Hercules Always Wins: The Mathematics of the Hydra Game

Danial Bulatbek Tom Rocks Math Essay Competition

1 The Monster in the Myth

In Greek mythology, the Hydra of Lerna was a serpentine monster with a fearsome property: cut off one of its heads and two would grow back in its place. Hercules, charged with slaying it as one of his Twelve Labours, quickly discovered that brute force alone would not win the day. The more heads he severed, the worse his situation became. It seemed, on the face of it, utterly hopeless.

Mathematicians Laurence Kirby and Jeff Paris had the delightful idea of turning this myth into a rigorous game. In their 1982 paper *Accessible Independence Results for Peano Arithmetic*, they introduced a precise mathematical version of the Hydra — and asked: can Hercules always win? The answer, as we will see, is a resounding yes. But the proof is stranger and more beautiful than you might expect, requiring us to venture into the realm of the infinite to prove something about a completely finite game.

2 The Rules of the Game

First, let us describe the modern mathematical Hydra. Rather than a simple many-headed beast, our Hydra has a *tree* structure — like a branching family tree, but drawn upside down. At the very bottom sits a single node called the **root** (the Hydra's body). Growing upward from the root are branches, and at the end of each branch is a **head**. Crucially, some heads are attached directly to the root, while others sit atop more complex branching structures called **necks**.

Hercules plays the game by cutting off heads one at a time. The Hydra responds according to two rules:

1. **Simple case:** If the head Hercules cuts is attached directly to the root, nothing grows back. The Hydra simply loses that head.
2. **General case:** If the head is further from the root, the Hydra responds. Let x be the node where the severed neck was attached, and let y be the node one step closer to the root from x. The Hydra then grows N new copies of the entire subtree rooted at x — and Hercules gets to choose how large N is! (In the standard version, N equals the current turn number.)

This is the twist that makes the game so counterintuitive. Every time Hercules cuts a head with a grandparent, the Hydra doesn't just replace the head — it *multiplies* entire branches of itself. The situation appears to get catastrophically worse with each move. How could Hercules possibly win?

3 A Small Example

Let us trace through a small example to get a feel for the game. Imagine a Hydra that looks like this: a root, with one node above it, and two heads on top of that node. We might draw it as:

$$Head_1 \ Head_2 \nwarrow \nearrow \text{Node} \ \uparrow Root$$

When Hercules cuts $Head_1$, it has a grandparent (the Root). So we descend one step from Node to Root. The Hydra then grows N copies of the subtree rooted at Node (which currently has only $Head_2$ on top of it). If this is turn 1 and $N = 1$, the Hydra after the first cut still has one node with one head, plus the original $Head_2$ remaining — that is, a node with two heads. Now Hercules cuts one of those, and so on. Despite looking dire initially, the game does eventually terminate.

What if Hercules is allowed to choose $N = 1{,}000{,}000$ on every turn? The Hydra would grow to an almost incomprehensible size. Yet — as we are about to prove — Hercules will still win in a finite number of steps. The key insight lies not in counting heads, but in something far more abstract.

4 Counting Beyond Infinity: Ordinal Numbers

To understand why Hercules must always win, we need a tool from the mathematical theory of infinity: **ordinal numbers**. While the natural numbers 0, 1, 2, 3, … describe the sizes of finite sets, ordinals extend this idea into the infinite.

After all the natural numbers, the first infinite ordinal is written ω (omega). It represents the order type of the sequence 0, 1, 2, 3, … Beyond that, we have $\omega + 1$, $\omega + 2$, …, then 2ω, 3ω, …, then ω^2, ω^3, all the way up to ω^ω, and beyond. The crucial property we will use is this:

There is no infinite strictly decreasing sequence of ordinals.

In other words, if you start at any ordinal and keep going strictly downwards, you must eventually reach zero. This is called being *well-ordered*. The ordinary natural numbers have this property too — any descending sequence of natural numbers must terminate. Ordinals extend this idea into the infinite, and this is the property that will guarantee Hercules his victory.

5 Assigning an Ordinal to a Hydra

Here is the key idea. We will assign an ordinal number to every possible Hydra, in such a way that each move *strictly decreases* that ordinal. Since there is no infinite descending sequence of ordinals, the game must eventually terminate — and the only way it can terminate is when every head has been removed, meaning Hercules has won.

The assignment is defined recursively, starting from the leaves and working down to the root:

3. Every **head** is assigned the ordinal *0*.
4. If a node has children with assigned ordinals $\alpha_1 \geq \alpha_2 \geq ... \geq \alpha_n$, then that node is assigned the ordinal $\omega^{\alpha_1} + \omega^{\alpha_2} + ... + \omega^{\alpha_n}$.
5. The ordinal of the **entire Hydra** is the ordinal assigned to the root.

Let us see this in action. A single head has ordinal 0. A node directly above the root, with one head on top of it, has ordinal $\omega^0 = 1$. The root, looking up at this node, therefore has ordinal $\omega^1 = \omega$. A node with two heads has ordinal $\omega^0 + \omega^0 = 2$, and a root looking up at it has ordinal $\omega^2 = \omega^2$. Already we can see how the ordinals grow rapidly with the depth of the Hydra.

6 Why Every Move Decreases the Ordinal

Now we must show that every move Hercules makes strictly decreases the ordinal. There are two cases.

Case 1: Hercules cuts a head attached directly to the root. This head contributes a term ω^0 = 1 to the root's ordinal. Removing the head removes this term, and nothing grows back. The ordinal strictly decreases. Easy!

Case 2: Hercules cuts a head not attached to the root. Let's say the head's parent is a node with ordinal $\alpha + 1$. This node contributes a term ω^(α+1) to its own parent's ordinal. When Hercules cuts the head, the parent's ordinal decreases from $\alpha + 1$ to α, and the Hydra grows N new copies of the subtree rooted at the parent.

So the term ω^(α+1) in the grandparent's ordinal is replaced by *(N + 1) · ω^α*. Is this smaller? Yes — and here is why this is so remarkable. No matter how large the finite number N is:

$$\omega\text{\^{}}(\alpha+1) > N \cdot \omega\text{\^{}}\alpha \quad \textit{for any finite } N.$$

This is the fundamental magic of ordinal arithmetic. The infinite absorbs the finite. Multiplying ω^α by any finite number whatsoever is still strictly less than ω^(α+1). In ordinary arithmetic, swapping $10^2 = 100$ for $99 \times 10 = 990$ would be an increase. But in the ordinal world, ω^2 is infinitely larger than $N \cdot \omega$ for any finite N. The Hydra's frantic regeneration, however wild, cannot compensate for the loss of even one level of depth.

Since every move decreases the ordinal, and since there is no infinite decreasing sequence of ordinals, the game must end after finitely many moves. Hercules wins — always, no matter what strategy he uses, and no matter how aggressively the Hydra regrows.

7 The Twist: A Proof Beyond Arithmetic

You might wonder: if the proof is so clear, why is this result particularly interesting? The answer lies in the logical foundations of mathematics itself.

Our proof above used ordinal numbers, which live in the realm of infinite set theory — a system much more powerful than ordinary arithmetic. This is surprising, because the Hydra game is completely finite: every Hydra has finitely many heads, and the game ends in a finite number of steps. Surely, you might think, we could prove a statement about a finite game using only finite methods?

Kirby and Paris showed that we cannot. Their landmark result states that any proof that Hercules always wins requires techniques strong enough to prove the consistency of Peano arithmetic — the standard axiomatic system for the natural numbers. By Gödel's Second Incompleteness Theorem, Peano arithmetic cannot prove its own consistency from within. Therefore, the statement *'Hercules always wins the Hydra game'* is **independent of Peano arithmetic**: it can be neither proved nor disproved using only the basic axioms of natural number arithmetic.

This puts the Hydra game in extraordinary company, alongside Gödel's own incompleteness statements and Goodstein's Theorem — another result by Kirby and Paris in the same paper — as one of the simplest, most natural mathematical statements that nevertheless lies beyond the reach of everyday arithmetic. The Hydra is not just a game; it is a window into the limits of mathematical proof.

8 How Long Does Hercules Take?

If Hercules always wins, one might ask: how many moves does it take? The answer is staggering. Even for modest-looking Hydras, the number of steps required grows faster than any function you can define within Peano arithmetic.

Consider the simplest non-trivial case: a Hydra that is just a chain of two nodes above the root, with a single head at the top. Define *Hydra(k)* as the number of steps needed to defeat a Hydra that is a path of length *k* using the standard strategy of always cutting the rightmost head with *N* copies at step *N*. Then *Hydra(1)* = 1, *Hydra(2)* is already in the billions, and *Hydra(4)* is a number that dwarfs Graham's number — itself once described as the largest number ever used in a serious mathematical proof.

Hercules wins, but he had better start early.

9 Conclusion

What I find most beautiful about the Hydra game is the contrast it embodies. On one side: a simple, playful game that a child could understand, named after a monster from ancient myth. On the other: a result sitting at the very edge of formal mathematics, telling us something profound about the limits of proof itself.

The key mathematical ingredient — that every move decreases the ordinal, and that ordinals are well-ordered — is elegant and not too hard to grasp. Yet this innocent observation implies something that cannot be captured within the usual axioms of arithmetic. It is as if the Hydra, in losing, has the last laugh: to prove that Hercules always wins, you must step outside the finite world entirely.

This is what mathematics does at its best. A question that begins with a game and a myth leads us, through careful reasoning, to the frontiers of logic and infinity. I hope this journey has been as exhilarating for you as it was for me — and I hope that if you ever find yourself facing a Hydra, you will remember: no matter how many heads it grows, you cannot lose.