What Numbers Can't Be Bought

On the Frobenius Problem, and the strange comfort of knowing exactly where impossibility ends

Tom Rocks Maths Essay Competition 2026 • Student Category (Age 16)

I want to start with something that bothered me for longer than I'd like to admit.

If you only have coins worth 3 pence and 5 pence, you can make 6 (two threes), you can make 8 (a three and a five), you can make 10, 11, 13. But you cannot make 7. You cannot make 4, or 1, or 2. Fine — those feel small, obviously unreachable. But 7 felt different to me. 7 is not that small. 7 is a prime number. 7 is the number of days in a week. And yet 7 pence is permanently, provably, forever out of reach if all you have are 3s and 5s.

What bothered me was not that 7 was impossible. It was that 8 was possible. And 9. And everything after that, without exception. Something *switches* at 8, and once it switches, it never switches back. I wanted to understand why.

That question — what is the largest number that cannot be formed from a fixed set of coins? — is the Frobenius Problem. It has been studied for well over a century. It is named after the German mathematician Ferdinand Georg Frobenius, though Frobenius himself never fully solved it, and neither has anyone since, at least not in the way you might hope.

The Two-Coin Case, Which Is Almost Too Clean

For two coin values — say a and b, with no common factor between them — there is a formula. It was known to Sylvester in the 1880s, and it is one of those results that feels slightly too good to be true:

$$g(a,\, b) \;=\; ab - a - b$$

For $a = 3$ and $b = 5$: $g = 15 - 3 - 5 = 7$. There it is. The formula just hands you the answer, and the answer is exactly what troubled me.

The condition that a and b share no common factor matters enormously. If both coins were even, every odd number would be unreachable — not just finitely many, but infinitely many. Coprimality is what guarantees that eventually, the integers close up. The formula is not just a trick; it encodes a real structural fact about how two coprime numbers mesh together as you add multiples of them.

There is also a result about *how many* numbers are unreachable, not just what the largest one is. For coprime a and b, the count is:

$$\frac{(a-1)(b-1)}{2}$$

For 3 and 5, that's $\frac{2\times 4}{2} = 4$. The four unreachable numbers are 1, 2, 4, and 7 — and if you check, yes, exactly four. What strikes me about this formula is its tidiness. Number theory often rewards you with expressions like this that look like they belong to geometry or combinatorics, not arithmetic. The boundary between the possible and the impossible turns out to have a precise, countable shape.

And there is a symmetry hiding inside this. For coprime a and b, and any integer n between 0 and $g(a, b)$: exactly one of n and $g(a, b) - n$ is representable. The unreachable numbers and the "gaps below g" mirror each other. I find this genuinely unsettling in the best way — there is no obvious reason from the problem's statement that obstruction should have a reflection, and yet it does.

Three Coins, and Why Everything Falls Apart

At this point a reasonable person might think: fine, for three coins it'll be messier, maybe two products instead of one, maybe an extra term. The mathematics does not care about being reasonable.

For three or more coin denominations, no closed-form formula is known. Not unknown-yet, like Fermat's Last Theorem once was. Provably not expressible as a simple formula — any general expression for $g(a, b, c)$ must be, in a precise sense, more complex than any polynomial in the inputs. The jump from two coins to three is not an incremental step. It is a wall.

I think this is worth pausing on. The problem with two variables yields to a six-symbol formula. The problem with three variables has resisted every attempt at a similar answer for over a hundred and forty years. Something fundamentally changes when you add a single extra coin. The mathematics does not scale.

> **Just to fix ideas.** Take coins of value 6, 9, and 20 pence. Can you make 43? No — you can verify by exhaustion that no combination works. Can you make 44? Yes: $6 \times 4 + 20 = 44$. Can you make every number from 44 onward? Yes. So $g(6, 9, 20) = 43$. This particular case became mildly famous because 6, 9, and 20 were (approximately) the original Chicken McNugget box sizes, making 43 the largest number of nuggets you could not order. I am aware this sounds like a joke. It is not a joke.

The McNugget case can be computed by hand without too much misery. But computing the Frobenius number in full generality, for arbitrary coin values and an arbitrary number of them, is an NP-hard problem. Meaning: we do not know of any efficient algorithm, and most mathematicians believe — though cannot prove — that none exists. The difficulty is intrinsic, not just a gap in our current methods.

What the Problem Is Really Asking

I have been describing this in terms of coins, but that framing undersells it. The Frobenius Problem is a question about the structure of the integers under addition. Given a set of generators, what does the additive semigroup they produce look like? Where does it become complete? What is the shape of its gaps?

This connects to algebraic geometry through the theory of numerical semigroups — the set of all numbers representable by your coins forms a semigroup under addition, and its properties control things like the singularities of certain algebraic curves. It connects to combinatorics, to scheduling theory, to coding theory. The problem appears across mathematics not because people are forcing it in, but because the underlying question — how does additive structure accumulate? — turns out to be load-bearing in many places.

There is a conjecture, still open, about the number of distinct numerical semigroups of a given genus (the number of gaps). The count seems to grow in a Fibonacci-like pattern. Nobody knows why. The integers, even in their most elementary behaviour, still contain things we do not understand.

The Thing That Actually Interests Me

Most mathematical problems have a clear hierarchy of difficulty: easy cases, hard cases, impossibly hard cases. The Frobenius Problem has a strange version of this. Two variables: completely solved, elegant formula, deep symmetry. Three variables: no formula, NP-hard, open conjectures. The gap between two and three is disproportionate in a way that still seems philosophically puzzling to me.

My instinct — and this is speculation, not a theorem — is that the two-variable case is solvable in closed form precisely because two coprime numbers tile the integers in an essentially one-dimensional way. Adding a third generator introduces a combinatorial explosion that one-dimensional intuition cannot contain. The geometry of the problem jumps dimension, and our algebraic tools have not kept up.

What I find most interesting is not the formula for two variables, beautiful as it is. It is the *contrast*: that a problem this naturally stated, this concrete, this old, can collapse into genuine intractability the moment you add one more coin. Mathematics is full of this phenomenon — the sense that just past the solvable case, something enormous is hiding. The Frobenius Problem makes you feel it with unusual clarity.

Back to 7

So: why does the switch happen at 8 for coins of 3 and 5?

Because $8 = 3 + 5$. And $9 = 3 + 3 + 3$. And $10 = 5 + 5$. Once you have three consecutive reachable numbers — 8, 9, 10 — you can get every subsequent number by adding 3 to a number you could already reach. The switch is not mysterious; it is a consequence of the smallest denomination being 3, meaning that any run of three consecutive reachable integers propagates forward indefinitely.

The formula $g(a, b) = ab - a - b$ tells you when that run begins. It tells you precisely how long the integers resist being covered, and exactly where resistance ends.

I started this essay bothered by 7. I am still bothered by it, in the sense that I find it genuinely strange that a number as ordinary as 7 can be structurally unreachable from two numbers as ordinary as 3 and 5. But I understand the structure of the strangeness now, and that is different from the strangeness just sitting there unexplained.

That, I think, is what mathematics does at its best. It does not make the strange less strange. It gives you a precise account of exactly why the strange thing is strange, and where the strangeness stops.

In this case, it stops at 7. Everything after is fine.

Further Reading

- Ramírez Alfonsín, J.L. (2005). *The Diophantine Frobenius Problem.* Oxford University Press. — The definitive reference; thorough and beautifully organised.
- Sylvester, J.J. (1884). Mathematical Questions. *Educational Times*, 41, 21.
- Bras-Amorós, M. (2008). Fibonacci-like behavior of the number of numerical semigroups of a given genus. *Semigroup Forum*, 76(2), 379–384.
- Rödseth, Ø.J. (1978). On a linear Diophantine problem of Frobenius. *Journal für die reine und angewandte Mathematik*, 301, 171–178.