A LESSON ABOUT PRIME NUMBERS

What would solving the Riemann Hypothesis mean?

This week a famous puzzle about prime numbers made the news. Before anyone tells you what it says, find the primes yourself — and notice a pattern that has kept mathematicians busy for more than 160 years.

Name _____ Class / year _____ Date _____

1. The sieve of Eratosthenes, up to 100

A **prime** is a whole number with exactly two factors: 1 and itself. So 7 is prime, but 9 is not (3 × 3 = 9). The number 1 is **not** a prime — it has only one factor. "Sieve" means a strainer. Around 240 BCE, Eratosthenes found a simple way to pour the numbers through and leave only the primes behind.

1. Cross out **1**. It is not a prime.
2. Circle **2**, then cross out every other multiple of 2 (4, 6, 8, …).
3. The next number still clear is **3**. Circle it, then cross out the other multiples of 3.
4. Do the same for **5**, then for **7**. You can stop after 7: 11 × 11 = 121, already past 100, so every composite number up to 100 has a prime factor of 7 or less.

Mistakes are useful. If you cross out a prime by accident, check its factors and put it back. Lots of people do this the first time they try a sieve.

1	2	3	4	5	6	7	8	9	10
11	12	13	14	15	16	17	18	19	20
21	22	23	24	25	26	27	28	29	30
31	32	33	34	35	36	37	38	39	40
41	42	43	44	45	46	47	48	49	50
51	52	53	54	55	56	57	58	59	60
61	62	63	64	65	66	67	68	69	70
71	72	73	74	75	76	77	78	79	80
81	82	83	84	85	86	87	88	89	90
91	92	93	94	95	96	97	98	99	100

When you have finished, the numbers still clear (and circled) should be the primes up to 100. How many did you find? _____

2. Count the primes in blocks of ten

The primes can look a bit scattered. Counting them in each block of ten makes the pattern easier to see. Use your sieve. List the primes in each row, then write how many there are.

Block	How many primes?	List them
1–10		
11–20		
21–30		
31–40		
41–50		
51–60		
61–70		
71–80		
81–90		
91–100		

Check: the ten counts should add up to the number of primes you found on page 1. Total: _____

3. What do you notice?

Q1. Which block has the most primes? Which has the fewest?

Q2. Do the primes stay evenly spread as the numbers get larger, or do they seem to thin out? Point to the evidence in your table.

Q3. Why might there be fewer primes later on? (Hint: as numbers get larger, they have more possible factors.)

Q4. A friend says “the first block has five primes: 1, 2, 3, 5 and 7.” What would you say to them, and why?

> The counts you should get, if the sieve is right, are **4, 4, 2, 2, 3, 2, 2, 3, 2, 1**. If your table does not match, go back to the grid and look together — that check is part of the work, not a sign you “can’t do primes”.

4. Predict, then check: a rule of thumb

The primes in your table look almost random. Yet there is a surprisingly simple rule underneath. The number of primes up to a number *x* is about

$$x \,/\, \ln(x)$$

ln is the natural logarithm (logarithm with base *e*). On a scientific calculator, look for the button marked ***ln***.

First, guess — before you calculate

You have just sieved the primes up to 100, so you already know the real answer. Use that knowledge later. For now, imagine you *didn't* have the list: what would you guess for "how many primes up to 100"? There is no wrong first try — a guess that is off is useful, because you have started thinking.

My guess for the number of primes up to 100 is _____ .

Now calculate the estimate, then compare

Q5. Work out 100 / ***ln***(100). Write down the value of ***ln***(100) that your calculator shows, then divide.

Estimate, to 1 decimal place: _____

Q6. From your sieve, the actual number of primes up to 100 is _____ .

Q7. How close was the estimate? How close was your guess? (Close is not the same as exact.)

Q8. A bigger check. Up to 1,000, the same rule gives 1000 / ***ln***(1000) ≈ **145**. There are actually **168** primes up to 1,000. Is the estimate still "in the right neighbourhood"? What has happened to the size of the gap?

5. What the Riemann Hypothesis says — in plain words

The rule above is called the **prime number theorem**. It tells us the primes get thinner, on average, in a predictable way. It does not say they are equally many in every block of ten — your table already shows they are not.

The **Riemann Hypothesis** (posed by Bernhard Riemann in 1859) is a sharper claim about that rule. In plain words:

> The primes are as evenly spread as possible. The error in the estimate stays small.

You do not need the formal version today — that uses later maths. The idea is enough: the primes are not chaotic, and the gap between the estimate and the real count stays under tighter control than the basic rule alone would guarantee.

Q9. In one or two sentences, explain what the Riemann Hypothesis is claiming, using your own words. (Try it without looking back, then check.)

6. Why it matters

A lot of other maths already assumes the Riemann Hypothesis is true. That is one reason a proof would be a big deal: many results would rest on firmer ground.

> **Would a proof break the internet?** No. Online security often uses very large primes, but solving the Riemann Hypothesis would **not** break encryption. The usual systems (such as RSA) depend on it being hard to factorise a huge number, not on this puzzle remaining unsolved.

In 2000 the Clay Mathematics Institute named the Riemann Hypothesis one of seven **Millennium Prize Problems**, with a prize of **US$1 million** for an accepted solution.

In November 2026, Levent Alpöge, a mathematician at Anthropic, released a computer-checked proof he worked on with the AI Claude. It is a Lean library of about 17 million lines, and it also covers the Generalized Riemann Hypothesis for Dirichlet L-functions. Independent machine checks by Andrew Sutherland (MIT) and Michael Rubinstein (University of Waterloo) report that the proof compiles using only Lean's standard axioms, so it **now looks very likely correct**. A human-readable write-up is still being drafted. The Clay Mathematics Institute still lists the problem as Active, and the US$1 million prize has **not** been awarded. A headline is not the same as a prize being paid.

Q10. A classmate says "they've solved it, so internet shopping is no longer safe." How would you answer them?

Q11. Why might a prize committee wait, even after a formal proof is in the news?

7. Extension — for older students, or anyone still curious

You may use a scientific calculator. The symbol $\pi(x)$ means "the number of primes less than or equal to x" (it is not the circle-number π).

E1. Complete this table. Estimates to the nearest whole number are fine. The actual counts are given, so you can compare.

x	Estimate $x / \ln(x)$	Actual $\pi(x)$	Gap (actual − estimate)
10		4	
100		25	
1,000		168	
1,000,000		78,498	

E2. The gap in the last column gets larger as x grows. Now look at the gap as a fraction of the actual count (gap divided by the actual count). What do you notice? (Check 100 and 1,000 carefully — the change is not perfectly smooth.)

E3. A student reads "the primes are as evenly spread as possible" and concludes that every block of ten should contain the same number of primes. Using page 2, explain why that is a misunderstanding.

E4. The prime number theorem (the "about $x / \ln(x)$" rule) was proved in 1896. The Riemann Hypothesis is a sharper claim about how small the error stays. In your own words, what is the difference between "we already know the average thinning" and "we want tighter control of the error"?

The formal statement uses the Riemann zeta function. We have stayed with the prime-counting picture on purpose — a door you can knock on later, not one you need today.

Answers — for teachers, or for checking afterwards

The 25 primes up to 100

2, 3, 5, 7, 11, 13, 17, 19, 23, 29, 31, 37, 41, 43, 47, 53, 59, 61, 67, 71, 73, 79, 83, 89, 97.

Primes in each block of ten

Block	How many?	List
1–10	4	2, 3, 5, 7 (1 is not a prime)
11–20	4	11, 13, 17, 19
21–30	2	23, 29
31–40	2	31, 37
41–50	3	41, 43, 47
51–60	2	53, 59
61–70	2	61, 67
71–80	3	71, 73, 79
81–90	2	83, 89
91–100	1	97

The ten counts are 4, 4, 2, 2, 3, 2, 2, 3, 2, 1 and add up to 25.

Q1. Most: 1–10 and 11–20 (four each). Fewest: 91–100 (one: 97).

Q2. They thin out overall. Early blocks have four primes; later blocks have two, three or one. It is not a perfectly smooth drop (41–50 and 71–80 rise slightly to three), which is worth noticing — thinning is a trend, not a straight line.

Q3. Larger numbers have more possible factors, so they are more likely to be composite. Accept any clear version of this idea.

Q4. 1 is not a prime, because a prime must have exactly two factors. The first block is 2, 3, 5 and 7 — four primes, not five. A friendly correction is enough; the friend has made a very common slip.

Q5. $\ln(100) = 2\ln(10) \approx 4.605170$, so $100 / \ln(100) \approx$ **21.7** (more precisely 21.71). The film rounds this to “about 22”. Accept 21.7, 21.71 or 22 if the working is shown.

Q6. 25.

Q7. The estimate is close but not exact — it is about 3 too low (or about 4 if they used 22). Guesses will vary; what matters is that they compare honestly.

Q8. Yes, 145 is still in the right neighbourhood of 168. The gap has grown (from about 3 to about 23). The estimate is a rule of thumb, not a perfect count.

Q9. Look for the two plain ideas: the primes are as evenly spread as possible, and the error in the $x / \ln(x)$ estimate stays small. Students do not need formal wording.

Q10. A proof would not break encryption. Security systems such as RSA rely on the difficulty of factorising large numbers, not on the Riemann Hypothesis remaining unsolved.

Q11. A computer-checked proof now looks very likely correct, but a human-readable write-up is still being drafted. The Clay Mathematics Institute still lists the problem as Active, and the prize has not been awarded. Checking a proof takes time. A news report is not the same as awarding the prize.

Extension

E1. Estimates: $10 / \ln(10) \approx$ **4.3** (nearest whole number 4); $100 / \ln(100) \approx$ **21.7** (22); $1{,}000 / \ln(1{,}000) \approx$ **144.8** (145); $1{,}000{,}000 / \ln(1{,}000{,}000) \approx$ **72,382**. Gaps (actual – estimate, using the more precise estimates): about **–0.3**, **3.3**, **23.2** and **6,116**. Accept sensible rounding.

E2. The absolute gap grows, but as a fraction of the actual count it tends to get smaller for large x (about 13% at 100, about 8% at 1,000,000). It is *not* perfectly smooth: at 1,000 the relative gap is about 14%, a little larger than at 100. That honesty is worth praising if a student spots it.

E3. “As evenly spread as possible” is about the error in the overall estimate staying small, not about every decade containing equally many primes. The table already shows the counts change (4, 4, 2, 2, 3, …, 1).

E4. We already know the average thinning (the prime number theorem). The Riemann Hypothesis is a claim that the wobble around that average stays unusually small. Accept any clear contrast between “the trend” and “tighter control of the error”.