

Far in the Tail the Moment Generating Function Bounds a Slow Sort Tighter; Near the Mean Variance Alone Does

— building the Chernoff bound for random-pivot quicksort from a recursion for the moment generating function, the chance of exceeding twice the mean is bounded within 20.97568 times the exact value, while for departures near the mean the bound is looser than Cantelli's —the separator is whether a departure is measured in multiples of the mean or of the standard deviation — [Paper 987]

Yuuki YAMAGISHI (山岸 勇輝)

Yamagishi Complexity Torus Lab, Japan

September 2026

YOI: <https://www.yamagishiyuuki.com/yoi/987>

— *From thinking to syncing.* —

Abstract

Variance alone bounds the chance that random-pivot quicksort runs slow, but the bound is far looser than reality (Paper 952). Exponential tail bounds using the shape of the distribution were only named there. For the number of comparisons of the random-pivot quicksort of Paper 952, this paper builds the Chernoff bound from a recursion for the moment generating function and compares it with the Cantelli bound of Paper 952 and with the exact tail of the distribution at $n = 100$. No new theorem or law is claimed. First, far in the tail, the bound from the moment generating function is tighter by orders of magnitude: for $n = 100$ the chance of exceeding twice the mean is exactly 4.780561×10^{-9} ; the Chernoff bound was **20.97568** times that and the Cantelli bound **1.748705** $\times 10^6$ times (Section 3). Second, for departures near the mean, the variance-only bound is tighter: for exceeding 1.25 times the mean (exactly 1.525455×10^{-2}), Cantelli's 1.188529×10^{-1} was smaller than Chernoff's 1.265041×10^{-1} (Section 3). Third, where the two bounds cross moves with n : the two were equal for exceeding **1.2598** times the mean at $n = 100$ and **1.1599** times at $n = 1000$, with Chernoff smaller further out (Section 4). Fourth, at $n = 1000$ the gap widens further: for exceeding twice the mean the bounds were 3.378363×10^{-3} for Cantelli and 9.840440×10^{-14} for Chernoff (Section 4). Fifth, the separator is whether a departure is measured in multiples of the mean or of the standard deviation: dividing the crossing departure by the mean gave a ratio of **1.6251** between $n = 100$ and $n = 1000$, but dividing by the standard deviation gave **2.8299** and **2.7462**, a ratio of **1.0305** (Section 5). **Keywords:** quicksort, Chernoff bound, moment generating function, Cantelli's inequality, tail probability, standard deviation.

Scope of this paper (scope note): No new theorem or law is claimed: Chernoff's method of bounding tail probabilities from the moment generating function, that the size of the chance of random-pivot quicksort's comparisons departing far from the mean has been studied (McDiarmid and Hayward), and that variance-only bounds are loose far in the tail, are known in probability and the analysis of algorithms. Only comparisons are counted: memory reads and writes, constant factors and recursion depth are not treated (as in Paper 927 and Paper 952). Two bounds only: only the Chernoff and Cantelli bounds are compared; other inequalities and the best bound are not treated. n is only 100 and 1000, and no tail shape for n growing without limit is claimed. The crossing is not generalised: the number of standard deviations at which the two bounds cross is for these two values of n , and is not claimed to hold for other algorithms or distributions. Relation to earlier papers: Paper 952, in its scope note, did not derive strong bounds, only naming exponential tail bounds using the shape of the distribution and giving a variance-only bound; this paper sits in that seat (no earlier occupant). Paper 952 showed that variance alone bounds the chance that random-pivot quicksort runs slow, but that the bound is far looser than reality. Here how much tighter a bound using the shape of the distribution is, and where it loses to the variance-only bound, are counted. What is added: building the Chernoff bound from a recursion for the moment generating function, computing the exact distribution at $n = 100$ by convolution and comparing both bounds with it, finding where the two bounds cross, comparing that point in units of the mean and of the standard deviation, and placing the separator on "multiples of the mean, or of the standard deviation".

1. Notation and the Question

Setting: as in Paper 952, a list of length n is sorted by quicksort with a uniformly random pivot, and C_n is the number of comparisons. With the pivot at rank K , $C_n = (n - 1) + C_{K-1} + C'_{n-K}$ (independent), so the moment generating function $M_n(\lambda) = \mathbb{E}e^{\lambda C_n}$ satisfies

$$M_n(\lambda) = e^{\lambda(n-1)} \frac{1}{n} \sum_{K=1}^n M_{K-1}(\lambda) M_{n-K}(\lambda), \quad M_0 = M_1 = 1$$

This was run in logarithms up to n , and with mean μ the Chernoff bound and the Cantelli bound (Paper 952)

$$P(C_n \geq \mu + t) \leq \min_{\lambda > 0} M_n(\lambda) e^{-\lambda(\mu+t)}, \quad P(C_n - \mu \geq t) \leq \frac{\sigma^2}{\sigma^2 + t^2}$$

were compared. λ was chosen near the minimum among 80 log-spaced points from 10^{-5} to 2 and refined by ternary search. At $n = 100$ the distribution of C_n itself was computed by convolution with the same recursion (total probability 1.000000000000, 0.2 seconds).

Question: how much tighter is a bound using the shape of the distribution than the variance-only bound, and where does it lose?

2. Mean and Variance

Computation 2.1: mean 647.850259 and variance 3538.266678 at $n = 100$ (the same from the exact distribution), and mean 10985.912663 and variance 409117.768339 at $n = 1000$. The slope of the log moment generating function at $\lambda = 0$ was 647.8503 at $n = 100$ and 10985.9126 at $n = 1000$, agreeing with the mean within a relative 10^{-4} .

3. Three Values at $n = 100$

Computation 3.1: chance of exceeding a multiple of the mean.

exceeding	Cantelli	Chernoff	exact	Cantelli over exact	Chernoff over exact
1.25 times mean	1.188529×10^{-1}	1.265041×10^{-1}	1.525455×10^{-2}	7.791305	8.292877
1.5 times mean	3.262105×10^{-2}	2.540512×10^{-3}	1.885716×10^{-4}	172.9903	13.47240
twice the mean	8.359791×10^{-3}	1.002755×10^{-7}	4.780561×10^{-9}	1.748705×10^6	20.97568

In all three, both bounds were at least the exact value. The λ giving Chernoff's minimum was 2.032720×10^{-2} , 2.719945×10^{-2} and 3.474357×10^{-2} .

Observation 3.2

far in the tail Chernoff was tighter, near the mean Cantelli was. Cantelli over exact widened from 7.791305 to 1.748705×10^6 , while Chernoff over exact only widened from 8.292877 to 20.97568: at 1.25 times the mean, Cantelli was smaller than Chernoff.

4. The Crossing and $n = 1000$

Computation 4.1: the two bounds at $n = 1000$.

exceeding	Cantelli	Chernoff
1.25 times mean	5.144673×10^{-2}	1.374939×10^{-2}
1.5 times mean	1.337787×10^{-2}	8.324893×10^{-6}
twice the mean	3.378363×10^{-3}	9.840440×10^{-14}

Computation 4.2: finding by bisection where the two bounds are equal gave 1.2598 times the mean at $n = 100$ and 1.1599 times at $n = 1000$; nearer the mean Cantelli was smaller, further out Chernoff was.

5. The Separator

Observation 5.1

the separator is whether a departure is measured in multiples of the mean or of the standard deviation. Dividing the crossing departure by the standard deviation gave 2.8299 at $n = 100$ and 2.7462 at $n = 1000$, a ratio of 1.0305: dividing by the mean gave a ratio of 1.6251. The same “twice the mean” is a departure of 10.8913 standard deviations at $n = 100$ and 17.1756 at $n = 1000$, further out in the tail the larger n . At $n = 1000$, 1.25 times the mean was already 4.2939 standard deviations out, beyond the crossing at 2.7462, and Chernoff was smaller in all three.

Observation 5.2

what is not said. The tail shape as n grows without limit, other inequalities or the best bound, whether the crossing multiple holds for other algorithms, and actual running time are not treated. What is shown is where, for two values of n , the bound using the shape of the distribution and the variance-only bound cross, and how far apart they grow in the far tail.

6. Roll Call and Closing

1. recursion for the moment generating function (Section 1): slope matches the mean.
2. at twice the mean Chernoff is **20.97568** times exact (Computation 3.1): Cantelli **1.748705**×10⁶ times.
3. at 1.25 times the mean Cantelli is smaller (Computation 3.1): variance alone is tighter near the mean.
4. crossing at **1.2598** and 1.1599 times the mean (Computation 4.2): moves with n .
5. **2.8299** and 2.7462 standard deviations, ratio 1.0305 (Observation 5.1): nearly the same.
6. multiples of the mean, or of the standard deviation (Observation 5.1): the separator; the scope note of Paper 952.

Closing: Far in the tail the moment generating function bounds a slow sort tighter; near the mean variance alone does. Building the moment generating function for the quicksort of Paper 952 by recursion, the Chernoff bound on exceeding twice the mean at $n = 100$ was 20.97568 times the exact value, far tighter than Cantelli’s 1.748705×10⁶ times, while at 1.25 times the mean Cantelli was smaller: the two crossed at 1.2598 times the mean (1.1599 times at

$n = 1000$), a departure of 2.8299 and 2.7462 standard deviations. The separator is whether a departure is measured in multiples of the mean or of the standard deviation. Placed among the earlier papers: to the tail Paper 952 bounded by variance alone, the bound using the shape of the distribution that Paper 952 only named was added. To be honest: two bounds and exact values were set side by side for two values of n . From where using the shape of the distribution pays, and no further.

References

1. Chernoff, H. (1952). A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Annals of Mathematical Statistics*.
 2. McDiarmid, C., & Hayward, R. (1996). Large deviations for Quicksort. *Journal of Algorithms*.
 3. Yamagishi, Y. (2026). *Quicksort with the First Element as Pivot Is Slowest on Sorted Input, and Choosing the Pivot at Random Makes Its Average Cost Independent of the Input* [Paper 927]. YOI: <https://www.yamagishiyuuki.com/yoi/927>
 4. Yamagishi, Y. (2026). *Variance Alone Bounds the Chance That Random-Pivot Quicksort Runs Slow, but the Bound Is Far Looser Than Reality* [Paper 952]. YOI: <https://www.yamagishiyuuki.com/yoi/952>
-

— *From thinking to syncing.* —

Yuuki YAMAGISHI (山岸 勇輝)

© 2026 Yamagishi Complexity Torus Lab, Japan · Licensed under Creative Commons Attribution 4.0 International (CC BY 4.0)