## A Variance of LogEstimator

We now bound the variance of our estimator by $O\left(\log ^{2} k\right)$. Recall that the output of LogEstimator is given by $\log (\boldsymbol{X} / t)-g\left(\boldsymbol{B}_{1}, \ldots, \boldsymbol{B}_{r}\right)$, where the function $g$ is bounded. Since the variance we seek is $O\left(\log ^{2} k\right)$, it suffices to show that the variance of $\log (\boldsymbol{X} / t)$ is $O\left(\log ^{2} k\right)$ with $i \sim \mathcal{D}$, since subtracting $g$ changes the estimate by at most a constant (see Lemma 2.3).
Lemma A.1. Let $\boldsymbol{i} \sim \mathcal{D}$ and $\boldsymbol{X}$ denote the number of independent trials from $\operatorname{Ber}\left(p_{\boldsymbol{i}}\right)$ before we see $t$ successes. Then, $\operatorname{Var}[\log (\boldsymbol{X} / t)]=O\left(\log ^{2} k\right)$.

Proof. Let $X_{\max }=2 k t$, and consider the random variable $\boldsymbol{X}^{\prime}=\min \left\{\boldsymbol{X}, X_{\max }\right\}$. Then

$$
\begin{aligned}
\operatorname{Var}[\log (\boldsymbol{X} / t)] & \leq \mathbf{E}\left[\left(\log (\boldsymbol{X} / t)-\log \left(\boldsymbol{X}^{\prime} / t\right)+\log \left(\boldsymbol{X}^{\prime} / t\right)\right)^{2}\right] \\
& \leq 2 \cdot \mathbf{E}\left[\left(\log (\boldsymbol{X} / t)-\log \left(\boldsymbol{X}^{\prime} / t\right)\right)^{2}\right]+2 \cdot \mathbf{E}\left[\log ^{2}\left(\boldsymbol{X}^{\prime} / t\right)\right] \\
& \leq 2 \cdot \mathbf{E}\left[\log ^{2}\left(\boldsymbol{X} / \boldsymbol{X}^{\prime}\right)\right]+2 \log ^{2}(2 k) \\
& \leq \frac{4}{\ln ^{2}(2)} \cdot \mathbf{E}\left[\left(\sqrt{\frac{\boldsymbol{X}}{\boldsymbol{X}^{\prime}}-1}\right)^{2}\right]+2 \log ^{2}(2 k)
\end{aligned}
$$

where we used that $\log \left(\boldsymbol{X}^{\prime} / t\right) \leq \log (2 k)$ always, and that $\log (z) \leq \sqrt{z-1} / \ln (2)$ for all $z \geq 1$. Then,

$$
\mathbf{E}\left[\frac{\boldsymbol{X}}{\boldsymbol{X}^{\prime}}-1\right] \leq \mathbf{E}\left[\frac{\boldsymbol{X}}{X_{\max }}\right]=\frac{1}{X_{\max }} \sum_{i=1}^{k} p_{i} \cdot \frac{t}{p_{i}}=\frac{t k}{X_{\max }}=2
$$

## B Omitted Details from Section 2

Proof of Claim 2.5. Notice that $\boldsymbol{X}$ is the number of trials from $\operatorname{Ber}\left(p_{i}\right)$ until we see $t$ successes. We now have the following string of equalities:

$$
\begin{aligned}
\mathbf{E}_{\boldsymbol{X}, \boldsymbol{B}_{1}, \ldots, \boldsymbol{B}_{r}}\left[\boldsymbol{\eta}-\log \left(\frac{1}{p_{i}}\right)\right] & =\mathbf{E}_{\boldsymbol{X}}[\log \boldsymbol{Y}]-\mathbf{E}_{\boldsymbol{B}_{1}, \ldots, \boldsymbol{B}_{r}}\left[g\left(\boldsymbol{B}_{1}, \boldsymbol{B}_{2}, \ldots, \boldsymbol{B}_{r}\right)\right] \\
& =\mathbf{E}_{\boldsymbol{X}}[f(\boldsymbol{Y})+h(\boldsymbol{Y})]-g\left(p_{i}, p_{i}^{2}, \ldots, p_{i}^{r}\right)=\mathbf{E}_{\boldsymbol{X}}[h(\boldsymbol{Y})]
\end{aligned}
$$

where we used the fact that $g$ is a linear function, and that $\mathbf{E}\left[\boldsymbol{B}_{\ell}\right]=p_{i}^{\ell}$ in order to substitute

$$
\mathbf{E}_{\boldsymbol{B}_{1}, \ldots, \boldsymbol{B}_{r}}\left[g\left(\boldsymbol{B}_{1}, \ldots, \boldsymbol{B}_{r}\right)\right]=g\left(p_{i}, p_{i}^{2}, \ldots, p_{i}^{r}\right)
$$

Furthermore, we divide $\log \boldsymbol{Y}=f(\boldsymbol{Y})+h(\boldsymbol{Y})$, where $f(z)$ is the degree- $r$ Taylor expansion of $\log z$ at 1 , and $h(z)=\log z-f(z)$ is the error in the degree- $r$ Taylor expansion of $\log (z)$, i.e.,

$$
h(z)=\log (z)-f(z)
$$

Finally, by construction of $g, \mathbf{E}[f(\boldsymbol{Y})]=g\left(p_{i}, p_{i}^{2}, \ldots, p_{i}^{r}\right)$, which gives the desired equality.
Verifying $\boldsymbol{Y}$ is subgamma. Recall that $\boldsymbol{X}$ is the number of independent draws from a $\operatorname{Ber}(p)$ distribution until we see $t$ successes. In other words, we may express $\boldsymbol{X}=\boldsymbol{X}_{1}+\cdots+\boldsymbol{X}_{t}$, where $\boldsymbol{X}_{i}$ is the number of draws of $\operatorname{Ber}(p)$ before we get a single success. Then, we always satisfy

$$
\mathbf{E}\left[\boldsymbol{X}_{i}\right]=\frac{1}{p} \quad \operatorname{Pr}\left[\boldsymbol{X}_{i}>\ell\right]=(1-p)^{\lceil\ell\rceil}<e^{-p \ell}
$$

This, in turn, implies that for any $r \geq 1$

$$
\left(\mathbf{E}\left[\left|\boldsymbol{X}_{i}-1 / p\right|^{r}\right]\right)^{1 / r} \leq\left(\mathbf{E}_{\boldsymbol{X}_{i}, \boldsymbol{X}_{i}^{\prime}}\left[\left|\boldsymbol{X}_{i}-\boldsymbol{X}_{i}^{\prime}\right|^{r}\right]\right)^{1 / r} \leq 2\left(\mathbf{E}\left[\left|\boldsymbol{X}_{i}\right|^{r}\right]\right)^{1 / r}=O(r / p)
$$

where the first line is by Jensen's inequality, and the second is by the triangle inequality and Hölder inequality. Finally, we use the tail bound on $\boldsymbol{X}_{i}$ to upper bound the expectation of $\left|\boldsymbol{X}_{i}\right|^{r}$. Then, we have

$$
\begin{aligned}
\mathbf{E}\left[e^{\lambda\left(\boldsymbol{X}_{i}-1 / p\right)}\right] & =1+\lambda \mathbf{E}\left[\boldsymbol{X}_{i}-1 / p\right]+\sum_{k=2}^{\infty} \frac{\lambda^{k}}{k!} \cdot \mathbf{E}\left[\left|\boldsymbol{X}_{i}-1 / p\right|\right] \\
& =1+\sum_{k=2}^{\infty} \frac{\lambda^{k}}{k!}(O(k / p))^{k} \leq 1+O\left(\lambda^{2} / p^{2}\right), \quad \text { when }|\lambda| \text { sufficiently smaller than } p \\
& \leq \exp \left(O\left(\lambda^{2} / p^{2}\right)\right)
\end{aligned}
$$

Then, since $\boldsymbol{X}_{1}, \ldots, \boldsymbol{X}_{t}$ are all independent, we have

$$
\mathbf{E}\left[e^{\lambda(\boldsymbol{X}-t / p)}\right] \leq \exp \left(O\left(\lambda^{2} t / p^{2}\right)\right) \Longrightarrow \mathbf{E}\left[e^{\lambda(\boldsymbol{Y}-1)}\right] \leq \exp \left(O\left(\lambda^{2} / t\right)\right)
$$

and this bound is valid whenever $|\lambda|$ is sufficiently smaller than $t$.

## C Omitted Proofs from Section 3

Proof of Lemma 3.1. The approach is to estimate

$$
\begin{equation*}
\mathbf{E}_{i \sim \mathcal{D}}\left[h_{t}\left(p_{\boldsymbol{i}}\right)\right]=\mathbf{E}_{\boldsymbol{i} \sim \mathcal{D}}\left[g\left(p_{\boldsymbol{i}}, p_{\boldsymbol{i}}^{2}, \ldots, p_{\boldsymbol{i}}^{r}\right)\right] \tag{10}
\end{equation*}
$$

There exists an algorithm using $O\left(\log (1 / \epsilon) / \epsilon^{2}\right)$ samples to estimate the above quantity: for $j \in$ $\left\{0, \ldots, O\left(1 / \epsilon^{2}\right)\right\}$, one takes a sample $\boldsymbol{i}_{j} \sim \mathcal{D}$ and uses $r=O(\log (1 / \epsilon))$ additional samples $s_{1}, \ldots, s_{r} \sim \mathcal{D}$ to define

$$
\boldsymbol{B}_{m}^{(j)} \stackrel{\text { def }}{=} \mathbb{1}\left\{\boldsymbol{s}_{1}=\cdots=\boldsymbol{s}_{m}=\boldsymbol{i}_{j}\right\} \quad \text { and } \quad \boldsymbol{Z}_{j}=g\left(\boldsymbol{B}_{1}^{(j)}, \ldots, \boldsymbol{B}_{r}^{(j)}\right)
$$

Then, let $\boldsymbol{Z}$ be the average of all $\boldsymbol{Z}_{j}$ 's, which is an unbiased estimate to $\mathbf{E}_{\boldsymbol{i} \sim \mathcal{D}}\left[g\left(p_{\boldsymbol{i}}, p_{\boldsymbol{i}}^{2}, \ldots, p_{\boldsymbol{i}}^{r}\right)\right]$. Since $g$ is bounded (from Lemma 2.3), the variance of $O\left(1 / \epsilon^{2}\right)$ such values is a large constant factor smaller than $\epsilon^{2}$. By Chebyshev's inequality, we estimate (10) to error $\pm \epsilon$ with probability at least 0.9. With that estimate, we will now use Lemma 2.4. Specifically, the entropy of $\mathcal{D}$ is exactly $\mathbf{E}_{\boldsymbol{i} \sim \mathcal{D}}\left[\log \left(1 / p_{\boldsymbol{i}}\right)\right]$, and we have

$$
\begin{aligned}
\left|\mathbf{E}_{\boldsymbol{i} \sim \mathcal{D}}\left[\log \left(1 / p_{i}\right)\right]-(\hat{H}-\boldsymbol{Z})\right| & \leq \epsilon+\left|\mathbf{E}_{\boldsymbol{i} \sim \mathcal{D}}\left[\log \left(1 / p_{\boldsymbol{i}}\right)\right]-(\hat{H}-\boldsymbol{Z})\right| \\
& \leq \epsilon+\mathbf{E}_{\boldsymbol{i} \sim \mathcal{D}}\left[\left|\log \left(\frac{1}{p_{\boldsymbol{i}}}\right)-\mathbf{E}\left[\boldsymbol{\eta}_{\boldsymbol{i}}\right]\right|\right] \leq 2 \epsilon
\end{aligned}
$$

where $\boldsymbol{\eta}_{\boldsymbol{i}}$ is the result of running LogEstimator $(\mathcal{D}, \boldsymbol{i})$.
Proof of Lemma 3.2. We note that since $\log (\cdot)$ is monotone increasing, we must have $H \geq \tilde{H}$. To see that it is not much larger, note that we always have $\log z=\ln (z) / \ln (2) \leq(z-1) / \ln (2)$, which means

$$
\begin{aligned}
H-\tilde{H} & =\mathbf{E}_{\boldsymbol{i}, \boldsymbol{X}}\left[\log \left(\boldsymbol{X} / \boldsymbol{X}^{\prime}\right)\right] \leq \frac{1}{\ln (2)} \mathbf{E}_{\boldsymbol{i}, \boldsymbol{X}}\left[\frac{\boldsymbol{X}}{\min \left\{\boldsymbol{X}, X_{\max }\right\}}-1\right] \leq \frac{1}{\ln (2)} \mathbf{E}_{\boldsymbol{i}, \boldsymbol{X}}\left[\frac{\boldsymbol{X}}{X_{\max }}\right] \\
& =\frac{1}{X_{\max } \cdot \ln (2)} \sum_{i=1}^{k} p_{i} \cdot \frac{t}{p_{i}}=\frac{t k}{X_{\max } \cdot \ln (2)}=\epsilon
\end{aligned}
$$

Proof of Lemma 3.4. Substituting the $r_{\ell}$ values into Lemma 3.3 ensures $\mathbf{E}\left[\right.$ Error $\left.^{2}\right] \leq \epsilon^{2} / 10$. Hence the estimator is within $\pm \epsilon$ of $\tilde{H}$ with probability 0.9 by Chebyshev's inequality.
For the intervals $\ell=\{1, \ldots, L-1\}$, we always spend $r_{\ell}$ tries to determine whether a sample falls within a particular interval. Note that we take one sample to determine $\boldsymbol{i} \sim \mathcal{D}$, and then we take at
most $b_{\ell}$ samples. Therefore, the sample complexity for these is

$$
\begin{aligned}
\sum_{\ell=1}^{L-1} r_{\ell} \cdot b_{\ell} & =\frac{80 t k}{\epsilon^{2}} \cdot \sum_{\ell=1}^{L-1} \frac{\log ^{2}\left(\log ^{(\ell-1)}(k) / \epsilon\right)}{\left(\log ^{(\ell)} k\right)^{3}}=\frac{80 t k}{\epsilon^{2}} \cdot \sum_{\ell=1}^{L-1} \frac{\left(3 \log ^{(\ell)}(k)+\log (1 / \epsilon)\right)^{2}}{\left(\log ^{(\ell)} k\right)^{3}} \\
& \leq k t \cdot O\left(\log ^{2}(1 / \epsilon) / \epsilon^{2}\right)
\end{aligned}
$$

where we used the fact that

$$
\sum_{\ell=1}^{L-1} \frac{1}{\left(\log ^{(\ell)} k\right)} \leq \frac{1}{1}+\frac{1}{\exp (1)}+\frac{1}{\exp (\exp (1))}+\frac{1}{\exp (\exp (\exp (1)))}+\ldots=O(1)
$$

Finally, it remains to bound the expected sample complexity of the bucket $L$. Here, we note

$$
r_{L}=\frac{O(1)}{\epsilon^{2}} \cdot \log ^{2}\left(\frac{\log ^{(L-1)} k}{\epsilon}\right) \leq O\left(\frac{\log ^{2}(1 / \epsilon)}{\epsilon^{2}}\right)
$$

Therefore, the expected sample complexity for interval $L$ is $r_{L} \cdot \sum_{i=1}^{k} p_{i} \cdot \frac{t}{p_{i}}=O\left(k \log ^{4}(1 / \epsilon) / \epsilon^{2}\right)$.

## D Conjectured Lower Bound

Recall that without a memory constraint the sample complexity is known to be $n=\Theta\left(\max \left\{\epsilon^{-1}\right.\right.$. $\left.k / \log (k / \epsilon), \epsilon^{-2} \log ^{2} k\right\}$ ) [VV17, VV11, JVHW15, WY16]. To prove a $\Omega\left(k / \epsilon^{2}\right)$ lower bound for the memory constrained version, we conjecture the following randomized process can be used to generate distributions over $[2 k]$ that look alike to any constant space algorithm that uses $o\left(k / \epsilon^{2}\right)$ samples but they have different entropies.
Suppose we have $k$ Bernoulli random variables with parameter $\alpha: Y_{1}, \ldots, Y_{k}$. And, we have $k$ Rademacher random variables $Z_{1}, \ldots, Z_{k}$ (that are +1 or -1 with probability $1 / 2$ ). We construct distribution $p$ in such a way that it is uniform over $k$ pairs of elements $(1,2),(3,4), \ldots,(2 k-1,2 k)$. However, conditioning on pair $(2 i-1,2 i)$, we may have a constant bias based on the random variable $Y_{i}$. And, we decide about the direction of the bias based on $Z_{i}$. More precisely, we set the probabilities in $p$ as follows:

$$
p_{2 i-1}=\frac{1+Y_{i} \cdot Z_{i} / 4}{2 k}, \quad p_{2 i}=\frac{1-Y_{i} \cdot Z_{i} / 4}{2 k} \quad \forall i \in[k]
$$

Now, it is not hard to show that if we generate two distributions as above with $\alpha=(1+\epsilon) / 2$ and $\alpha=(1-\epsilon) / 2$, then their entropies are $\Theta(\epsilon)$ separated with a constant probability. Thus, any algorithm that can estimate the entropy has to distinguish $\alpha=(1+\epsilon) / 2$ from $\alpha=(1-\epsilon) / 2$. Intuitively, to learn $\alpha$, we would require to determine $\Omega\left(1 / \epsilon^{2}\right)$ many of $Y_{i}$ 's. Since we have only a constant words of memory, we cannot perform the estimation of the $Y_{i}$ 's in parallels. Thus, any natural algorithm would require to draw $\Omega\left(k / \epsilon^{2}\right)$ samples.

