakages, we further condition the min-entropy on $\Lambda_1(\mathsf{S}_1), \ldots, \Lambda_t(\mathsf{S}_t)$ for possibly different leakage functions $\Lambda_1, \ldots, \Lambda_t$.

5 Robustness

We show that our anonymity definition is robust in the sense that, if the source distributions $\mathcal{S}$ and $\mathcal{S}'$ are close (in Rényi divergence), then the anonymity of a ring sampler with respect to S is close to that with respect to S'. This allows us to analyze ring samplers with respect to some distribution $\mathcal{S}$ which is easier to deal with, and get a guarantee of the anonymity of the sampler with respect to the real distribution $\mathcal{S}'$ assuming that it is close enough to $\mathcal{S}$.

Robustness also allows us to reason about the anonymity of a ring sampler against active attackers who attempt to perturb the signer distribution from $\mathcal{S}$ to $\mathcal{S}'$. Suppose we have deduced that the anonymity of a ring sampler with respect to $\mathcal{S}$ (against a passive adversary) is high. Assuming that no adversary could influence $\mathcal{S}$ too much, *i.e.*, $\mathcal{S}'$ is not too far away from $\mathcal{S}$, then by robustness the anonymity of the ring sampler with respect to $\mathcal{S}'$ is also high. Such an assumption could be realistic, *e.g.*, in the cryptocurrency setting where we anyway assume that the majority of the users are honest for the consensus protocol to function.

Theorem 5.1 (Robustness)**.** *For any* S *and* S' *with* $\mathsf{Supp}(\mathsf{S}) \subseteq \mathsf{Supp}(\mathsf{S}')$*, any* Π *and* Λ*, and any* $\varepsilon \geq 0$*, if* $D_\infty(\mathsf{S} \| \mathsf{S}') \leq \varepsilon$*, then*

$$\alpha(\mathsf{S}, \Pi, \Lambda) \geq \alpha(\mathsf{S}', \Pi, \Lambda) - \varepsilon.$$

6 Analysis of Natural Samplers

We formalize the the uniform, mimicking, and partitioning samplers, and analyze their anonymity. To understand the fundamental strengths and weaknesses of the samplers, we focus on 1-signer distributions, assume that the ring samplers are not subverted, and that no side-channel leakage is present. In cases where a close form of anonymity is unavailable, we provide lower bounds.

6.1 Uniform Samplers

A natural (yet generally bad) way to select rings is to just sample them uniformly at random. Formally, for each $1 \leq k \leq n \leq N$, we define the uniform sampler $\Pi_{\mathrm{Rand},k,n}$ as follows:

$\Pi_{\mathrm{Rand},k,n}(S \subseteq_{\leq k} [N])$: Sample $R \subseteq_n [N]$ uniformly at random subject to $S \subseteq R$.

Theorem 6.1 (Uniform Sampler)**.** *Let* $\mathcal{S}$ *be a* 1*-signer distribution,* E_i *be the* i*-th most probable event in* S *and*

$$\rho_i = \begin{cases} \Pr[E_i] & i \in [|\mathsf{Supp}(\mathsf{S})|] \\ 0 & i \in [N] \setminus [|\mathsf{Supp}(\mathsf{S})|]. \end{cases}$$

Then

$$\alpha(\mathsf{S}, \Pi_{\mathrm{Rand},1,n}) = -\lg \left(\frac{\sum_{i=n-1}^{N-1} \binom{i}{n-1} \rho_{N-i}}{\binom{N-1}{n-1}} \right). \quad (1)$$

Let us say a few words about the uniform sampler. Suppose $\mathcal{S}$ is the uniform distribution over $[N]$, we have $\rho_i = 1/N$ for all $i \in [N]$. Then by the "hockey-stick" identity, we have $\alpha(\mathsf{S}, \Pi_{\mathrm{Rand},1,n}) = \lg n$ which is optimal. This aligns with our expectation that when $\mathcal{S}$ is uniform the best way to sample a ring is to just sample uniformly.

Next we examine the scenario where $\mathcal{S}$ is far from uniform. For example, suppose that $\rho_i = 2\rho_{i-1}$ for all i. In this case, $\sum_{i=n-1}^{N-1} \binom{i}{n-1} \rho_{N-i}$ is dominated by the first few terms as ρ_i diminishes exponentially as i decreases. We can therefore expect that $\alpha(\mathsf{S}, \Pi_{\mathrm{Rand},1,n})$ is far from $\lg n$.

6.2 Mimicking Samplers

Another natural strategy of ring sampling is to mimic the true source distribution $\mathcal{S}$. Suppose that $\hat{\mathcal{S}}$ is an estimate of the true source distribution and is efficiently sampleable. We formalize this strategy as the $\hat{\mathcal{S}}$-mimicking sampler $\Pi^{\hat{\mathcal{S}}}_{\mathsf{Mimic},k,n}$ with size parameter (k, n) as follows:

$\Pi^{\hat{\mathcal{S}}}_{\mathsf{Mimic},k,n}(S \subseteq_k [N])$: Let $S_1 := S$. For $i \in [n] \setminus \{1\}$, sample $S_i \leftarrow\!\!\$\, \hat{\mathcal{S}}$. Output $R := \bigcup_{i \in [n]} S_i$.

Note that $\Pi^{\hat{\mathcal{S}}}_{\mathsf{Mimic},k,n}$ is a (k, kn)-ring sampler. In case $\hat{\mathcal{S}} = \mathcal{S}$ and $\mathcal{S}$ is a k-signer distribution, we write $\Pi^{\mathcal{S}}_{\mathsf{Mimic},k,n}$ as $\Pi_{\mathsf{Mimic},k,n}$ and call it the mimicking sampler.

We remark that $\Pi^{\hat{\mathcal{S}}}_{\mathsf{Mimic},k,n}$ is defined as above for easier analysis. It does not always produce a ring of size kn due to collisions from sampling with replacement, *i.e.*, it might happen that $S_i \cap S_j \neq \emptyset$ for $i \neq j$. Therefore the anonymity of $\Pi^{\hat{\mathcal{S}}}_{\mathsf{Mimic},k,n}$ cannot be optimal among all kn-ring samplers. The anonymity can only increase by

padding the ring to contain kn users with any strategy independent of S. In the special case where $k = 1$, one can continue to populate the ring with samples from $\hat{\mathcal{S}}$, until the ring size reaches n.

Despite the above suboptimality, in the case $N \gg n$, sampling with replacement is a reasonable approximation of sampling without replacement. It is therefore reasonable to expect that if the mimicking sampler has access to the true source distribution $\mathcal{S}$, its anonymity should be close to optimal. In the following, we give an evidence that this is the case.

To facilitate the analysis of $\Pi_{\mathsf{Mimic},k,n}$, we define a very similar algorithm $\overline{\Pi}_{\mathsf{Mimic},k,n}$ which treats the S_i's as multisets (sets with possibly repeated elements) and replaces the union operation with multiset sum[12]:

> $\overline{\Pi}_{\mathsf{Mimic},k,n}(S \subseteq_{\le k} [N])$: Let $S_1 := S$. For $i \in [n]\backslash\{1\}$, sample $S_i \leftarrow\!\!\$\ \mathcal{S}$. Output $X := \sum_{i\in[n]} S_i$.

Clearly, $\Pi_{\mathsf{Mimic},k,n}$ is a function of $\overline{\Pi}_{\mathsf{Mimic},k,n}$ (which removes all duplicated elements from the latter). Furthermore, let $\vec{x} \in \mathbb{N}_0^N$ be the characteristic vector of X. Then, if $\mathcal{S}$ is a 1-signer distribution, then the characteristic vectors $\vec{\mathsf{x}}$ of $\overline{\Pi}_{\mathsf{Mimic},k,n}(\mathsf{S})$ have a multinomial distribution with weights given by $\mathcal{S}$.

Theorem 6.2 (Mimicking Sampler)**.** *Let $\mathcal{S}$ be a 1-signer distribution. Let $\vec{\mathsf{x}} = (\mathsf{x}_i)_{i=1}^N$ be the characteristic vector of $\overline{\Pi}_{\mathsf{Mimic},1,n}(\mathsf{S})$.*

$$\alpha(\mathsf{S}, \Pi_{\mathsf{Mimic},1,n}) \ge \lg n - \lg \mathbb{E}[\max_i \mathsf{x}_i]. \tag{2}$$

Furthermore, assuming that $H_\infty(\mathsf{S}) \ge \lg n$, we have

$$\alpha(\mathsf{S}, \Pi_{\mathsf{Mimic},1,n}) \ge \lg(\sqrt{n} - 1) \approx \frac{\lg n}{2}. \tag{3}$$

Our proof of the theorem in Appendix C uses a bound of Aven [1], which is loose in some cases as it does not take into account the correlations between random variables. Nevertheless, we are able to show a non-trivial lower bound of (roughly) $\frac{1}{2}\lg n$, which is only a constant fraction away from optimal.

We emphasize that although Theorem 6.2 shows that the optimal anonymity is always almost achievable up to a constant factor, the result is mostly of theoretical interest, because it requires the knowledge of an estimation $\hat{\mathcal{S}}$ of the signer distribution $\mathcal{S}$. Even if it is possible to obtain a reasonable estimation $\hat{\mathcal{S}}$ of $\mathcal{S}$, a questionable assumption, $\mathcal{S}$ may change over time, *e.g.*, due to economic bubbles and recessions, and depends on the free will of users. For a good and practical sampler we recommend the partitioning sampler in Section 6.3.

Remark 2. *Attentive readers might observe the following peculiar phenomenon: Suppose the real signer happens to be Alice who has very low signing probability according to $\mathcal{S}$. It is likely that the mimicking sampler produces a ring in which all members except Alice have high signing probabilities, making Alice stand out. This is paradoxical since the mimicking sampler is close to optimal.*

The answer to the riddle is that the sampled ring could be, with similar probability (not the same due to potential collision), the result of someone else in the ring being the real signer, and picking Alice as a ring member.

With the same reasoning, the mimicking sampler naturally resists timing attacks described in Section 2.1, which assumes that the signing probability of a signer depends on its age (c.f. Section 7.1). Indeed, the event that a young signer ending up in the ring could be with similar (high) probabilities the result of him being the signer or him being chosen by another signer.

6.3 Partitioning Samplers

Another natural idea for ring sampling is to put signers with similar signing probabilities into the same ring. We first abstract this idea as the family of partitioning samplers. We then propose a practical partitioning strategy which also provides other security features.

6.3.1 Abstract Description

A set P of sets (called chunks) is said to be a partition of $[N]$ if $\bigcup_{C\in P} C = [N]$, $C \cap C' = \emptyset$ for all $C, C' \in P$ with $C \ne C'$, and $C \ne \emptyset$ for all $C \in P$. Fix a size parameter $n \in [N]$. Let $\mathcal{P}$ be a distribution over the partitions of the set of signers $[N]$ where each chunk is of size at least n. Intuitively, the partitions P chosen by a partitioning sampler should have support containing only partitions where signers in each chunk have similar signing probabilities. We will only use this assumption in the anonymity analysis but not in the construction: The construction works for all distributions of partitions.

Given any distribution $\mathcal{P}$ of partitions, size parameters k and (optionally) n, we define the partitioning

12 For example, $\{a, a, b, c\} + \{b, c, c\} = \{a, a, b, b, c, c, c\}$.

sampler $\Pi_{\mathrm{Part},\mathsf{P},k,n}$ ($\Pi_{\mathrm{Part},\mathsf{P},k}$ if n is not given) as follows:

> $\Pi_{\mathrm{Part},\mathcal{P},k,n}(S \subseteq_{\leq k} [N])$: Let $P \leftarrow\!\!\$\ \mathcal{P}$ be a partition of $[N]$. For each $s \in S$, let $C_s \in P$ be the unique chunk such that $s \in C_s$. Sample $R_s \subseteq_n C_s$ uniformly subject to $s \in R_s$. (Note that we assumed $|C| \geq n$ for all $C \in P$ for all $P \in \mathsf{Supp}(\mathsf{P})$.) Output $R := \bigcup_{s\in S} R_s$.

> $\Pi_{\mathrm{Part},\mathcal{P},k}(S \subseteq_{\leq k} [N])$: Let $P \leftarrow\!\!\$\ \mathcal{P}$ be a partition of $[N]$. For each $s \in S$, let $C_s \in P$ be the unique chunk such that $s \in C_s$. Output $R := \bigcup_{s\in S} C_s$.

Clearly $\Pi_{\mathrm{Part},\mathcal{P},k,n}$ is a (k, kn)-ring sampler. Due to the potential collision of chunks, *i.e.*, there exist distinct $s, s' \in S$ such that $s, s' \in C$ for some $C \in P$, the partitioning sampler cannot be optimal with respect to k-signer distributions where $k > 1$. Although collisions can be made rare if the partition is fine-grained and random enough, the anonymity can only increase by padding the ring to size kn, similar to our suggestion for the mimicking sampler.

We analyze the anonymity of $\Pi_{\mathrm{Part},\mathcal{P},1,n}$ and $\Pi_{\mathrm{Part},\mathcal{P},1}$ with respect to any 1-signer distribution S. We start with the simple case where the support of P is a singleton, *i.e.*, $\mathsf{P} \equiv P$ for some partition P of $[N]$.

Theorem 6.3 (Partitioning Sampler)**.** *Let $\mathcal{S}$ be a 1-signer distribution. Let $n \in [N]$. Let $\mathsf{P} \equiv P$ for some partition P of $[N]$ such that $|C| \geq n$ for all $C \in P$. For each $C \in P$, let μ_C be the mean of $\mathsf{Pr}[\mathsf{S} = \{s\}]$ over all $s \in C$, i.e., $\mu_C := |C|^{-1} \sum_{s\in C} \mathsf{Pr}[\mathsf{S} = \{s\}]$. Suppose that for all $C \in P$, all $s \in C$, it holds that $|\mathsf{Pr}[\mathsf{S} = \{s\}] - \mu_C| \leq \varepsilon_C$ for some $\varepsilon_C \geq 0$. Let $\varepsilon_P := \sum_{C\in P} |C|\varepsilon_C$. Then*

$$\alpha(\mathsf{S}, \Pi_{\mathrm{Part},\mathcal{P},1,n}) \geq \lg n - \lg(\varepsilon_P + 1)$$

and

$$\alpha(\mathsf{S}, \Pi_{\mathrm{Part},\mathcal{P},1}) \geq \lg n - \lg(\varepsilon_P + 1).$$

We next show that the anonymity can only be better with a larger support of P, condition on that all partitions in the support of P satisfy the above constraints.

Corollary 6.1. *Let $\mathcal{S}$ be a 1-signer distribution and $n \in [N]$. Suppose for each partition P in the support of P, for all $C \in P$, all $s \in C$, it holds that $|\mathsf{Pr}[\mathsf{S} = \{s\}] - \mu_C| \leq \varepsilon_C$ for some $\varepsilon_C \geq 0$, and $|C| \geq n$. Let $\varepsilon_P := \sum_{C\in P} |C|\varepsilon_C$ and let $\varepsilon_{\mathsf{P}} := \sum_P \mathsf{Pr}[\mathsf{P} = P]\,\varepsilon_P$. Then*

$$\alpha(\mathsf{S}, \Pi_{\mathrm{Part},\mathcal{P},1,n}) \geq \lg n - \lg(\varepsilon_{\mathsf{P}} + 1)$$

and

$$\alpha(\mathsf{S}, \Pi_{\mathrm{Part},\mathcal{P},1}) \geq \lg n - \lg(\varepsilon_{\mathsf{P}} + 1).$$

If the size parameter n is given, we observe that if all signers in a chunk have identical signing probabilities, then $\varepsilon_{\mathsf{P}} = 0$ and the anonymity is optimal, *i.e.*, $\lg n$.

6.3.2 Suggested Instantiations

We suggest concrete strategies for partitioning the universe of signers in a realistic cryptocurrency setting.

In the simple case where the signers can be clustered into chunks according to signing probabilities, such that each chunk is of size $\geq n$ and consists of signers with the same signing probability, the collection of these chunks form a natural partition P which satisfies the conditions in Theorem 6.3 with $\varepsilon_P = 0$. The partitioning sampler $\Pi_{\mathrm{Part},\mathcal{P},1,n}$ (with n given) therefore achieves optimal anonymity. The above requirements can be met, *e.g.*, when we assume that the signing probability depends only on the "age" of the signer, and there are enough signers of the same age.[13] For a more detailed discussion about age, we refer to Section 7.

The above requires to partition the universe such that each chunk is of size $\geq n$. In reality it could happen that some chunks are of size $< n$. Taking Monero as an example, if we consider all output accounts in the same (blockchain) block to have the same age, and assume that the signing probability depends on the age, then there are on average around 13 accounts [14] in one such chunk, which is insufficient for a ring size of $n > 13$. To resolve this issue, a natural approach (Approach 1) is to group several chunks into a bigger chunk, such that the latter is of size $\geq n$. Assuming that the signing probability of signers in consecutive chunks are similar, the resulting value of ε_P would still be quite close to 0, and hence the partitioning sampler is still close to optimal.

Another issue is about the anonymity of the partitioning sampler after several rings sampled by signers in the same chunk are observed. In the ideal case, where each chunk is of size exactly n (as in the suggested ring sampler of [25]), no extra information about the signers can be extracted even after seeing multiple rings – the

13 Depending on the coarseness of the definition of "age", signers (*e.g.*, accounts in a cryptocurrency) of the same age might not spawn simultaneously. In such case the ring sampler should treat as if the youngest signers are not in the universe until all of them have spawned. Also, the youngest signers should wait until all their fellows have spawned before signing.

14 See Table 3 for detailed numbers

signers are essentially running the "all sampler" treating the chunk as the universe. In reality however where the chunk size is often greater than n, graph analysis can potentially be performed to extract non-trivial information about the signers, especially when the sampler is supplied with bad randomness or even subverted (as in the black marble attacks described in Section 2.1).

To avoid the potential risk of graph analysis, an idea is to enforce the chunk size of n. Assuming that the universe size $N = \ell \cdot T$ for some ℓ which is a multiple of n,[15] and assuming signers that are close in age have similar signing probabilities, we partition $[N]$ in the following recursive manner (Approach 2). Let $N' = \ell(T-1)$. Suppose that the subset $[N']$ of signers were already partitioned. Immediately after the universe size advances from N' to $N = N' + \ell$, we partition the ℓ new signers into chunks each of size n uniformly at random using public randomness (*e.g.*, derived by hashing the state of the blockchain up to current time). Unioning these chunks with the original partition of $[N']$ gives a partition of $[N]$. Once the partition is sampled and fixed, the ring sampler is deterministic. The case $\ell = n$ coincides with the sampler suggested in [25].

Below, we highlight some interesting properties of our instantiation, the first two of which are outside our model for anonymity.

Obliviousness to Signer Distribution Unlike the mimicking sampler, the partitioning sampler is oblivious to the real signer distribution and does not require knowledge of a close estimate of it. This provides an easy way to create partitions for any universe as long as the assumptions about similar probabilities are met, *i.e.* we always can set ℓ to the next multiple of n greater or equal the mean block size.

Trade-off between Waiting Time and Anonymity Both suggested approaches above require the younger signers to wait until enough of them have spawned to be able to use the sampler. The waiting time increases with the ring size n and hence with anonymity. For reference, in Table 3 we report the average waiting time until ℓ accounts have spawned in Monero for some values of ℓ.

Security against Deanonymization Attacks Having near-optimal anonymity (with respect to reasonable signer distributions), our partitioning sampler instantiation is more secure against the deanonymization attacks mentioned in Section 2.1, *e.g.*, than the uniform sampler, according to Theorem 4.4.

The passive security can be seen intuitively. Exploitation of transaction times is prevented as the ring is fixed as soon as the whole chunk of accounts is available. Graph structure analysis is confined as the induced bipartite graph now consists of disconnected subgraphs, each corresponding to a chunk. In the extreme case where each chunk is of size n (as suggested in [25]), each subgraph is balanced and complete, hence no information can be inferred from graph analysis. Correlation between output accounts of the same transaction is not useful, since the number of signer is restricted to 1.

For active security (of the signer-distribution-influencing kind), *e.g.*, against black marble attacks, Theorem 5.1 guarantees that the partitioning sampler is near optimal with respect to a slightly tempered signer distribution. We remark however that the near-optimality is in a global sense (c.f., Section 4.2.3). Locally, a targetted black marble attack in which the attacker introduces a large number of black marbles within a single chunk of a partition may cause rings sampled from that chunk to contain mostly black marbles, hence provide little anonymity. A simple mitigation strategy is to enforce a large chunk size, so an effective attack becomes expensive. Another idea is to use a hybrid between the mimicking and partitioning sampler – first sample a subring using the mimicking sampler, and then run the partitioning sampler on each member of the subring.

7 Empirical Evaluation

We study empirically the anonymity of the uniform and mimicking samplers with respect to several signer distributions. We skip the partitioning sampler as it is optimal for all distributions that we consider (with the appropriate parameters).

7.1 Signer Distributions

Uniform Distribution As a reference, we first consider the uniform distribution $\mathsf{U}_{[N]}$ over $[N]$. $\mathsf{U}_{[N]}$ is the easiest to build a good ring sampler for, in the sense that the simple uniform sampler is optimal for $\mathsf{U}_{[N]}$. While $\mathsf{U}_{[N]}$ is unrealistic in the cryptocurrency context, it might decently model the reality in "one-shot" ap-

15 If not, then as before we treat the youngest signers as if they were not part of the universe, until there are enough of them to make the universe size a multiple of n.

plications of ring samplers, *e.g.*, secret disclosure, especially when not much side-channel information is known about the potential signers by the adversary.

Monero Distribution To obtain more realistic distributions, Möser et al. [12] analyzed the empirical distribution of the age of transaction outputs/accounts. The age here refers to the difference between the spent time and the creation time of a transaction output/account (measured in blocks). While this information is supposedly hidden in a privacy-preserving cryptocurrency such as Monero, Möser et al. [12] analyzed the transaction graph of Monero in the pre-RingCT era, and successfully deanonymized a lot of transactions. For these deanonymized transactions, Möser et al. [12] "heuristically determined" that the logarithm of the age of accounts matches a gamma distribution. In our terminology, we call such an age distribution the "log-gamma" distribution, which has the PDF

$$\mathsf{Pr}\left[\mathsf{age} = t\right] \propto (\ln t)^{a-1} t^{-b}$$

for some shape parameter $a > 0$ and rate parameter $b > 0$, and has support $\mathsf{Supp}(\mathsf{age}) = (0, \infty)$. The parameters of the log-gamma distribution fitted by Möser et al. [12] are $a = 19.28$ and $b = 1.61$ respectively.

Subsequently, the log-gamma distribution is used in the ring sampler of Monero in the following way. First, an age is sampled from the log-gamma distribution. Rejection sampling is employed so that age ≤ 10 blocks are rejected. Then, an account is chosen uniformly at random from all accounts having the sampled age. This process is repeated until the ring is populated to a desired size. This can be viewed as an S-mimicking sampler, where the age of S has log-gamma distribution, and signers of the same age have equal signing probability.

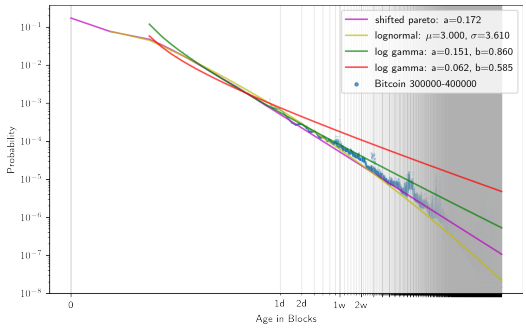

Fig. 1. Empirical Bitcoin age distribution in blocks and fitted PDFs

Baseline Distribution Modeled after Bitcoin Since the graph analysis tools used by Möser et al. [12] are not publicly available, we could not replicate their results for Monero. Nevertheless, we re-examine the age distribution of Bitcoin transaction outputs created within the 300,000-400,000 block period. In Figure 1 is a log-log plot of the probability density functions (PDFs) of the empirical age distribution of Bitcoin, a fitted log-normal distribution, a fitted (shifted) Pareto distribution, and two fitted log-gamma distributions. Pale and dark vertical lines mark days and weeks respectively.

The log-normal and Pareto distributions are chosen because the log-log plot of the age distribution looks almost like a straight line. The log-gamma distribution is include since it was the distribution of choice of Möser et al. [12]. The log-normal distribution has the PDF

$$\mathsf{Pr}\left[\mathsf{age} = t\right] \propto t^{-\left(\frac{\ln t - 2\mu}{2\sigma^2} + 1\right)}$$

for some parameters $\mu, \sigma > 0$. For a fixed shift of 1 (to shift the support from $[1, \infty)$ to $[0, \infty)$), the (shifted) Pareto distribution has the PDF

$$\mathsf{Pr}\left[\mathsf{age} = t\right] \propto (t+1)^{-(a+1)}$$

for some shape parameter $a > 0$.

In Table 2 we summarize the fitting range, parameters, and the root-mean-square error (RMSE) of the fitted distributions. We emphasize that only the (shifted) Pareto distribution has the correct support, *i.e.*, $[0, \infty)$, while the support of log-normal and log-gamma is $(1, \infty)$. For this reason, although the log-gamma distribution fitted to the range $[10 : 197394]$ has the lowest RMSE, it is not necessarily better than the others since the magnitude of the probability decreases rapidly in t.

Distribution	Fitting Range	Parameters	RMSE
Log-normal	$[1 : 197394]$	$(\mu, \sigma) = (3.000, 3.610)$	2.271×10^{-5}
(Shifted) Pareto	$[1 : 197394]$	$a = 0.172$	1.987×10^{-5}
Log-gamma	$[2 : 197394]$	$(a, b) = (0.062, 0.585)$	4.580×10^{-5}
Log-gamma	$[10 : 197394]$	$(a, b) = (0.151, 0.860)$	4.542×10^{-6}

Table 2. Parameters of fitted distributions

While irrelevant to this work, the periodicity of the Bitcoin distribution is interesting – the local maxima align with the daily marks. This phenomenon is even clearer when the age is measured in minutes (Figure 4) or seconds, where unfortunately some transaction outputs appear to have negative age due to the variation of system time in different machines. It seems difficult to de-noise the data and fitting distributions to noisy data seems less meaningful.

Based on the observation on Bitcoin data, we propose to use the discretization of (shifted) Pareto distributions, *i.e.*, (shifted) zeta distributions, as a baseline for signer distributions. More precisely, our baseline family is parameterized by $(\ell, T, a) \in \mathbb{N}^2 \times (0, \infty)$, where ℓ is the number of signers of the same age, T is the size of the age range, and a is the parameter of the (shifted) zeta distribution. The universe size is $N = \ell T$. For $i \in [N]$, $\mathsf{Pr}[\mathsf{S} = \{i\}] \propto (t+1)^{-(a+1)}$ where $t = \left\lfloor \frac{i-1}{\ell} \right\rfloor$ is the age of the signer.

7.2 Evaluation Results

Uniform Sampler Figure 2 plots the anonymity of the uniform sampler (Equation (1)) with respect to the uniform distribution $\mathcal{U}_{[42198964]}$ (blue) and the baseline distributions, with $a = 0.172$, $\ell = 1024$ and different values of T, against ring size in linear-log scale. We also plotted the upper bound $\lg n$ (orange). We observed that the anonymity for the uniform distribution is independent of the actual universe size. For the baseline distribution, the anonymity is independent of ℓ. As shown by the overlapping blue and orange lines in Figure 2, the uniform sampler is optimal for the uniform distribution. The anonymity with respect to the baseline distributions drifts away from the optimum as T increases.

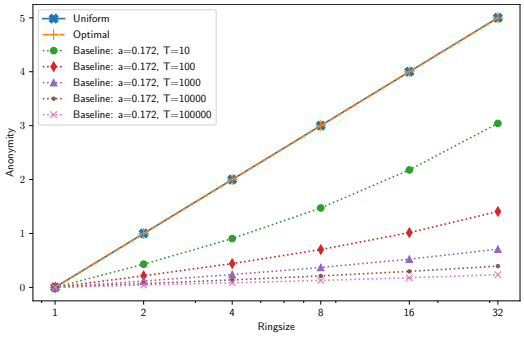

Fig. 2. Selected results for anonymity of the uniform sampler, full data is available in Table 4

Mimicking Sampler Figure 3 plots the lower bound of the anonymity of the mimicking sampler (Inequality (2)) with respect to the uniform distributions $\mathcal{U}_{[N]}$, where $N \in \{160, 6400, 10240000\}$, and the baseline distributions, with $a = 0.172$ and $(\ell, T) \in \{16, 64, 1024\} \times \{100, 10000\}$, against ring size in linear-log scale. We also plotted the upper bound $\lg n$ and the global lower bound $\lg(\sqrt{n} - 1)$ (Inequality (3)). To evaluate the term $\mathbb{E}[\max_i \mathsf{x}_i]$ in Inequality (2), we have implemented the algorithm in [3].

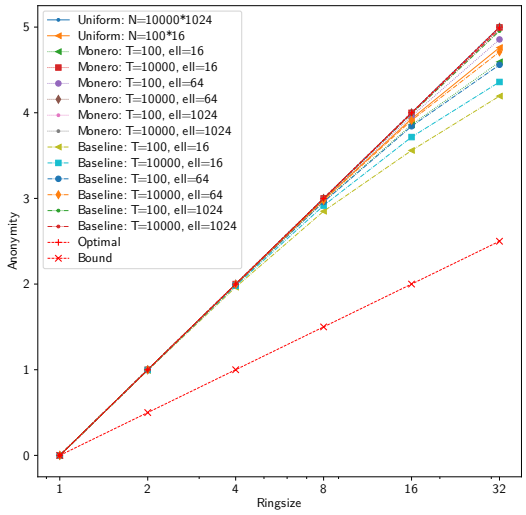

Fig. 3. Selected results for anonymity of the mimicking sampler. Full data is available in Tables 5 and 7

Figure 3 shows that for a uniform distribution the mimicking sampler is nearly optimal while getting even closer for larger N. For the baseline (with $a = 0.172$) and Monero distributions, the anonymity approaches to the optimum as ℓ and T increases, with the effect of ℓ being much more significant.

Partitioning Sampler We remind that the partitioning sampler achieves the optimal anonymity of $\lg n$ as long as each chunk in each possible partition has size at least n and contains users with equal signing probability. Both assumptions are satisfied by the uniform distribution and the baseline distribution for $\ell \geq n$.

Acknowledgments

This work is supported by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) as part of the Research and Training Group 2475 “Cybercrime and Forensic Computing” (number 393541319/GRK2475/1-2019) and by the state of Bavaria at the Nuremberg Campus of Technology (NCT).

References

[1] T. Aven. "Upper (Lower) Bounds on the Mean of the Maximum (Minimum) of a Number of Random Variables." In: *Journal of Applied Probability* 22.3 (1985), pp. 723–728. ISSN: 00219002. URL: http://www.jstor.org/stable/3213876.

[2] M. Backes, S. Meiser, and M. Slowik. "Your Choice MATor(s)." In: *PoPETs* 2016.2 (Apr. 2016), pp. 40–60.

[3] C. J. Corrado. "The Exact Distribution of the Maximum, Minimum and the Range of Multinomial/Dirichlet and Multivariate Hypergeometric Frequencies." In: *Statistics and Computing* 21.3 (July 2011), 349–359. ISSN: 0960-3174. DOI: 10.1007/s11222-010-9174-3. URL: https://doi.org/10.1007/s11222-010-9174-3.

[4] Y. Deng, J. Pang, and P. Wu. "Measuring Anonymity with Relative Entropy." In: *FAST 2006*. Ed. by T. Dimitrakos, F. Martinelli, P. Y. A. Ryan, and S. A. Schneider. Vol. 4691. Lecture Notes in Computer Science. Springer, 2006, pp. 65–79. DOI: 10.1007/978-3-540-75227-1_5. URL: https://doi.org/10.1007/978-3-540-75227-1_5.

[5] C. Díaz, S. Seys, J. Claessens, and B. Preneel. "Towards Measuring Anonymity." In: *PET 2002*. Ed. by R. Dingledine and P. F. Syverson. Vol. 2482. LNCS. Springer, Heidelberg, Apr. 2002, pp. 54–68. DOI: 10.1007/3-540-36467-6_5.

[6] Y. Dodis, R. Ostrovsky, L. Reyzin, and A. Smith. "Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other Noisy Data." In: *SIAM J. Comput.* 38.1 (Mar. 2008), pp. 97–139. ISSN: 0097-5397. DOI: 10.1137/060651380.

[7] M. Edman, F. Sivrikaya, and B. Yener. "A Combinatorial Approach to Measuring Anonymity." In: *ISI 2007*. IEEE, 2007, pp. 356–363. DOI: 10.1109/ISI.2007.379497. URL: https://doi.org/10.1109/ISI.2007.379497.

[8] J. Geddes, R. Jansen, and N. Hopper. "How Low Can You Go: Balancing Performance with Anonymity in Tor." In: *PETS 2013*. Ed. by E. De Cristofaro and M. K. Wright. Vol. 7981. LNCS. Springer, Heidelberg, July 2013, pp. 164–184. DOI: 10.1007/978-3-642-39077-7_9.

[9] A. Kumar, C. Fischer, S. Tople, and P. Saxena. "A Traceability Analysis of Monero's Blockchain." In: *ESORICS 2017, Part II*. Ed. by S. N. Foley, D. Gollmann, and E. Snekkenes. Vol. 10493. LNCS. Springer, Heidelberg, Sept. 2017, pp. 153–173. DOI: 10.1007/978-3-319-66399-9_9.

[10] J. K. Liu, V. K. Wei, and D. S. Wong. "Linkable Spontaneous Anonymous Group Signature for Ad Hoc Groups (Extended Abstract)." In: *ACISP 04*. Ed. by H. Wang, J. Pieprzyk, and V. Varadharajan. Vol. 3108. LNCS. Springer, Heidelberg, July 2004, pp. 325–335. DOI: 10.1007/978-3-540-27800-9_28.

[11] A. Mackenzie, S. Noether, and M. C. Team. *Improving Obfuscation in the CryptoNote Protocol*. Tech. rep. URL: https://ww.getmonero.org/resources/research-lab/pubs/MRL-0004.pdf.

[12] M. Möser, K. Soska, E. Heilman, K. Lee, H. Heffan, S. Srivastava, K. Hogan, J. Hennessey, A. Miller, A. Narayanan, and N. Christin. "An Empirical Analysis of Traceability in the Monero Blockchain." In: *PoPETs* 2018.3 (July 2018), pp. 143–163. DOI: 10.1515/popets-2018-0025.

[13] R. E. Newman, I. S. Moskowitz, P. F. Syverson, and A. Serjantov. "Metrics for Trafic Analysis Prevention." In: *PET 2003*. Ed. by R. Dingledine. Vol. 2760. LNCS. Springer, Heidelberg, Mar. 2003, pp. 48–65. DOI: 10.1007/978-3-540-40956-4_4.

[14] S. Noether, A. Mackenzie, and the Monero Research Lab. "Ring Confidential Transactions." In: *Ledger* 1 (2016), pp. 1–18. ISSN: 2379-5980. DOI: 10.5195/ledger.2016.34.

[15] S. Noether, S. Noether, and A. Mackenzie. *A Note on Chain Reactions in Traceability in CryptoNote 2.0*. Tech. rep. URL: https://ww.getmonero.org/resources/research-lab/pubs/MRL-0001.pdf.

[16] R. L. Rivest, A. Shamir, and Y. Tauman. "How to Leak a Secret." In: *ASIACRYPT 2001*. Ed. by C. Boyd. Vol. 2248. LNCS. Springer, Heidelberg, Dec. 2001, pp. 552–565. DOI: 10.1007/3-540-45682-1_32.

[17] A. Serjantov and G. Danezis. "Towards an Information Theoretic Metric for Anonymity." In: *PET 2002*. Ed. by R. Dingledine and P. F. Syverson. Vol. 2482. LNCS. Springer, Heidelberg, Apr. 2002, pp. 41–53. DOI: 10.1007/3-540-36467-6_4.

[18] R. Shokri, G. Theodorakopoulos, G. Danezis, J.-P. Hubaux, and J.-Y. Le Boudec. "Quantifying Location Privacy: The Case of Sporadic Location Exposure." In: *PETS 2011*. Ed. by S. Fischer-Hübner and N. Hopper. Vol. 6794. LNCS. Springer, Heidelberg, July 2011, pp. 57–76. DOI: 10.1007/978-3-642-22263-4_4.

[19] R. Shokri, G. Theodorakopoulos, J.-Y. Le Boudec, and J.-P. Hubaux. "Quantifying Location Pri-

vacy." In: *2011 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, May 2011, pp. 247–262. DOI: 10.1109/SP.2011.18.

[20] G. Tóth and Z. Hornák. "Measuring Anonymity in a Non-adaptive, Real-Time System." In: *PET 2004*. Ed. by D. M. Martin Jr. and A. Serjantov. Vol. 3424. LNCS. Springer, Heidelberg, May 2004, pp. 226–241. DOI: 10.1007/11423409_14.

[21] F. Tramèr, D. Boneh, and K. Paterson. "Remote Side-Channel Attacks on Anonymous Transactions." In: *USENIX Security 2020*. Ed. by S. Capkun and F. Roesner. USENIX Association, Aug. 2020, pp. 2739–2756.

[22] I. Wagner and D. Eckhoff. "Technical Privacy Metrics: A Systematic Survey." In: *ACM Comput. Surv.* 51.3 (June 2018). ISSN: 0360-0300. DOI: 10.1145/3168389. URL: https://doi.org/10.1145/3168389.

[23] D. A. Wijaya, J. Liu, R. Steinfeld, and D. Liu. "Monero Ring Attack: Recreating Zero Mixin Transaction Effect." In: *TrustCom/BigDataSE 2018*. IEEE, 2018, pp. 1196–1201. DOI: 10.1109/TrustCom/BigDataSE.2018.00165.

[24] S. Xu and M. Yung. "Accountable ring signatures: a smart card approach." In: *Smart Card Research and Advanced Applications VI*. Springer, 2004, pp. 271–286.

[25] J. Yu, M. H. A. Au, and P. Esteves-Verissimo. "Rethinking untraceability in the CryptoNote-style blockchain – The Sun Tzu survival problem." In: *CSF 2019*. Ed. by S. Delaune and L. Jia. United States of America: IEEE, 2019, pp. 94–107.

[26] L. von Ahn and N. J. Hopper. "Public-Key Steganography." In: *EUROCRYPT 2004*. Ed. by C. Cachin and J. Camenisch. Vol. 3027. LNCS. Springer, Heidelberg, May 2004, pp. 323–341. DOI: 10.1007/978-3-540-24676-3_20.

A Waiting Time in Monero

We perform an empirical analysis of the waiting time in Monero and report the findings in Table 3. Monero creates a new block every two minutes. For this evaluation we considered a transaction log from Monero starting on February 21st 2019 [16] and collected 36,000 blocks. Currently, the blocked time that a user has to wait in Monero in order to spend from its account is around 20 minutes.

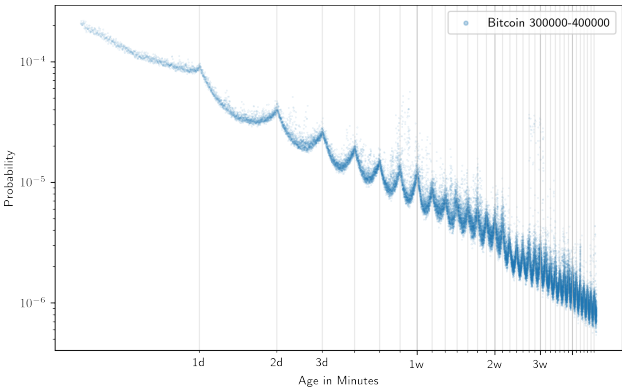

Fig. 4. Empirical Bitcoin age distribution in minutes

B Implication to Ring Signatures

We define a generalized version of ring signatures which aim to capture its different variants. We also define a simulation-based notion of anonymity, which is equivalent to the classic indistinguishability-based notion, but synergizes better with the notion of ring samplers.

Definition B.1 (Ring Signatures)**.** *A ring signature scheme* Σ *is a tuple of* PPT *algorithms* $(\mathsf{KGen}, \mathsf{Sig}, \mathsf{Vf})$ *with the following syntax:*

$(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KGen}(1^\lambda)$*: The key generation algorithm generates a public verification key* pk *and a secret signing key* sk*.*

$\sigma \leftarrow \mathsf{Sig}\left(\{\mathsf{pk}_i\}_{i=1}^n, \{\mathsf{sk}_i\}_{i\in I}, m\right)$*: The sign algorithm inputs a set of public keys* $\{\mathsf{pk}_i\}_{i=1}^n$ *called the ring, a set of secret keys* $\{\mathsf{sk}_i\}_{i\in I}$ *corresponding to* pk_i *for* $i \in I$ *for some* $I \subseteq [n]$*, and a message* $m \in \mathcal{M}$ *for some message space* $\mathcal{M}$*. It outputs a signature* σ*.*

$b \leftarrow \mathsf{Vf}\left(\{\mathsf{pk}_i\}_{i=1}^n, m, \sigma\right)$*: The verify algorithm inputs a ring* $\{\mathsf{pk}_i\}_{i=1}^n$*, a message* m*, and a signature* σ*. It outputs a bit* b *deciding if* σ *is a valid signature of* m *with respect to the ring.*

Σ *is correct if for any* $\lambda, N \in \mathbb{N}$*, any* $J \subseteq I \subseteq_n [N]$*, any* $(\mathsf{pk}_i, \mathsf{sk}_i) \in \mathsf{KGen}(1^\lambda)$ *for* $i \in [N]$*, any* $m \in \mathcal{M}$*, and any* $\sigma \in \mathsf{Sig}\left(\{\mathsf{pk}_i\}_{i\in I}, \{\mathsf{sk}_i\}_{i\in J}, m\right)$*, it holds that* $\mathsf{Vf}\left(\{\mathsf{pk}_i\}_{i\in I}, m, \sigma\right) = 1$*.*

16 Corresponding to Blockhash

size ℓ	50	100	150	200	250	300	350	400	450	500	550	600
mean	3.83	7.66	11.49	15.32	19.15	22.98	26.81	30.64	34.47	38.30	42.13	45.96
stdev	2.76	4.40	5.80	7.17	8.40	9.55	10.79	11.91	13.06	14.22	15.41	16.16

Table 3. Mean waiting time (in blocks) until ℓ new accounts have been accumulated

Distribution	T	ℓ	n=1	n=2	n=4	n=8	n=16	n=32
Baseline	100	16	$7.37 \cdot 10^{-15}$	$2.15 \cdot 10^{-1}$	$4.40 \cdot 10^{-1}$	$6.99 \cdot 10^{-1}$	$1.01 \cdot 10^{0}$	$1.40 \cdot 10^{0}$
Baseline	100	64	$-4.16 \cdot 10^{-15}$	$2.15 \cdot 10^{-1}$	$4.41 \cdot 10^{-1}$	$7.01 \cdot 10^{-1}$	$1.01 \cdot 10^{0}$	$1.40 \cdot 10^{0}$
Baseline	100	256	$1.32 \cdot 10^{-13}$	$2.15 \cdot 10^{-1}$	$4.41 \cdot 10^{-1}$	$7.01 \cdot 10^{-1}$	$1.02 \cdot 10^{0}$	$1.41 \cdot 10^{0}$
Baseline	100	1,024	$-6.54 \cdot 10^{-13}$	$2.15 \cdot 10^{-1}$	$4.41 \cdot 10^{-1}$	$7.01 \cdot 10^{-1}$	$1.02 \cdot 10^{0}$	$1.41 \cdot 10^{0}$
Baseline	1,000	16	$2.64 \cdot 10^{-14}$	$1.18 \cdot 10^{-1}$	$2.37 \cdot 10^{-1}$	$3.69 \cdot 10^{-1}$	$5.23 \cdot 10^{-1}$	$7.08 \cdot 10^{-1}$
Baseline	1,000	64	$-8.33 \cdot 10^{-15}$	$1.18 \cdot 10^{-1}$	$2.37 \cdot 10^{-1}$	$3.69 \cdot 10^{-1}$	$5.23 \cdot 10^{-1}$	$7.08 \cdot 10^{-1}$
Baseline	1,000	256	$2.40 \cdot 10^{-13}$	$1.18 \cdot 10^{-1}$	$2.37 \cdot 10^{-1}$	$3.69 \cdot 10^{-1}$	$5.23 \cdot 10^{-1}$	$7.09 \cdot 10^{-1}$
Baseline	1,000	1,024	$1.62 \cdot 10^{-12}$	$1.18 \cdot 10^{-1}$	$2.37 \cdot 10^{-1}$	$3.69 \cdot 10^{-1}$	$5.23 \cdot 10^{-1}$	$7.09 \cdot 10^{-1}$
Baseline	10,000	16	$9.90 \cdot 10^{-14}$	$7.00 \cdot 10^{-2}$	$1.39 \cdot 10^{-1}$	$2.12 \cdot 10^{-1}$	$2.96 \cdot 10^{-1}$	$3.94 \cdot 10^{-1}$
Baseline	10,000	64	$1.71 \cdot 10^{-13}$	$7.00 \cdot 10^{-2}$	$1.39 \cdot 10^{-1}$	$2.12 \cdot 10^{-1}$	$2.96 \cdot 10^{-1}$	$3.94 \cdot 10^{-1}$
Baseline	10,000	256	$-3.24 \cdot 10^{-13}$	$7.00 \cdot 10^{-2}$	$1.39 \cdot 10^{-1}$	$2.12 \cdot 10^{-1}$	$2.96 \cdot 10^{-1}$	$3.94 \cdot 10^{-1}$
Baseline	10,000	1,024	$-1.28 \cdot 10^{-12}$	$7.00 \cdot 10^{-2}$	$1.39 \cdot 10^{-1}$	$2.12 \cdot 10^{-1}$	$2.96 \cdot 10^{-1}$	$3.94 \cdot 10^{-1}$
Baseline	100,000	16	$2.44 \cdot 10^{-13}$	$4.35 \cdot 10^{-2}$	$8.53 \cdot 10^{-2}$	$1.30 \cdot 10^{-1}$	$1.79 \cdot 10^{-1}$	$2.35 \cdot 10^{-1}$
Baseline	100,000	64	$9.80 \cdot 10^{-13}$	$4.35 \cdot 10^{-2}$	$8.53 \cdot 10^{-2}$	$1.30 \cdot 10^{-1}$	$1.79 \cdot 10^{-1}$	$2.35 \cdot 10^{-1}$
Baseline	100,000	256	$-5.27 \cdot 10^{-12}$	$4.35 \cdot 10^{-2}$	$8.53 \cdot 10^{-2}$	$1.30 \cdot 10^{-1}$	$1.79 \cdot 10^{-1}$	$2.35 \cdot 10^{-1}$
Baseline	100,000	1,024	$4.29 \cdot 10^{-12}$	$4.35 \cdot 10^{-2}$	$8.53 \cdot 10^{-2}$	$1.30 \cdot 10^{-1}$	$1.79 \cdot 10^{-1}$	$2.35 \cdot 10^{-1}$
Uniform			$0.00 \cdot 10^{0}$	$1.00 \cdot 10^{0}$	$2.00 \cdot 10^{0}$	$3.00 \cdot 10^{0}$	$4.00 \cdot 10^{0}$	$5.00 \cdot 10^{0}$

Table 4. Anonymity for Uniform Sampler

Let Λ be a (possibly probabilistic and stateful) leakage function, and $\varepsilon = \varepsilon(\lambda) \in [0,1]$. Σ is (Λ,ε)-anonymous if for any $N \in \mathbb{N}$ and any stateful PPT *adversary $\mathcal{A}$, there exists a stateful* PPT *simulator $\mathcal{S}$, such that*

$$\left| \begin{matrix} \mathsf{Pr}\left[\mathsf{Real}_{\Sigma,N,\mathcal{A}}(1^\lambda) = 1\right] \\ -\mathsf{Pr}\left[\mathsf{Ideal}_{\Sigma,\Lambda,N,\mathcal{A},\mathcal{S}}(1^\lambda) = 1\right] \end{matrix} \right| < \varepsilon$$

where the experiments $\mathsf{Real}_{\Sigma,N,\mathcal{A}}$ and $\mathsf{Ideal}_{\Sigma,\Lambda,N,\mathcal{A},\mathcal{S}}$ are defined in Figure 5.

Typically, ε is considered to be a negligible function in λ. It is easy to see that the above simulation-based definition is equivalent to the classical indistinguishability-based definition. If Σ is simulation-based anonymous, then the indistinguishability-based anonymity can be proven by a standard hybrid argument where we hop from the "0" experiment to the ideal experiment and then to the "1" experiment. Conversely, if Σ is indistinguishability-based anonymous, then we can construct a simulator who generates signatures by running the sign algorithm on an appropriate subset of keys.

The classic ring signatures corresponds to the setting where the number of signers is $|I| = n \equiv 1$ and there is no leakage, *e.g.*, Λ always return the empty string. To capture, for example, linkable ring signatures, the leakage function Λ should reveal which members of J have issued a signature before (due to linkability) and the carnality of J (due to the number of linkability tags).

We formally capture the concrete anonymity of a ring signature scheme coupled with a ring sampler with respect to a signer distribution as follows.

Definition B.2. *Let $\delta > 0$. Let Σ be a ring signature scheme, Π be a ring sampler, and* S *be a signer distribution. The compound system (Σ,Π) is said to have concrete anonymity δ (the smaller the better) with respect to* S*, if for all* PPT *adversary, it holds that*

$$\mathsf{Pr}\left[\mathsf{ConcreteAnon}_{\Sigma,\Pi,\mathsf{S},\mathcal{A}}(1^\lambda) = 1\right] \leq \delta$$

where $\mathsf{ConcreteAnon}_{\Sigma,\Pi,S,\mathcal{A}}$ is defined in Figure 6.

In the following theorem, we relate the concrete anonymity of (Σ,Π) to the anonymity of Σ and Π.

Distribution	T	ℓ	n=1	n=2	n=4	n=8	n=16	n=32
Baseline	100	16	$0.00 \cdot 10^0$	$9.94 \cdot 10^{-1}$	$1.96 \cdot 10^0$	$2.85 \cdot 10^0$	$3.56 \cdot 10^0$	$4.19 \cdot 10^0$
Baseline	1,000	16	$0.00 \cdot 10^0$	$9.96 \cdot 10^{-1}$	$1.98 \cdot 10^0$	$2.90 \cdot 10^0$	$3.66 \cdot 10^0$	$4.30 \cdot 10^0$
Baseline	10,000	16	$0.00 \cdot 10^0$	$9.97 \cdot 10^{-1}$	$1.98 \cdot 10^0$	$2.92 \cdot 10^0$	$3.72 \cdot 10^0$	$4.36 \cdot 10^0$
Baseline	100	64	$0.00 \cdot 10^0$	$9.98 \cdot 10^{-1}$	$1.99 \cdot 10^0$	$2.96 \cdot 10^0$	$3.84 \cdot 10^0$	$4.56 \cdot 10^0$
Baseline	1,000	64	$0.00 \cdot 10^0$	$9.99 \cdot 10^{-1}$	$1.99 \cdot 10^0$	$2.97 \cdot 10^0$	$3.89 \cdot 10^0$	$4.66 \cdot 10^0$
Baseline	10,000	64	$0.00 \cdot 10^0$	$9.99 \cdot 10^{-1}$	$2.00 \cdot 10^0$	$2.98 \cdot 10^0$	$3.91 \cdot 10^0$	$4.71 \cdot 10^0$
Baseline	100	256	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$2.99 \cdot 10^0$	$3.96 \cdot 10^0$	$4.84 \cdot 10^0$
Baseline	1,000	256	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$2.99 \cdot 10^0$	$3.97 \cdot 10^0$	$4.89 \cdot 10^0$
Baseline	10,000	256	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$2.99 \cdot 10^0$	$3.98 \cdot 10^0$	$4.91 \cdot 10^0$
Baseline	100	1,024	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$3.99 \cdot 10^0$	$4.95 \cdot 10^0$
Baseline	1,000	1,024	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$3.99 \cdot 10^0$	$4.97 \cdot 10^0$
Baseline	10,000	1,024	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$3.99 \cdot 10^0$	$4.98 \cdot 10^0$

Table 5. Anonymity for Mimicking Sampler and Baseline Distribution

Distribution	T	ℓ	n=1	n=2	n=4	n=8	n=16	n=32
Monero	100	16	$0.00 \cdot 10^0$	$9.99 \cdot 10^{-1}$	$1.99 \cdot 10^0$	$2.96 \cdot 10^0$	$3.86 \cdot 10^0$	$4.59 \cdot 10^0$
Monero	1,000	16	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$3.99 \cdot 10^0$	$4.96 \cdot 10^0$
Monero	10,000	16	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$4.00 \cdot 10^0$	$5.00 \cdot 10^0$
Monero	100	64	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$2.99 \cdot 10^0$	$3.96 \cdot 10^0$	$4.86 \cdot 10^0$
Monero	1,000	64	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$4.00 \cdot 10^0$	$4.99 \cdot 10^0$
Monero	10,000	64	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$4.00 \cdot 10^0$	$5.00 \cdot 10^0$
Monero	100	256	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$3.99 \cdot 10^0$	$4.96 \cdot 10^0$
Monero	1,000	256	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$4.00 \cdot 10^0$	$5.00 \cdot 10^0$
Monero	10,000	256	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$4.00 \cdot 10^0$	$5.00 \cdot 10^0$
Monero	100	1,024	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$4.00 \cdot 10^0$	$4.99 \cdot 10^0$
Monero	1,000	1,024	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$4.00 \cdot 10^0$	$5.00 \cdot 10^0$
Monero	10,000	1,024	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$4.00 \cdot 10^0$	$5.00 \cdot 10^0$

Table 6. Anonymity for Mimicking Sampler and Monero Distribution

Theorem B.3. *Let Σ be an (Λ, ε)-anonymous ring signature scheme. Let Π be a ring sampler. Let S be a signer distribution. Let $\delta := |\mathsf{Supp}(\Lambda(\mathsf{S}))| \cdot 2^{-\alpha[\Pi,\mathsf{S}]} + \varepsilon$. Then the system (Σ, Π) has concrete anonymity δ with respect to S.*

Proof. It suffices to prove that

$$\begin{aligned}&\Pr\left[\mathsf{ConcreteAnon}_{\Sigma,\Pi,\mathsf{S},\mathcal{A}}(1^\lambda) = 1\right]\\ \leq& 2^{-H_\infty(\mathsf{S}|\Pi(\mathsf{S}),\Lambda(\mathsf{S}))} + \varepsilon.\end{aligned}$$

The theorem then follows immediately from the chain rule of min-entropy. For the above inequality, we can rewrite it as

$$\begin{aligned}&\Pr\left[\mathsf{ConcreteAnon}_{\Sigma,\Pi,\mathsf{S},\mathcal{A}}(1^\lambda) = 1\right]\\ \leq& \mathsf{Guess}\left(\mathsf{S}|\Pi(\mathsf{S}),\Lambda(\mathsf{S})\right) + \varepsilon.\end{aligned}$$

Since Σ is (Λ, ε)-anonymous, there exists a PPT simulator $\mathcal{S}$, such that

$$\left|\begin{matrix}\Pr\left[\mathsf{Real}_{\Sigma,N,\mathcal{A}}(1^\lambda) = 1\right]\\ -\Pr\left[\mathsf{Ideal}_{\Sigma,\Lambda,N,\mathcal{A},\mathcal{S}}(1^\lambda) = 1\right]\end{matrix}\right| < \varepsilon.$$

We therefore consider a modified experiment $\mathsf{ConcreteAnon}'_{\Sigma,\Pi,\mathsf{S},\mathcal{A},\mathcal{S}}$ which is almost identical to $\mathsf{ConcreteAnon}_{\Sigma,\Pi,\mathsf{S},\mathcal{A}}$, except that the keys pk_i are produced by $\mathcal{S}(1^\lambda)$, and the signature σ is produced by $\mathcal{S}(R, \Lambda(S), m)$. Due to the (Λ, ε)-anonymity of Σ, we have

$$\left|\begin{matrix}\Pr\left[\mathsf{ConcreteAnon}_{\Sigma,\Pi,\mathsf{S},\mathcal{A}}(1^\lambda) = 1\right]\\ -\Pr\left[\mathsf{ConcreteAnon}'_{\Sigma,\Pi,\mathsf{S},\mathcal{A},\mathcal{S}}(1^\lambda) = 1\right]\end{matrix}\right| < \varepsilon.$$

Note that in $\mathsf{ConcreteAnon}'_{\Sigma,\Pi,\mathsf{S},\mathcal{A},\mathcal{S}}$ the only information about S available to $\mathcal{A}$ is $R = \Pi(S)$ and

Distribution	T	ℓ	n=1	n=2	n=4	n=8	n=16	n=32
Uniform	100	16	$0.00 \cdot 10^0$	$9.99 \cdot 10^{-1}$	$2.00 \cdot 10^0$	$2.98 \cdot 10^0$	$3.93 \cdot 10^0$	$4.76 \cdot 10^0$
Uniform	1,000	16	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$3.99 \cdot 10^0$	$4.97 \cdot 10^0$
Uniform	10,000	16	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$4.00 \cdot 10^0$	$5.00 \cdot 10^0$
Uniform	100,000	16	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$4.00 \cdot 10^0$	$5.00 \cdot 10^0$
Uniform	100	64	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$3.98 \cdot 10^0$	$4.93 \cdot 10^0$
Uniform	1,000	64	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$4.00 \cdot 10^0$	$4.99 \cdot 10^0$
Uniform	10,000	64	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$4.00 \cdot 10^0$	$5.00 \cdot 10^0$
Uniform	100,000	64	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$4.00 \cdot 10^0$	$5.00 \cdot 10^0$
Uniform	100	256	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$4.00 \cdot 10^0$	$4.98 \cdot 10^0$
Uniform	1,000	256	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$4.00 \cdot 10^0$	$5.00 \cdot 10^0$
Uniform	10,000	256	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$4.00 \cdot 10^0$	$5.00 \cdot 10^0$
Uniform	100,000	256	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$4.00 \cdot 10^0$	$5.00 \cdot 10^0$
Uniform	100	1,024	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$4.00 \cdot 10^0$	$5.00 \cdot 10^0$
Uniform	1,000	1,024	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$4.00 \cdot 10^0$	$5.00 \cdot 10^0$
Uniform	10,000	1,024	$0.00 \cdot 10^0$	$1.00 \cdot 10^0$	$2.00 \cdot 10^0$	$3.00 \cdot 10^0$	$4.00 \cdot 10^0$	$5.00 \cdot 10^0$

Table 7. Anonymity for Mimicking Sampler and Uniform Distribution

$\underline{\mathsf{Real}_{\Sigma,N,\mathcal{A}}(1^\lambda)}$
$(\mathsf{pk}_i, \mathsf{sk}_i) \leftarrow \mathsf{KGen}(1^\lambda)\, \forall\, i \in [N]$
$b \leftarrow \mathcal{A}^{\mathsf{Sig}\mathcal{O}_{\mathsf{Real}}}\left(\{\mathsf{pk}_i\}_{i=1}^N\right)$
return b

$\underline{\mathsf{Ideal}_{\Sigma,\Lambda,N,\mathcal{A},\mathcal{S}}(1^\lambda)}$
$\{\mathsf{pk}_i\}_{i=1}^N \leftarrow \mathcal{S}(1^\lambda)$
$b \leftarrow \mathcal{A}^{\mathsf{Sig}\mathcal{O}_{\mathsf{Ideal}}}\left(\{\mathsf{pk}_i\}_{i=1}^N\right)$
return b

$\underline{\mathsf{Sig}\mathcal{O}_{\mathsf{Real}}(I, J, m)}$
$\sigma \leftarrow \mathsf{Sig}\left(\{\mathsf{pk}_i\}_{i\in I}, \{\mathsf{sk}_i\}_{i\in J}, m\right)$
return σ

$\underline{\mathsf{Sig}\mathcal{O}_{\mathsf{Ideal}}(I, J, m)}$
$\sigma \leftarrow \mathcal{S}\left(I, \Lambda(J), m\right)$
return σ

Fig. 5. Anonymity experiments of ring signatures

$\underline{\mathsf{ConcreteAnon}_{\Sigma,\Pi,\mathsf{S},\mathcal{A}}(1^\lambda)}$
$(\mathsf{pk}_i, \mathsf{sk}_i) \leftarrow \mathsf{KGen}(1^\lambda)\, \forall\, i \in [N]$
$m \leftarrow \mathcal{A}\left(\{\mathsf{pk}_i\}_{i=1}^N\right)$
$S \leftarrow\!\!\$\, \mathsf{S}$
$R \leftarrow \Pi(S)$
$\sigma \leftarrow \mathsf{Sig}\left(\{\mathsf{pk}_i\}_{i\in R}, \{\mathsf{sk}_i\}_{i\in S}, m\right)$
$S^* \leftarrow \mathcal{A}(R, \sigma)$
return $(S = S^*)$

Fig. 6. Concrete anonymity experiments of ring signatures

$L = \Lambda(S)$. We therefore have

$$\begin{aligned}&\mathsf{Pr}\left[\mathsf{ConcreteAnon}'_{\Sigma,\Pi,\mathsf{S},\mathcal{A},\mathcal{S}}(1^\lambda) = 1\right]\\ \leq\ &\mathsf{Guess}\left(\mathsf{S}|\Pi(\mathsf{S}), \Lambda(\mathsf{S})\right).\end{aligned}$$

The claim then follows. □

C Proofs

We restate and prove all lemmas and theorems whose proofs do not fit into the main body.

Lemma 4.1. *For any k-signer distribution S, any n-ring sampler Π, and any leakage function Λ,*

$$\alpha(\mathsf{S}, \Pi, \Lambda) \leq \lg \sum_{i=1}^{k} \binom{n}{i}.$$

In particular, for $k = 1$ we have

$$\alpha(\mathsf{S}, \Pi, \Lambda) \leq \lg n.$$

Proof. By monotonicity, $\alpha(\mathsf{S}, \Pi, \Lambda) \leq \alpha(\mathsf{S}, \Pi)$. It therefore suffices to show that, for each $R \subseteq_{\leq n} [N]$,

$$\max_S \mathsf{Pr}\left[\mathsf{S} = S | \Pi(\mathsf{S}) = R\right] \geq \binom{n}{k}^{-1}.$$

Fix $R \subseteq_{\leq n} [N]$. For all $S \in \mathsf{Supp}(\mathsf{S})$ and $S \not\subseteq R$, we have $\mathsf{Pr}\left[\mathsf{S} = S | \Pi(\mathsf{S}) = R\right] = 0$. Therefore

$$\sum_{S: S \in \mathsf{Supp}(\mathsf{S}) \wedge S \subseteq R} \mathsf{Pr}\left[\mathsf{S} = S | \Pi(\mathsf{S}) = R\right] = 1.$$

Note that

$$|\{S : S \in \mathsf{Supp}(\mathsf{S}) \wedge S \subseteq R\}| \leq \sum_{i=1}^{k} \binom{n}{i}.$$

The desired result follows from the pigeonhole principle. □

Lemma C.1. *For any* S *and* S' *with* $\mathsf{Supp}(\mathsf{S}) \subseteq \mathsf{Supp}(\mathsf{S}')$*, any (probabilistic) function* $\Phi : \{0,1\}^* \to \{0,1\}^*$*, and any* $\varepsilon \geq 0$*, if* $D_\infty(\mathsf{S}\|\mathsf{S}') \leq \varepsilon$*, then for any* $Y \in \mathsf{Supp}(\Phi(\mathsf{S}))$*,*

$$\max_{S \in \mathsf{Supp}(\mathsf{S})} \mathsf{Pr}\left[(\mathsf{S}, \Phi(\mathsf{S})) = (S, Y)\right] \leq 2^{\varepsilon} \cdot \max_{S \in \mathsf{Supp}(\mathsf{S}')} \mathsf{Pr}\left[(\mathsf{S}', \Phi(\mathsf{S}')) = (S, Y)\right].$$

Proof. By assumption, we have

$$\lg \max_{S \in \mathsf{Supp}(\mathsf{S}')} \frac{\mathsf{Pr}[\mathsf{S} = S]}{\mathsf{Pr}[\mathsf{S}' = S]} = D_\infty(\mathsf{S}\|\mathsf{S}') \leq \varepsilon$$
$$\frac{\mathsf{Pr}[\mathsf{S} = S']}{\mathsf{Pr}[\mathsf{S}' = S']} \leq \max_{S \in \mathsf{Supp}(\mathsf{S}')} \frac{\mathsf{Pr}[\mathsf{S} = S]}{\mathsf{Pr}[\mathsf{S}' = S]} \leq 2^{\varepsilon}$$

for all $S' \in \mathsf{Supp}(\mathsf{S}')$. Fix $Y \in \mathsf{Supp}(\Phi(\mathsf{S}))$. Let $S^* \in \mathsf{Supp}(\mathsf{S})$ be such that

$$\mathsf{Pr}\left[(\mathsf{S}, \Phi(\mathsf{S})) = (S^*, Y)\right] = \max_{S \in \mathsf{Supp}(S)} \mathsf{Pr}\left[(\mathsf{S}, \Phi(\mathsf{S})) = (S, Y)\right].$$

We have

$$\begin{aligned}
&\max_{S \in \mathsf{Supp}(\mathsf{S})} \mathsf{Pr}\left[(\mathsf{S}, \Phi(\mathsf{S})) = (S, Y)\right] \\
=&\mathsf{Pr}\left[(\mathsf{S}, \Phi(\mathsf{S})) = (S^*, Y)\right] \\
=&\mathsf{Pr}\left[(\Pi(S^*), \Lambda(S^*)) = Y\right] \mathsf{Pr}\left[\mathsf{S} = S^*\right] \\
\leq&2^{\varepsilon} \mathsf{Pr}\left[(\Pi(S^*), \Lambda(S^*)) = Y\right] \mathsf{Pr}\left[\mathsf{S}' = S^*\right] \\
=&2^{\varepsilon} \mathsf{Pr}\left[(\mathsf{S}', \Pi(\mathsf{S}'), \Lambda(\mathsf{S}')) = (S^*, Y)\right] \\
\leq&2^{\varepsilon} \max_{S \in \mathsf{Supp}(\mathsf{S}')} \mathsf{Pr}\left[(\mathsf{S}', \Pi(\mathsf{S}'), \Lambda(\mathsf{S}')) = (S, Y)\right].
\end{aligned}$$ □

Theorem 5.1 (Robustness)**.** *For any* S *and* S' *with* $\mathsf{Supp}(\mathsf{S}) \subseteq \mathsf{Supp}(\mathsf{S}')$*, any* Π *and* Λ*, and any* $\varepsilon \geq 0$*, if* $D_\infty(\mathsf{S}\|\mathsf{S}') \leq \varepsilon$*, then*

$$\alpha(\mathsf{S}, \Pi, \Lambda) \geq \alpha(\mathsf{S}', \Pi, \Lambda) - \varepsilon.$$

Proof. By Lemma C.1,

$$\max_{S \in \mathsf{Supp}(\mathsf{S})} \mathsf{Pr}\left[(\mathsf{S}, \Phi(\mathsf{S})) = (S, Y)\right] \leq 2^{\varepsilon} \cdot \max_{S \in \mathsf{Supp}(\mathsf{S}')} \mathsf{Pr}\left[(\mathsf{S}', \Phi(\mathsf{S}')) = (S, Y)\right].$$

for any (probabilistic) function Φ and any $Y \in \mathsf{Supp}(\Phi(\mathsf{S}))$. Therefore

$$\begin{aligned}
&2^{-H_\infty(\mathsf{S}|\Phi(\mathsf{S}))} \\
=&\sum_{Y \in \mathsf{Supp}(\Phi(\mathsf{S}))} \max_{S \in \mathsf{Supp}(\mathsf{S})} \mathsf{Pr}\left[(\mathsf{S}, \Pi(\mathsf{S})) = (S, Y)\right] \\
\leq&2^{\varepsilon} \sum_{Y \in \mathsf{Supp}(\Phi(\mathsf{S}))} \max_{S \in \mathsf{Supp}(\mathsf{S}')} \mathsf{Pr}\left[(\mathsf{S}', \Pi(\mathsf{S}')) = (S, Y)\right] \\
\leq&2^{\varepsilon} \sum_{Y \in \mathsf{Supp}(\Phi(\mathsf{S}'))} \max_{S \in \mathsf{Supp}(\mathsf{S}')} \mathsf{Pr}\left[(\mathsf{S}', \Pi(\mathsf{S}')) = (S, Y)\right] \\
=&2^{-(H_\infty(\mathsf{S}'|\Phi(\mathsf{S}')) - \varepsilon)}.
\end{aligned}$$

Recall that

$$\alpha(\mathsf{S}, \Pi, \Lambda) = H_\infty(\mathsf{S}|\Pi(\mathsf{S}), \Lambda(\mathsf{S})).$$

By setting $\Phi = (\Pi, \Lambda)$, we have the desired result. □

Theorem 6.1 (Uniform Sampler)**.** *Let* $\mathcal{S}$ *be a* 1*-signer distribution,* E_i *be the* i*-th most probable event in* S *and*

$$\rho_i = \begin{cases} \mathsf{Pr}[E_i] & i \in [|\mathsf{Supp}(\mathsf{S})|] \\ 0 & i \in [N] \setminus [|\mathsf{Supp}(\mathsf{S})|]. \end{cases}$$

Then

$$\alpha(\mathsf{S}, \Pi_{\mathrm{Rand},1,n}) = -\lg\left(\frac{\sum_{i=n-1}^{N-1} \binom{i}{n-1} \rho_{N-i}}{\binom{N-1}{n-1}}\right). \quad (1)$$

Proof. Let $\Pi = \Pi_{\mathrm{Rand},1,n}$. Recall that

$$\begin{aligned}
&\alpha(\mathsf{S}, \Pi) = H_\infty(\mathsf{S}|\Pi(\mathsf{S})) \\
=&-\lg\left(\sum_R \max_S \mathsf{Pr}\left[\Pi(\mathsf{S}) = R | \mathsf{S} = S\right] \mathsf{Pr}\left[\mathsf{S} = S\right]\right).
\end{aligned}$$

We examine the value $\mathsf{Pr}[\Pi(\mathsf{S}) = R|\mathsf{S} = S]$. Fix $R \subseteq [N]$. If $S \not\subseteq R$, we have $\mathsf{Pr}[\Pi(\mathsf{S}) = R|\mathsf{S} = S] = 0$. On the other hand, if $S \subseteq R$, we have $\mathsf{Pr}[\Pi(\mathsf{S}) = R|\mathsf{S} = S] = \binom{N-1}{n-1}^{-1}$ since there are $\binom{N-1}{n-1}$ many $R \supseteq S$.

Note that

$$\begin{aligned}
&\sum_R \max_S \mathsf{Pr}\left[\Pi(\mathsf{S}) = R | \mathsf{S} = S\right] \mathsf{Pr}\left[\mathsf{S} = S\right] \\
=&\sum_R \max_{S \in \mathsf{Supp}(\mathsf{S}) \wedge S \subseteq R} \mathsf{Pr}\left[\Pi(\mathsf{S}) = R | \mathsf{S} = S\right] \mathsf{Pr}\left[\mathsf{S} = S\right] \\
=&\binom{N-1}{n-1}^{-1} \sum_R \max_{S \in \mathsf{Supp}(\mathsf{S}) \wedge S \subseteq R} \mathsf{Pr}\left[\mathsf{S} = S\right]
\end{aligned}$$

Consider each term of the sum

$$\sum_R \max_{S \in \mathsf{Supp}(\mathsf{S}) \wedge S \subseteq R} \mathsf{Pr}\left[\mathsf{S} = S\right].$$

For each $i \in [|\mathsf{Supp}(\mathsf{S})|]$, let S_i be the i-th most probable event in S (note that $|\mathsf{S}| \equiv 1$). For $i \in [N] \setminus [|\mathsf{Supp}(\mathsf{S})|]$, let $S_i = \emptyset$. In either case we have $\Pr[\mathsf{S} = S_i] = \rho_i$. For each R, if $S_1 \subseteq R$, then we have $\max_{S \in \mathsf{Supp}(\mathsf{S}) \wedge S \subseteq R} \Pr[\mathsf{S} = S] = \rho_1$. Note that there are $\binom{N-1}{n-1}$ many such R. Else, if $S_1 \not\subseteq R$ but $S_2 \subseteq R$, then we have $\max_{S \in \mathsf{Supp}(\mathsf{S}) \wedge S \subseteq R} \Pr[\mathsf{S} = S] = \rho_2$. Continuing in this way, we can conclude that

$$\begin{aligned}
&\sum_R \max_{S \subseteq R} \Pr[\mathsf{S} = S] \\
=&\binom{N-1}{n-1}\rho_1 + \binom{N-2}{n-1}\rho_2 + \ldots + \binom{n-1}{n-1}\rho_{N-n+1} \\
=&\sum_{i=n-1}^{N-1} \binom{i}{n-1}\rho_{N-i}.
\end{aligned}$$

The theorem statement follows. □

Theorem 6.2 (Mimicking Sampler)**.** *Let $\mathcal{S}$ be a 1-signer distribution. Let $\vec{\mathsf{x}} = (\mathsf{x}_i)_{i=1}^N$ be the characteristic vector of $\overline{\Pi}_{\mathsf{Mimic},1,n}(\mathsf{S})$.*

$$\alpha(\mathsf{S}, \Pi_{\mathsf{Mimic},1,n}) \geq \lg n - \lg \mathbb{E}[\max_i \mathsf{x}_i]. \tag{2}$$

Furthermore, assuming that $H_\infty(\mathsf{S}) \geq \lg n$, we have

$$\alpha(\mathsf{S}, \Pi_{\mathsf{Mimic},1,n}) \geq \lg(\sqrt{n} - 1) \approx \frac{\lg n}{2}. \tag{3}$$

Proof. For each S and $\vec{x}$, note that if $S \not\subseteq X$, then $\Pr[\mathsf{S} = S \wedge \vec{\mathsf{x}} = \vec{x}] = 0$. On the other hand, suppose $S \subseteq X$. Since $\mathcal{S}$ is a 1-signer distribution, $S = \{i\}$ for some $i \in [N]$. We have

$$\Pr[\mathsf{S} = S \wedge \vec{\mathsf{x}} = \vec{x}] = \frac{x_i(n-1)!}{x_1! \ldots x_N!} \prod_{j \in [N]} \Pr[\mathsf{S} = \{j\}]^{x_j}$$

Since $\vec{\mathsf{x}}$ is multinomially distributed,

$$\Pr[\vec{\mathsf{x}} = \vec{x}] = \frac{n!}{x_1! \ldots x_N!} \prod_{j \in [N]} \Pr[\mathsf{S} = \{j\}]^{x_j}.$$

Therefore

$$\begin{aligned}
\Pr[\mathsf{S} = S | \vec{\mathsf{x}} = \vec{x}] =& \frac{\Pr[\mathsf{S} = S \wedge \vec{\mathsf{x}} = \vec{x}]}{\Pr[\vec{\mathsf{x}} = \vec{x}]} \\
=& \frac{x_i}{n} \\
\max_S \Pr[\mathsf{S} = S | \vec{\mathsf{x}} = \vec{x}] =& \frac{\max_i x_i}{n} \\
\sum_{\vec{x}} \Pr[\vec{\mathsf{x}} = \vec{x}] \max_S \Pr[\mathsf{S} = S | \vec{\mathsf{x}} = \vec{x}] =& \frac{\mathbb{E}[\max_i \mathsf{x}_i]}{n} \\
H_\infty(\mathsf{S}|\vec{\mathsf{x}}) =& \lg n - \lg \mathbb{E}[\max_i \mathsf{x}_i].
\end{aligned}$$

Recall that $\Pi_{\mathsf{Mimic},1,n}$ is a function of $\overline{\Pi}_{\mathsf{Mimic},1,n}$. By a data processing inequality (Lemma 3.3), we have

$$\begin{aligned}
\alpha(\mathsf{S}, \Pi_{\mathsf{Mimic},1,n}) =& H_\infty(\mathsf{S}|\mathsf{R}) \\
\geq& H_\infty(\mathsf{S}|\vec{\mathsf{x}}) \\
=& \lg n - \lg \mathbb{E}[\max_i \mathsf{x}_i].
\end{aligned}$$

Furthermore, assume that $H_\infty(\mathsf{S}) \geq \lg n$, or in other words $\max_S \Pr[\mathsf{S} = S] \leq 1/n$. We recall the upper bound

$$\mathbb{E}[\max_i \mathsf{x}_i] \leq \max_i \mu_i + \sqrt{\frac{N-1}{N} \sum_i \sigma_i^2}$$

by Aven [1], where $\mu_i = \mathbb{E}[\mathsf{x}_i]$ and $\sigma_i^2 = \mathrm{Var}[\mathsf{x}_i]$.

For $i \in [N]$, denote $p_i := \Pr[\mathsf{S} = i]$. Substituting $\mu_i = np_i \leq n \cdot 1/n = 1$ and $\sigma_i^2 = np_i(1-p_i)$, we have

$$\begin{aligned}
\mathbb{E}[\max_i \mathsf{x}_i] &\leq 1 + \sqrt{\frac{N-1}{N} \sum_i np_i(1-p_i)} \\
&= 1 + \sqrt{\frac{N-1}{N} \cdot n \cdot (1 - \sum_i p_i^2)} \\
&\leq 1 + \sqrt{\frac{N-1}{N} \cdot n \cdot (1 - 1/N)} \\
&= 1 + \frac{N-1}{N}\sqrt{n} \leq \sqrt{n} + 1
\end{aligned}$$

where in the second inequality we applied the Cauchy-Schwarz inequality on $(p_1, \ldots, p_N)$ and $(1, \ldots, 1)$.

Consequently, we have

$$\begin{aligned}
H_\infty(\mathsf{S}|\mathsf{R}) &\geq \lg n - \lg(\sqrt{n} + 1) = \lg \frac{n - 1 + 1}{\sqrt{n} + 1} \\
&= \lg(\sqrt{n} - 1 + \frac{1}{\sqrt{n} + 1}) \\
&\geq \lg(\sqrt{n} - 1)
\end{aligned}$$

□

Theorem 6.3 (Partitioning Sampler)**.** *Let $\mathcal{S}$ be a 1-signer distribution. Let $n \in [N]$. Let $\mathsf{P} \equiv P$ for some partition P of $[N]$ such that $|C| \geq n$ for all $C \in P$. For each $C \in P$, let μ_C be the mean of $\Pr[\mathsf{S} = \{s\}]$ over all $s \in C$, i.e., $\mu_C := |C|^{-1} \sum_{s \in C} \Pr[\mathsf{S} = \{s\}]$. Suppose that for all $C \in P$, all $s \in C$, it holds that $|\Pr[\mathsf{S} = \{s\}] - \mu_C| \leq \varepsilon_C$ for some $\varepsilon_C \geq 0$. Let $\varepsilon_P := \sum_{C \in P} |C| \varepsilon_C$. Then*

$$\alpha(\mathsf{S}, \Pi_{\mathrm{Part},\mathcal{P},1,n}) \geq \lg n - \lg(\varepsilon_P + 1)$$

and

$$\alpha(\mathsf{S}, \Pi_{\mathrm{Part},\mathcal{P},1}) \geq \lg n - \lg(\varepsilon_P + 1).$$

Proof. In the following, we write $\Pi = \Pi_{\mathrm{Part},\mathsf{P},1,n}$, $\Pi' = \Pi_{\mathrm{Part},\mathsf{P},1}$, $\mathsf{R} := \Pi(\mathsf{S})$, and $\mathsf{R}' := \Pi'(\mathsf{S})$. We first analyze the anonymity of Π.

For any $C \in P$, and any R we have

$$\mathsf{Pr}\left[\mathsf{R} = R | \mathsf{S} = \{s\}\right] = \begin{cases} \binom{|C|-1}{n-1}^{-1} & s \in R \subseteq_n C \\ 0 & \text{otherwise.} \end{cases}$$

Since for all $C \in P$, all $s \in C$, $|\mathsf{Pr}\left[\mathsf{S} = \{s\}\right] - \mu_C| \leq \varepsilon_C$, we have $\max_{s \in R} \mathsf{Pr}\left[\mathsf{S} = \{s\}\right] \leq \mu_C + \varepsilon_C$. Therefore

$$\begin{aligned}
&2^{-\alpha(\Pi,S)} = 2^{-H_\infty(S|\Pi(S))} \\
&= \sum_R \max_s \mathsf{Pr}\left[\mathsf{R} = R | \mathsf{S} = \{s\}\right] \mathsf{Pr}\left[\mathsf{S} = \{s\}\right] \\
&= \sum_{C \in P} \sum_{R \subseteq_n C} \binom{|C|-1}{n-1}^{-1} \max_{s \in R} \mathsf{Pr}\left[\mathsf{S} = \{s\}\right] \\
&\leq \sum_{C \in P} \sum_{R \subseteq_n C} \binom{|C|-1}{n-1}^{-1} (\mu_C + \varepsilon_C) \\
&= \sum_{C \in P} \binom{|C|}{n} \binom{|C|-1}{n-1}^{-1} (\mu_C + \varepsilon_C) \\
&= \sum_{C \in P} \frac{|C|}{n} (\mu_C + \varepsilon_C) = \frac{\varepsilon_P + 1}{n}.
\end{aligned}$$

In a similar fashion, we analyze the anonymity of Π'. For any $C \in P$ we have

$$\mathsf{Pr}\left[\mathsf{R}' = B | \mathsf{S} = \{s\}\right] = \begin{cases} 1 & s \in C \\ 0 & \text{otherwise.} \end{cases}$$

Recall that for all $C \in P$, all $s \in C$, it holds that $|\mathsf{Pr}\left[\mathsf{S} = \{s\}\right] - \mu_C| \leq \varepsilon_C$. Therefore

$$\begin{aligned}
&2^{-\alpha(\Pi',S)} = 2^{-H_\infty(S|\Pi'(S))} \\
&= \sum_{C \in P} \max_s \mathsf{Pr}\left[\mathsf{R}' = B | \mathsf{S} = \{s\}\right] \mathsf{Pr}\left[\mathsf{S} = \{s\}\right] \\
&= \sum_{C \in P} \max_{s \in C} \mathsf{Pr}\left[\mathsf{S} = \{s\}\right] \\
&\leq \sum_{C \in P} \mu_C + \varepsilon_C = \sum_{C \in P} \frac{|C|}{|C|} (\mu_C + \varepsilon_C) \\
&\leq \sum_{C \in P} \frac{|C|}{n} (\mu_C + \varepsilon_C) = \frac{\varepsilon_P + 1}{n}.
\end{aligned}$$

□

Corollary 6.1. *Let $\mathcal{S}$ be a 1-signer distribution and $n \in [N]$. Suppose for each partition P in the support of P, for all $C \in P$, all $s \in C$, it holds that $|\mathsf{Pr}\left[\mathsf{S} = \{s\}\right] - \mu_C| \leq \varepsilon_C$ for some $\varepsilon_C \geq 0$, and $|C| \geq n$. Let $\varepsilon_P := \sum_{C \in P} |C| \varepsilon_C$ and let $\varepsilon_{\mathsf{P}} := \sum_P \mathsf{Pr}\left[\mathsf{P} = P\right] \varepsilon_P$. Then*

$$\alpha(\mathsf{S}, \Pi_{\mathrm{Part},\mathcal{P},1,n}) \geq \lg n - \lg(\varepsilon_{\mathsf{P}} + 1)$$

and

$$\alpha(\mathsf{S}, \Pi_{\mathrm{Part},\mathcal{P},1}) \geq \lg n - \lg(\varepsilon_{\mathsf{P}} + 1).$$

Proof. We prove the result for $\Pi_{\mathrm{Part},\mathsf{P},1,n}$ by direct calculation. As before, we write $\Pi = \Pi_{\mathrm{Part},\mathsf{P},1,n}$ and $\mathsf{R} = \Pi(\mathsf{S})$.

$$\begin{aligned}
&2^{-\alpha(\Pi,S)} = 2^{-H_\infty(S|\Pi(S))} \\
&= \sum_R \max_s \mathsf{Pr}\left[\mathsf{R} = R | \mathsf{S} = \{s\}\right] \mathsf{Pr}\left[\mathsf{S} = \{s\}\right] \\
&= \sum_R \max_s \sum_P \begin{pmatrix} \mathsf{Pr}\left[\mathsf{R} = R | \mathsf{S} = \{s\} \wedge \mathsf{P} = P\right] \\ \cdot \mathsf{Pr}\left[\mathsf{P} = P\right] \cdot \mathsf{Pr}\left[\mathsf{S} = \{s\}\right] \end{pmatrix} \\
&= \sum_P \mathsf{Pr}\left[\mathsf{P} = P\right] \sum_R \max_s \begin{pmatrix} \mathsf{Pr}\left[\mathsf{R} = R | \mathsf{S} = \{s\} \wedge \mathsf{P} = P\right] \\ \cdot \mathsf{Pr}\left[\mathsf{S} = \{s\}\right] \end{pmatrix} \\
&= \sum_P \mathsf{Pr}\left[\mathsf{P} = P\right] \sum_{C \in P} \sum_{R \subseteq_n C} \binom{|C|-1}{n-1}^{-1} \max_{s \in R} \mathsf{Pr}\left[\mathsf{S} = \{s\}\right] \\
&\leq \sum_P \mathsf{Pr}\left[\mathsf{P} = P\right] \sum_{C \in P} \sum_{R \subseteq_n C} \binom{|C|-1}{n-1}^{-1} (\mu_C + \varepsilon_C) \\
&= \sum_P \mathsf{Pr}\left[\mathsf{P} = P\right] \sum_{C \in P} \frac{|C|}{n} (\mu_C + \varepsilon_C) \\
&= \sum_P \mathsf{Pr}\left[\mathsf{P} = P\right] \frac{\varepsilon_P + 1}{n} = \frac{\varepsilon_{\mathsf{P}} + 1}{n}.
\end{aligned}$$

The analogy for $\Pi_{\mathrm{Part},\mathsf{P},1}$ is similar and is omitted. □